The Everyday AI Collection

150 step-by-step guides for real, everyday tasks

Staying safe · Money & admin · Home & family · Photos, video & music · Learning & getting started · Work & email · Understand anything · Automate the boring bits

AI World HQ · https://aiworldhq.com

Free edition · CC BY-NC 4.0 · free to share with credit

Before you start — what this is, honestly

This collection is free, and every guide in it is also free on our website. This file just gathers them in one place, in reading order, so you can print it, keep it on your phone, or read it offline. Our website has no paywall either — every guide links back to its online page, which is always the newer version.

AI **Every guide in this pack was written by AI** and checked by automated fact and quality gates. **Our editor spot-checks published guides every day, but not every text is read by a person.** Every guide page on https://aiworldhq.com carries the same notice — we say it openly because you deserve to know who wrote what you are reading.

Apps change. Buttons move and menus get renamed. If a step does not match what you see, the guide's online version is the newer one — the link is under every title.

Free to share. This collection is licensed under Creative Commons Attribution-NonCommercial 4.0 (CC BY-NC 4.0): you may copy, print and share it — with your class, your library or your team — as long as you credit AI World HQ with a link to aiworldhq.com and do not sell it.

What's inside

1. Staying safe

1. How to check if a suspicious message is a scam by pasting it into an AI
2. How to Spot an AI Scam Call, Fake Email, or Deepfake Voice Today
3. Turn Off AI Chat History and Data Saving in ChatGPT, Claude, and Gemini
4. What You Should Never Share With an AI Assistant: A Privacy Guide
5. How to use AI safely with your personal information
6. How to Spot a Phishing Email or Scam Text Before You Click
7. How to Spot AI Scams, Deepfakes, and Fake Content
8. How to Spot AI-Generated Scam Messages Before You Click or Reply
9. How to Spot a Deepfake Video or Voice Clone Before You Share It

2. Money & admin

1. Build a Simple Personal Budget Spreadsheet Using Any AI Assistant
2. Decode Any Confusing Phone, Internet, or Medical Bill in 10 Minutes
3. How to read and summarize a long lease or insurance policy with Gemini
4. How to Use ChatGPT to Explain a Confusing Bill or Invoice
5. How to Decode a Confusing Medical Bill with AI
6. How to Read Through Insurance Policies and Bank Contracts with Claude

3. Home & family

1. Build a Smart Packing List for Any Trip in Just a Few Minutes
2. Plan a Week's Meals from What's Already in Your Fridge
3. Plan a Family Weekend Getaway with DeepSeek: A Step-by-Step Guide
4. Create a soothing bedtime story routine with Alexa
5. How to Plan a Kids' Birthday Party Without Losing Your Mind
6. How to use Qwen Chat to plan a weekend road trip in the US without getting lost
7. Turn messy phone photos into a family album with Google Photos AI
8. Organize your favorite recipes with Perplexity Q&A
9. Create a personalized weekly meal plan based on your diet with Gemini
10. How to write heartfelt birthday cards, wedding toasts and thank-you notes with Copilot
11. Plan a 3-day weekend road trip with Meta AI in 15 minutes
12. Translate a foreign-language recipe or instruction manual with Hugging Face's free translator
13. Build a Simple Weekly Meal Plan With Grok Based on What's in Your Fridge
14. Plan a Budget-Friendly City Vacation in One Afternoon
15. Plan Your Weekly Meals and Shopping List with Any AI Assistant
16. Plan a Weekend Trip From a One-Line Idea With Cohere
17. Use ChatGPT to Plan a Road Trip with Stops, Fuel Breaks, and Drive Time
18. Creating a Travel Packing List for Any Trip with AI Guidance
19. How to scale a family recipe up or down for any number of guests
20. Plan a Relaxing Weekend Getaway From Scratch in One Sitting
21. Brainstorm birthday gift ideas with Le Chat for any budget and personality
22. How to create a custom bedtime story starring your child with any AI
23. How to organize a chaotic shared family calendar in one afternoon
24. Drafting a Basic Family-Tree Outline with AI: A Beginner's Guide
25. How to Troubleshoot Your Home Wi-Fi Step by Step with an AI Assistant

4. Photos, video & music

1. Generate unique profile pictures for social media with Midjourney
2. How to generate AI images in Picsart: a beginner's step-by-step guide
3. Auto-write captions for your photos using Hugging Face's free image captioning tool
4. How to remove a photo background cleanly with Meta AI on your phone
5. Create a multi-room music playlist with Alexa+
6. Sharpen Your Photos with AI Using Upscayl
7. How to create and edit images in ChatGPT
8. How to Sort and Clean Up a Messy Phone Photo Library with Gemini
9. Safely Rename a Batch of Photos with a GitHub Copilot Script
10. How to Turn a Long Podcast Transcript Into Shareable Notes With Claude

11. How to Summarize a Long YouTube Video in 60 Seconds with Any AI
12. Sharpen Old Videos with NVIDIA AI Super-Resolution Online
13. Create Royalty-Free AI Music Loops for Podcasts and Videos
14. Remove Distractions from Photos with Google Photos Magic Eraser
15. How to use Kling AI video templates in Picsart Flow
16. How to create images with AI: a beginner's first try
17. How to Edit Your Phone Photos with Simple Text Instructions in Meta AI
18. How to Generate Images with Google Gemini: A Beginner's Guide
19. Create Your First AI Image with Midjourney: A Beginner Walkthrough
20. How to Create AI Logos with Recraft V4.1 Vector
21. How to Create Custom Images with Copilot for Invites and Social Posts
22. Use Gemini on Your Phone to Pull Editable Text from Any Photo or Screenshot
23. How to Turn Your Saved Articles and Notes into a Personal Podcast with NotebookLM
24. How to Design a Simple Logo for Your Side Project with Midjourney

5. Learning & getting started

1. Practice a new language or presentation with Claude for instant feedback
2. Build a personal study-buddy chatbot in 15 minutes using Hugging Face Spaces (no code)
3. Build a simple weekly study plan with Meta AI in 10 minutes
4. Getting Started with Mistral's Le Chat: Your Private AI Guide
5. How to Translate and Understand Other Languages with AI
6. Use ChatGPT to Practice a Foreign Language Without Feeling Awkward
7. How to Build a 30-Day Plan to Learn Almost Any New Skill from Scratch
8. How to turn your messy study notes into practice quiz questions
9. Boost Your Learning: Get Quick Explanations and Study Guides with Copilot
10. How to use Qwen Chat to create personalized study flashcards for any subject
11. Turn a Reading into a Study Quiz with GitHub Copilot
12. How to Create a Simple Weekly Study Plan with Any AI Assistant
13. Turn Your Messy Class Notes into Study Flashcards with Cohere
14. AI Made Simple: A Gentle First Guide for Beginners
15. How to pick the right free AI app when you're just getting started
16. How to study and learn faster with AI – a beginner’s step-by-step guide
17. Using Perplexity AI for Quick Study Summaries and Exam Prep
18. Use Alexa+ to Pick Up Common Phrases in Another Language
19. The main AI assistants explained simply (and what each is best at)
20. How to use Meta AI inside WhatsApp – step-by-step for beginners
21. A Beginner's Guide to Using ChatGPT Voice for Hands-Free Tasks
22. Getting started with DeepSeek for everyday help
23. Using Claude to Understand Complex Topics for Study or Work
24. How to Use ChatGPT Voice Mode to Practice Conversations in a New Language
25. Organize your work with Claude Projects – a step-by-step beginner’s guide

6. Work & email

1. Automate Email Drafts with AI for Faster Inbox Management
2. How to Get Quick Summaries from Long Documents with Claude
3. Turn Messy Notes into a Polite Meeting Agenda in Five Minutes
4. Build a simple review form for your small business with GitHub Copilot
5. How to Clean Up a Flooded Email Inbox Using AI in One Afternoon
6. Check your resume for ATS-friendly keywords using Hugging Face’s Inference API
7. Write a Polite Complaint Email to a Company with DeepSeek
8. Drafting Professional Presentations from Scratch Using Microsoft Copilot in PowerPoint
9. Digitize Handwritten Notes into Editable Text with AI
10. How to turn a long PDF contract into a plain-English summary in 5 minutes
11. Unlock Document Insights with Google NotebookLM
12. Summarize Any Long Document or Article with AI
13. How to Prepare for a Job Interview Using a Pasted Job Post
14. Write a Resume From Scratch (Even With No Work Experience)
15. How to Organize a Messy Email Inbox Into Simple Folders in One Sitting
16. Practice job interview questions with Qwen before the real interview
17. How to write a polite but firm complaint email that actually gets a reply

18. Translate and Explain a Confusing Email or Letter with Perplexity
19. How to Use Gemini Deep Research for Thorough Reports
20. How to Draft a Polite Follow-up Email After a Job Interview
21. How to practice for a job interview with an AI mock interviewer tonight
22. Polish Your Resume Bullet Points with Le Chat in 10 Minutes
23. Pull Names, Emails, and Phone Numbers Out of Messy Text with Cohere
24. Use Alexa+ as a Hands-Free Meeting Note-Taker and Task Builder at Work
25. How to build or improve your CV with AI

7. Understand anything

1. How to Ask GitHub Copilot to Explain a Confusing Error Message
2. Compare two job offers side-by-side with Le Chat in plain English
3. Quickly Research Any Topic and Summarize Web Pages with Microsoft Copilot in Edge
4. Fact-Check Online Claims with Perplexity's Sources
5. How to Summarize a Long Article or PDF into Key Bullet Points
6. Translate a product description into Spanish and French with Cohere
7. How to compare two rental listings side by side with weighted pros and cons
8. How to use Gemini Deep Research to plan a big purchase decision
9. Catch Up Fast: Summarize Long Chats in WhatsApp and Messenger with Meta AI
10. Summarize a long PDF into 5 bullet points with Claude
11. Summarize Long Articles into Key Bullet Points with Grok on X
12. Translate and Explain Any Restaurant Menu with Qwen on Your Phone
13. Quickly gauge public mood with Hugging Face Sentiment Analysis
14. Unpack Complex Ideas with DeepSeek's Simple Explanations
15. How to Ask Gemini to Explain Your Bank Statement in Plain English
16. How to Research a Medical Procedure or New Medication with Perplexity
17. Research and Draft Short Social Media Captions with Perplexity
18. Compare Products with Perplexity Before You Buy
19. How to Compare Phone Plans and Pick the Right One Without the Jargon
20. Use Copilot to explain code you found online in plain English
21. Summarize a long chat thread into 3 bullet points with Copilot
22. Read the Room on X: Summarize Public Opinion with Grok
23. Ask Copilot to compare two product specs and list the differences
24. Summarize Long Research PDFs with Cohere's Clever AI

8. Automate the boring bits

1. Create a tiny AI weather-helper widget for your blog with Hugging Face widgets
2. How to Build a Simple Weekly Schedule in Google Sheets with Gemini
3. Build a Simple Q&A Chatbot for Your Website with Cohere
4. Set up named medication reminders with Alexa+ for yourself or an older relative
5. Build a Basic AI Chatbot Brain Using OpenAI's API – Quick Start
6. Create a personalized workout routine with DeepSeek's multilingual chat
7. Automate Git Commands with GitHub Copilot Suggestions
8. How to Add an AI Step to Your Zapier Workflow
9. How to Plan a Simple Weekly Workout Routine That Fits Your Real Life
10. Organize Your Daily Schedule by Voice with Alexa+
11. Set Up an Alexa+ Routine to Automate Your Everyday Tasks
12. How to Add an Alibaba Cloud AI Chatbot to Your WordPress Site

9. Every prompt and example in this pack

Staying safe

How to check if a suspicious message is a scam by pasting it into an AI

6 steps · Read online

> **In short:** Paste a sketchy message into any AI assistant — ChatGPT, Gemini, Claude, Copilot, or Le Chat — and ask it to flag the red flags. You'll get a second opinion in under a minute, but never paste real passwords, and always verify with the company first.

Scam messages have gotten good. They use real company logos, your real name, sometimes even your real address. Before you click anything, copy the message into an AI assistant and ask for a second pair of eyes. This guide walks you through that in six steps — even if you've never used an AI assistant before. One honest warning up front: the AI is a useful first check, not a verdict, and it cannot actually click the link or call the company for you. You still have to do that part.

Before you start

- An AI assistant you can sign into — ChatGPT, Gemini, Claude, Microsoft Copilot, and Le Chat all work, and their free plans are fine for this.
- The suspicious message itself, ready to copy (an email, a text, a chat screenshot, or a DM).
- About 5 minutes in total.
- Important: never paste real passwords, full card numbers, or government ID numbers, even if the message contains them — Step 2 shows you how to handle that.

Step 1 — Open a fresh chat in any AI assistant

Open the website or app of any AI assistant — ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or whichever you already use. Start a brand-new conversation rather than dropping the message into an older thread. On a phone, the app icon usually looks like a small chat bubble or a starry logo. On a computer, just go to the assistant's homepage and click into the empty chat box. A clean thread means the AI won't get distracted by older context and will focus entirely on your message. If you don't already have an account, the free sign-up usually takes under a minute and only needs an email address. 💬 Example: open ChatGPT, Gemini, Claude, Microsoft Copilot, or Le Chat — whichever one you already have an account on. You'll know it worked when you see an empty chat box with a placeholder like "Message…" or "Ask anything" and a send button (often an arrow or paper-plane icon).

Step 2 — Redact personal details before you paste

Before you copy the message, take 30 seconds to hide any sensitive bits it contains. Cover real passwords, full credit card numbers, your Social Security or national ID number, your home address, and any one-time login codes. A quick search-and-replace in any notes app works fine. Think of it like covering your hand on a debit-card PIN pad at the store — the AI doesn't need those details to spot a scam, and fewer real numbers in the chat mean less risk if your account is ever hacked. If the message has no personal details at all, you can skip ahead to Step 3. 💬 Example: turn "my SSN is 123-45-6789, please call 555-0123" into "my SSN is [REDACTED], please call [REDACTED]". You'll know it worked when the message still reads naturally but no real numbers, codes, or addresses remain.

Step 3 — Paste the message and ask the AI to flag red flags

Tap or click inside the empty chat box and paste the redacted message. Then add a short prompt asking the assistant to look at it as if it were a security expert. Press the send button (usually an arrow ↗ or paper-plane icon, but apps vary). After a few seconds the AI will start typing a reply — usually a short list of warning signs plus a one-line verdict like "this looks like a phishing attempt" (a fake message designed to steal your login or money) or "this appears legitimate, but verify the sender". If nothing appears after a few seconds, look for a small "stop" or square icon — the response sometimes needs a moment to begin. Apps change often, so if your screen looks a little different, just look for any button that submits your message. 💬 Example: "Here's a message I just received. Can you act as a security expert and tell me whether it looks like a scam? List the specific red flags you see, and tell me whether I should click any links in it. Message: [paste here]" You'll know it worked when you see a written response that names specific warning signs, such as urgency, suspicious links, an odd sender, or requests for personal info.

Step 4 — Read the AI's answer with a critical eye

Treat the AI's reply like a friend's opinion, not a doctor's diagnosis. Read each red flag it names and ask yourself: does this actually match the message I received? Many scams lean on urgency ("act in the next 24 hours"), impersonate a brand you trust, include a link that almost looks right, ask you to click or log in, or demand payment in gift cards or cryptocurrency. If the AI flags something that isn't actually in your message, ignore that point — assistants sometimes guess wrong, and you are still the expert on your own inbox. Conversely, if something feels off to you but the AI didn't mention it, trust your gut and move to the next step. Be especially wary if the AI says "this looks fine" but the message asks you to click a link, send money, or share a code — that's the moment to get a second human opinion. 💬 Example: if the AI says "the sender domain doesn't match the company", check that yourself — does the "from" address really end in the company's real website, or in something like "support-service.example.com"? You'll know it worked when you can repeat back, in your own words, two or three specific reasons the AI gave for its verdict.

Step 5 — Verify any links or sender names yourself

Don't click anything yet. If the message contains a link, hover over it on a computer or long-press it on a phone — you'll see the real destination. Real companies use their own short domain (the brand's actual homepage), while scam links often hide behind words like "secure-login" or "verify-account" on a totally different address. The simple rule: if the link doesn't end on the brand's real site, it's almost certainly fake. If the message claims to be from a person you know, contact that person through a channel you already trust — a phone call or a separate chat — and never reply to the suspicious message itself. If it claims to be from a company like your bank, look up the official number on their real website (type the address yourself, don't use any number inside the message) and ask them directly. 💬 Example: a "USPS" tracking link points to `usps-tracking.example.com/pay-now` — that is NOT usps.com, so it's a fake. You'll know it worked when you have either confirmed the message is genuine through a trusted channel, or you've concluded it's a scam.

Step 6 — Decide what to do, and don't act in a rush

Now choose one of three actions. If it's a scam: delete it, block the sender, and — if it pretended to be a company — report it to that company through their real website. If it's genuine but unexpected: still don't click links inside it; go to the company's real site yourself and log in there. If you're not sure: leave it alone for an hour, then re-read it cold; urgency is the scammer's favorite tool, and a slow second look usually reveals more. Tell a friend or family member if you're worried — a second human pair of eyes still catches things AI misses, especially with voice or video scams. The whole point of using an AI here is to slow down, not to rush. 💬 Example: text a family member — "Hey, I just got a weird message about my Amazon account, does this sound legit to you? [paste]" You'll know it worked when you take one concrete action in the next few minutes — delete, report, or verify through a trusted channel — instead of leaving the message sitting unread in your inbox.

Common mistakes

- **Pasting real passwords or full card numbers.** The AI doesn't need them, and storing them in a chat creates risk if your account is ever hacked. Fix: redact first (Step 2), or describe the message in your own words instead of pasting it.
- **Trusting the AI's verdict without checking.** AI assistants sometimes miss obvious scams or flag clean messages by mistake. Fix: always cross-check the specific red flags it names against the actual message (Step 4) and verify with the real company (Step 5).
- **Clicking the suspicious link "just to see what happens".** That's exactly what the scammer wants. Fix: hover or long-press to read the real URL, and never log in through a link inside an unexpected message — go to the company's website yourself.

What this means for you

- **In everyday life:** the next time a "your package couldn't be delivered" text or a "your account will be locked" email lands, you can get a clear second opinion in under a minute instead of guessing whether to tap it.
- **For your business or work:** if you handle invoices or vendor payments, this same 5-minute check on unusual payment-change requests can stop a business email compromise (a scam where a criminal impersonates a vendor or boss to redirect a payment) — one of the most expensive scams going.
- **If you're just getting started:** pick the AI assistant you already use, bookmark this guide, and try it the next time something feels even slightly off. Practice on a junk email first if you want to build confidence.

Try it now

Find one message in your inbox you've been unsure about, redact any personal details, paste it into your AI assistant, and ask it to flag the red flags. Five minutes, one fresh chat, one safer decision.

How to Spot an AI Scam Call, Fake Email, or Deepfake Voice Today

6 steps · Read online

> **In short:** AI scams now sound like your family, write like a real person, and spoof your bank's phone number. A five-second pause plus one private family word catches almost all of them — here's how to set both up today.

By the end of this guide you'll know how to recognize the three most common AI scams circulating right now, what one-second check stops each one, and what to do the moment you realize something is wrong. This is written for anyone who uses a phone and email — no tech background needed, and nothing here requires buying or installing anything.

Before you start

- A phone (mobile or landline) and an email account you check regularly.
- About 10 minutes total, mostly to set up one protection you can use for years.
- One or two trusted people — a family member or close friend — for the family safe-word step.

Step 1 — Listen for the giveaway in a "family emergency" voice call

When someone calls claiming to be a relative in trouble, the very first move is to **hang up** and dial that person's number yourself — the one already saved in your phone. What you'll usually see on screen is a familiar name on the caller ID, sometimes with a stranger's voice asking "guess who this is?" to make you fill in the blank yourself. After you hang up and redial, a real relative will answer or call you back; a scam line often disconnects, plays a generic voicemail, or doesn't ring at all. If your caller ID shows nothing or a random number, the same rule applies — never continue the call. Voice cloning (an AI that copies someone's voice from a few seconds of audio) is good enough now that even adult children of the "caller" can't tell the difference in the first few seconds. Example: "Grandma, it's me, I was in an accident, please send money" → hang up → dial your grandchild's saved number. You'll know it worked when the real person answers your second call and confirms they never phoned you.

Step 2 — Catch a fake customer-service agent

Any time someone calls YOU about a problem with your account — bank, phone, delivery, internet — hang up and call the company back using the number printed on your last statement or shown on the company's official site. What you'll usually see in the original call is an "agent" who already knows your name, the last four digits of your card, or recent transactions; the voice sounds professional and there's often office noise in the background to feel legit. A real company will let you hang up and call back; a scammer will push you to "stay on the line" or refuse to give you a number you can verify. If the caller insists they can't take a callback because the case is "time-sensitive," that's actually the biggest giveaway. Example: "Hi, this is Maria from your bank's fraud department" → hang up → look up the bank's number on your statement → dial it yourself. You'll know it worked when the real company's recorded greeting answers and the human you reach has no record of the previous call.

Step 3 — Spot the AI-written phishing email

Before clicking any link in a message that asks you to act, **hover over the sender's address and any link** to see where they actually point. What you'll see is an email that looks like it came from a delivery company, a streaming service, or your bank — the wording is fluent, friendly, and free of the typos that old scams used to have; there's often a small urgency like "act within 24 hours" to nudge you past thinking. Hovering reveals a different domain — something like `usps-tracking.example.com/pay-now` instead of a real postal address, or a long random-looking string of letters. On a phone you can long-press a link to preview it; if you can't tell where it goes, don't tap. AI-written phishing is convincing because the grammar is now perfect — the giveaway is no longer the spelling, it's the destination. Example: "Your package is held at customs, pay a small fee to release it" → hover the link → see it points to a domain you've never heard of → delete. You'll know it worked when you can name the actual destination of the link before deciding whether to click.

Step 4 — Use the one-second check that works for all three

When any message or call feels urgent, scary, or too lucky to be true, pause for five seconds and ask yourself one question: **"Would I be in trouble if I did nothing for the next ten minutes?"** What you'll notice in those five seconds is your own hands wanting to act fast — and the caller or email pushing you to "do it now." Most scams collapse under a ten-minute pause: the "agent" hangs up and moves on, the email can sit unread overnight, and your "grandchild" is still safe wherever they actually are. Some scams add a fake countdown timer to fight this exact habit; remember that a real emergency from a real company will still be a real emergency ten minutes from now, and a real loved one will understand a short wait. Example: "Pay this fee now or your account closes in five minutes" → pause → ask the ten-minute question → almost always the answer is no. You'll know it worked when the urgency melts away and you make your next move from calm instead of panic.

Step 5 — Set up a family safe word today

Pick one unusual word with one or two people you trust, and agree that any "emergency" call must include this word before any money moves. A short group chat, a text, or even an in-person conversation is enough — no app or password manager needed. What this gives you is a private test

that even a perfect voice clone can't pass: the scammer simply doesn't know your word, and your real relative does. If your family is scattered, agree on a private question with a private answer instead — "what did we name the cat at grandma's house?" works better than a single word when kids are involved. Write the word down once, share it with the people who need it, and don't post it anywhere online. 💬 Example: Text your sister: "If anyone calls saying they're me and needs money, ask them for our safe word. The word is platypus." You'll know it worked when at least two people in your circle can recite the word without looking it up.

Step 6 — Know what to do in the moment you realize it's a scam

Hang up, close the email, or stop replying — and then **report it**, even if you didn't lose a cent. Your heart will still be beating fast, the "agent" will be gone, and the email will still be sitting in your inbox; that's normal. Reporting helps block the next person, and in the US you can forward phishing emails to reportphishing@apwg.org and file a report at the FTC's reportfraud.ftc.gov, or with the FBI's Internet Crime Complaint Center at ic3.gov. If you DID send money or share a one-time code, call your bank's fraud line immediately — the number on the back of your card — and ask for a recall; speed matters more than embarrassment here. 💬 Example: Realized mid-call that the "bank agent" was fake → hung up → called the number on my debit card → told them what happened → account flagged for monitoring. You'll know it worked when the call is over, you've told one official channel, and your accounts and passwords are unchanged.

Common mistakes

- **Trusting caller ID.** Scammers can spoof any name or number. The display showing "Daughter" or "Your Bank" proves nothing — hang up and dial back on a number you saved yourself.
- **Acting on the "ten-minute countdown."** That timer is part of the scam. Real companies and real family emergencies don't disappear while you take a breath; pausing is always safe.
- **Sharing a one-time code "to verify yourself."** No legitimate company will ever ask for a code they just sent you. The instant someone asks for that string of numbers, the conversation is a scam — full stop.
- **Skipping the family safe word because "it'll never happen to us."** The setup takes one minute today and saves potentially life-changing money later. Treat it like a smoke detector: boring to install, priceless when it goes off.

What this means for you

- **In everyday life:** the next time your phone rings with a panicked voice, you'll know to hang up first and call back on a number you trust — instead of sending money in the first thirty seconds.
- **For your business or work:** any team that handles invoices, supplier emails, or wire transfers should adopt a private code word the same way families do; deepfake voice scams have already hit finance and HR departments, not just grandmothers.
- **If you're just getting started:** do just one thing today — text one person in your family the safe word from Step 5. That single message is more protection than any antivirus app on your phone.

Try it now

Right now, open your messaging app and text one trusted person: *"If I ever call you in a panic and ask for money or codes, ask me for our safe word. The word is [pick something unusual, like 'platypus' or 'silver-spoon']."* That one text takes under two minutes and is the single highest-value defense in this whole guide.

Turn Off AI Chat History and Data Saving in ChatGPT, Claude, and Gemini

6 steps · Read online

> **In short:** You can stop ChatGPT, Claude, and Gemini from using your conversations to train future AI models by flipping a privacy toggle in each service's settings — but the exact menu names change often, and turning off training does not erase the conversations those companies have already saved.

By default, ChatGPT, Claude, and Gemini save your chats and may use them to train future versions of themselves. This guide shows you how to switch that off in each service — on web and mobile — and explains what each toggle actually changes, and what it quietly does not.

Before you start

- A free or paid **personal** account on each AI service you want to adjust. Work or school accounts are usually managed by your employer and may not let you change these settings.
- A web browser is the easiest place to do this. Phone apps work too, but menus shift around more often.
- Set aside about **15 minutes total** — roughly five minutes per service.
- Make sure you are signed in to each service before you begin, so the settings actually save.

Step 1 — Open ChatGPT's Data Controls

Go to **chatgpt.com** in a browser and sign in if you are not already. Click your **profile picture or initials** in the top-right corner — a small menu appears. Look for an option like **Settings** (sometimes shown as a gear icon ⚙) and click it. In the Settings panel that opens on the left side of the page, look for a section called **Data Controls** and click that. If your screen looks different, try a section called **Privacy** instead — older versions used that name.

If you are on a phone, tap the **two lines or three dots** in the top corner of the ChatGPT app, then tap your name at the top of the menu that slides up. The same Settings → Data Controls path should appear.

You'll know it worked when you see a list of toggles — look for ones mentioning "Improve", "training", or "history".

Step 2 — Turn off ChatGPT's training toggle

On the Data Controls page, find the toggle labeled something like **"Improve model for everyone"** or **"Chat history & training"** — the exact label changes often, so read each toggle's description. Click the toggle so it slides to the **off** position (it usually turns gray or loses its color). You may see a small popup asking you to confirm — click **Confirm** or **Save**. The page may refresh, or you may need to close and reopen it for the change to show.

If you do not see a training toggle at all, your account may be on a plan or in a region where the option is hidden. In that case, leave the page and check back in a few weeks — settings are added over time.

You'll know it worked when the toggle stays gray (off) after a page refresh and no longer shows an "on" indicator. Even after turning this off, OpenAI may still keep your most recent chats for up to 30 days for safety monitoring — so avoid pasting anything truly sensitive in the meantime.

Step 3 — Open Claude's privacy settings

Go to **claude.ai** and sign in. Click your **profile picture or initials** — usually the bottom-left on web or the top-right on mobile. From the menu that pops up, choose **Settings**. Inside Settings, look for a tab or section called **Privacy**, **Account**, or **Profile** — the exact name varies. Within that section, hunt for a toggle that mentions **"help improve Claude"** or **"use my chats for training"**.

On the Claude mobile app, the same path usually works: tap your **profile icon**, then **Settings**, then look for the privacy section. If your app opens straight into a chat, look for a **gear icon** ⚙ in the top corner instead.

You'll know it worked when you see a privacy section with at least one on/off toggle related to training or improving the assistant.

Step 4 — Turn off Claude's training toggle

Click the toggle next to **"Help improve Claude"** (or whichever similar phrase appears on your screen) so it slides to the **off** position. Some users see a confirmation popup — accept it. The change usually saves instantly without a separate Save button, but if you see one, click it.

A few honest notes here. On some free Claude accounts, the option to opt out of training may be limited or unavailable — Anthropic has updated this policy more than once. If you do not see the toggle at all, check whether you are signed in with a **work or school email**; those accounts are managed by an admin and cannot be changed here. If you pay for Claude, you may see the toggle in a slightly different place — the same Settings area, but possibly under a **Privacy** sub-tab.

You'll know it worked when the toggle is gray and stays gray after you refresh the page. If the setting does not seem to stick, sign out, sign back in, and check again — sometimes it takes a moment to sync.

Step 5 — Open Gemini's activity controls

Gemini handles this a little differently because it lives inside your Google account. Go to **gemini.google.com** or open the Gemini app. Click your **profile picture** in the top corner. Look for **Settings** or **Activity** — on the web, you may see a direct link labeled **"Gemini Apps Activity"** that takes you straight into your Google account controls. Click it.

On mobile, tap your **profile picture** in the top corner of the Gemini app, then look for **Gemini Apps Activity** in the menu that appears. If it is not there, your app version may have moved it — try searching your phone's settings app for **"Gemini"** instead, or open **myactivity.google.com** in a browser.

You'll know it worked when a Google page opens showing a toggle labeled **"Gemini Apps Activity"** and a short description of what it does.

Step 6 — Turn off Gemini's Apps Activity

On the Gemini Apps Activity page, click the toggle at the top so it switches to **off**. A small confirmation panel may pop up explaining what will change — read it, then confirm. You may also see a separate option to **delete past activity**; if you want a clean slate, click that too and follow the prompts.

This single Google toggle covers the Gemini web app, the Gemini mobile app, and any other place Gemini saves your chats to your Google account. Disabling it does not delete your past conversations automatically — you need to do that separately if you want them gone. And remember: turning off Gemini Apps Activity does not affect your Google Search history or YouTube history, which are controlled by separate toggles.

You'll know it worked when the toggle is off and a line of text confirms activity is paused. You can verify by starting a new chat, then returning to this page — the activity should not appear there.

Common mistakes

- **Looking for a "delete my data" button.** None of these toggles erase what the company already has on file. If you want old chats removed, each service has a separate **Delete** or **Clear conversations** option — and even then, companies often keep copies for legal or safety reasons for a while.
- **Assuming one toggle fixes everything.** In Gemini especially, there are **multiple** Google activity controls. Turning off Gemini Apps Activity does not affect Search history, YouTube history, or other Google products. Check each one separately if you care about all of them.
- **Trusting the toggle forever.** Companies update their settings and policies regularly. The opt-out you flipped today may quietly disappear or change meaning in six months — set a calendar reminder every few months to come back and check.

What this means for you

- **In everyday life:** Turning these off is a good habit if you ever paste anything personal into an AI chat — names, addresses, health questions, work emails. It does not make you invisible, but it stops your words from being used to teach the next model.
- **For your business or work:** If you handle customer data, contracts, or anything confidential, do not rely on these toggles alone. Many companies explicitly say opt-outs do not cover legal requests or internal abuse detection. Use a separate AI account with synthetic data, or stick to enterprise plans that promise stronger privacy.
- **If you're just getting started:** Pick the service you use the most, open its settings, and turn off the training toggle right now. That single move is the most useful thing you can do today — you can come back for the other two later.

Try it now

Open the AI assistant you used most recently, click your profile picture, find the privacy or data settings, and turn off the training toggle. It takes about a minute, and you'll see exactly what your settings look like today.

What You Should Never Share With an AI Assistant: A Privacy Guide

6 steps · Read online

By the end of this guide you'll know the three buckets of personal data that belong strictly offline, how to spot any of those details in your chat history, and what to do if you've already let an AI see them. It's written for anyone who chats with tools like ChatGPT, Gemini, Claude or Copilot.

Before you start

- A device with internet access (phone, tablet or computer).
- An account on at least one AI assistant (e.g., ChatGPT, Gemini, Claude – the free tier is enough).
- About 5 minutes of uninterrupted time.

Step 1 — Spot the three privacy buckets

Open the AI assistant you use most and start a fresh conversation. On screen you'll see a blank text box waiting for your prompt, and a "Send" button on the right. When you type, the assistant will start generating a reply beneath your prompt. If the layout looks different, look for a field that says "Ask anything" or a similar text-input area. 💬 Example: "What types of personal information should I never share with you?" You'll know it worked when the assistant lists three groups such as (1) identifiers (name, address, ID numbers), (2) financial details (bank accounts, credit-card numbers), and (3) health or biometric data.

Step 2 — Scan your recent chats for those buckets

Go to the chat-history screen – usually accessed by tapping a clock-icon or a menu labeled "History". You'll see a list of past conversations, each summarized by its first line. Selecting a conversation opens the full exchange, where you can scroll through every message you sent. If you don't see a clock-icon, look for a three-line menu and choose "Recent chats" instead. 💬 Example: type "Show me my last five messages" (many assistants support a simple request like this). You'll know it worked when the assistant displays the recent prompts you entered, giving you a chance to review them.

Step 3 — Delete or redact any message that contains private info

In the opened conversation locate the message that holds an identifier, financial detail or health data. Most assistants let you tap and hold the message, revealing options such as "Delete" or "Remove". Choose "Delete" and confirm when a pop-up asks "Are you sure?". If the option isn't obvious, look for a trash-can icon or a "Clear" button at the bottom of the chat. 💬 Example: after selecting the message, type "Delete this message". You'll know it worked when the message disappears from the chat view and the history no longer shows that content.

Step 4 — Turn off chat-history saving (if your plan allows)

Open the settings panel – usually reached by tapping your profile picture or a gear-shaped icon in the top-right corner. The settings screen will list toggles like "Save chat history" or "Conversation storage". Switch the toggle off; a brief note may appear saying the change applies to future chats only. On a free plan some assistants only let you "clear history" rather than disable it permanently, so look for a button that says "Clear all chats". 💬 Example: "Turn off chat history". You'll know it worked when the toggle changes color or moves to the "off" position, and the app confirms "Your chats will not be saved".

Step 5 — Verify that nothing sensitive is being stored

Start a new conversation and type a harmless query such as "What's the capital of the US?". The assistant will reply below the prompt. Then open the history again and check that only this new, non-personal question appears. If older, private messages still show up, repeat Steps 2–4 for any that remain. 💬 Example: "Show my recent chats". You'll know it worked when the list only contains the recent, innocuous question you just asked.

Step 6 — Request a full data-deletion if you think something slipped through

If you suspect you shared something the assistant cannot delete yourself (for example a piece of text buried deep in a long conversation), look for a "Help" or "Support" link in the settings menu. Most providers have a contact form or email address labeled "Data-privacy request". Send a short message stating "Please delete all personal data associated with my account [your email]". On some platforms you may need to verify your identity first. 💬 Example: "Dear Support, please remove any personal identifiers from my account." You'll know it worked when you receive a confirmation email or a system message saying "Your request has been received".

Common mistakes

- **Sharing full addresses or ID numbers:** If you copy-paste your home address, the assistant now knows a location you might not want stored. Fix it by deleting the message and turning off history as shown in Steps 3–4.

- **Posting bank or credit-card numbers:** Even a partial number can be used for fraud. Remove the message immediately and contact your bank to monitor for any unusual activity.
- **Assuming the assistant deletes data for you:** Deleting a message only removes it from your view; the provider may still keep a copy on its servers. Always use the built-in “Clear all chats” option or submit a formal deletion request (Step 6).

What this means for you

- **In everyday life:** You can safely ask an AI to brainstorm a grocery list without worrying that your phone number or home address is being stored forever.
- **For your business or work:** Confidential client numbers and contract figures stay out of the AI’s training data, reducing the risk of accidental leaks.
- **If you’re just getting started:** The easiest first move is to turn off chat-history saving – it’s a single toggle that protects all future conversations.

Try it now

Open the AI assistant you use most, find the settings gear, and toggle **off** the “Save chat history” option. In under two minutes you’ll have taken a solid step toward keeping your private information out of AI-generated archives.

How to use AI safely with your personal information

4 steps · Read online

Want to use smart AI tools without accidentally sharing your private life with the world? This guide will show you exactly how to protect your personal details while still getting the most out of AI. You'll finish knowing how to thoughtfully interact with AI, manage privacy settings, and clean up sensitive information.

Before you start

- **An active account** on any popular AI tool like ChatGPT, Claude, or Gemini. A free plan is usually sufficient for this guide.
- **A computer or smartphone** with internet access. The steps apply whether you're using a web browser or a mobile app.
- **Rough total time:** About 10-15 minutes to read through and apply the steps to your chosen AI tool.

Step 1 — Understand the "postcard rule"

Before typing anything, pause and consider your message. The core idea for using AI safely is simple: treat everything you type into a chat window like writing on a postcard. If you wouldn't want a stranger reading it at the post office, do not send it to an AI. Most AI systems use your conversations for **data training** – this is where the AI company uses past chats to teach their computer models how to write and think better. Think of it like letting a student read your old letters so they can learn to write their own, but on a massive, anonymised scale.

You'll know it worked when you instinctively rephrase potentially sensitive information before hitting send. This habit helps the AI tool to assist you effectively without needing your private data. While AI can draft an email, it cannot magically remove personal information you've already typed.

Example: Instead of typing: *"My name is Jane Smith, I live at 24 Smith Street and need help budgeting my $1,500 weekly salary. My email is jane.smith@example.com."* Rephrase it to: *"Help me budget a weekly salary of $1,500 for a typical household in a major city."*

Step 2 — Organize your settings and opt out of training

Almost all major AI platforms allow you to turn off data training. When you switch this off, the company still processes your requests to give you an answer, but they promise not to use your chat history to teach future versions of their system. This setting usually applies to future chats, not ones you've already had.

To do this, open your AI tool. Look for a gear icon, your profile picture, or a menu with three horizontal lines, usually located in a corner of the screen (often bottom-left or top-right). Tap or click on this to open a menu. From the menu, find and select a button labeled "Settings", "Privacy", or "Data Controls". Within this section, look for a toggle switch or checkbox that controls how your data is used for "model improvement" or "training". Toggle this option *off*. What happens next is a confirmation that your setting has been saved, or the toggle will simply show it's now off.

If it looks different: App layouts vary between AI tools and change over time. If you don't see exact labels, explore sections related to 'Privacy & Safety', 'Data Management', or 'Chat History'. The key is to find an option that controls how your data is used for "model improvement" or "training". You might need to scroll down or look through a few sub-menus.

Example: In ChatGPT, open Settings (tap your profile picture or the gear icon), then look for "Privacy" or "Data Controls" and toggle off any option related to using your conversations for model improvement.

You'll know it worked when the toggle switch clearly shows it is in the "off" position (it might turn gray, or have a circle on the left side).

Step 3 — Sanitise your text before copying and pasting

"Sanitising" means cleaning your text to remove any specific details like names, phone numbers, email addresses, or financial figures *before* you feed them to the AI. This is your strongest defense against data leaks and ensures the AI doesn't learn sensitive facts about you. This action happens in your mind or a separate document, before you ever type into the AI chat.

When you review your text and replace real details with generic placeholders, the AI will still understand the context perfectly and can help you with your task. What happens is your input becomes general, focusing on the problem rather than the specifics of your personal life.

If it looks different: This step is about your pre-AI preparation, so what you see will be your own document, notes app, or even just your thoughts. No AI screen will look different here, as you're doing the sanitisation yourself.

Example: If you want the AI to polish a professional email, change the real names and figures to generic labels like this: Original: *"Hi Sarah, can you please transfer the $450 refund for the project at 12 High St to my Commonwealth account by Friday?"* Sanitised version: *"Hi [Recipient Name], can you please transfer the [Refund Amount] refund for the project at [Project Address] to my [Bank Name] account by [Date]?"*

You'll know it worked when you read through your prepared text and can't find any specific detail that would directly identify you, another person, or a specific location or account number.

Step 4 — Use temporary chats for sensitive tasks

If you need to discuss a particularly sensitive topic — like brainstorming a tricky work situation, or sorting out a personal issue — look for a temporary chat option. Some tools call this "Temporary Chat", "Incognito Mode", or "Ephemeral Chat". It works like write-on, wipe-off ink: once you close the window, the conversation is deleted from your history and generally won't be used for model training, even if your main training setting is on.

To activate this, open your AI tool and look near the top of your chat window or near the "New Chat" button for an option to start a temporary or incognito chat. Tap or click this before you begin your conversation. What happens is a new chat session will open, often with a visual indicator (like a different background color or an "Incognito" label) to show it's temporary. This temporary chat is usually designed to be forgotten once you close it.

If it looks different: Not all AI tools offer a specific "temporary chat" feature, especially on free plans. If you can't find it, don't worry – simply ensure your main data training setting (from Step 2) is switched off, and always sanitise your information (Step 3).

Example: In Claude, look for a "Temporary Chat" button next to "New Chat" in the sidebar. Click it to start a new, unsaved conversation.

You'll know it worked when you see a clear visual cue that your chat is temporary (e.g., a message like "This chat will not be saved," or a different interface color).

Common mistakes

- **Uploading complete PDF documents without checking:** Many people upload entire utility bills, tax documents, or medical records for the AI to summarize. These files often contain private tax file numbers, account passwords, or personal health information.
  - **The fix:** Always open and review documents yourself *before* uploading. If there's sensitive info, either redact it (black it out or delete it) or copy-paste only the specific, non-sensitive parts you need the AI to process.
- **Treating the AI like a diary:** Do not use the chat box to vent about specific people by name, or share highly emotional, identifiable personal secrets. Even if you've opted out of training, the raw content might still be viewable by human reviewers at the AI company for safety checks.
  - **The fix:** Use generic terms like "a friend," "a colleague," or "a family member" instead of names. Focus on the situation or problem rather than the specific individuals involved. If it's truly private, talk to a trusted person instead.
- **Assuming "deleted" means instant removal:** When you delete a chat from your history sidebar, the AI company may still keep a copy on their secure servers for a short period (days or weeks) to check for harmful content, or for legal compliance. It's not always an instant, permanent deletion across all systems.
  - **The fix:** The best way to protect sensitive data is to *never put it into the AI in the first place*. Rely on the postcard rule (Step 1) and sanitising your text (Step 3) to prevent the data from ever being recorded.

What this means for you

- **In everyday life:** You can safely write personal emails, plan holidays, or draft letters to your local council without worrying about your home address leaking, provided you've followed these steps.
- **For your business or work:** You can use AI to summarize meeting notes, brainstorm marketing plans, or draft reports, provided you swap out client names, proprietary business secrets, or identifiable employee information.
- **If you're just getting started:** The easiest first move is to build a "pause habit"—take three seconds to look over your prompt before hitting the send button. Ask yourself: "Would I be okay with this being on a postcard?"

Try it now

Open your favorite AI tool right now. Navigate to the **Settings** menu, find the **Data Controls** or **Privacy** tab, and toggle off the option for **Model Training** or **Improve the model for everyone**. This simple action takes less than a minute and instantly secures your future chats from being used for broad model training.

How to Spot a Phishing Email or Scam Text Before You Click

6 steps · Read online

> **In short:** Copy any suspicious email or scam text into an AI assistant with a short "fraud detector" prompt, and it lists the red flags — fake urgency, off-brand sender addresses, sketchy links — so you can decide whether to delete, block, or report the message before you click anything.

That pang of "wait, something feels off" you get when a strange email or text arrives — the AI can confirm it (or calm you down) in about 30 seconds. This guide shows you exactly what to type, and how to read the answer so you don't miss the warning signs. It works with ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek, or any other assistant you already use. One honest note up front: an AI can spot a lot, but it can't actually click the link for you or verify whether a sender is real — the final "yes" or "no" is still your call.

Before you start

- An account on any AI assistant you already use (free versions work fine — ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or DeepSeek all do this).
- The suspicious message itself, opened on your screen but with no links tapped or replies sent.
- Roughly 2 minutes.

Step 1 — Read the message calmly before doing anything

Open the suspicious email or text, and do exactly one thing: read it top to bottom with your finger off the mouse. Glance for three things — a sender or number you don't recognize, a tone that pushes you to act NOW, and any link you weren't expecting. You'll see the message exactly as you normally would; nothing changes yet. Example: if it's "Your package is delayed, confirm here: bit.ly/x9zq", just read the words. Don't tap the link. If it looks different: in some messaging apps long-pressing a link shows a small URL preview — that's actually useful, but still don't open it. Reading is the goal here. You'll know it worked when you can describe the message in one short sentence without having clicked anything.

Step 2 — Open any AI assistant you already use

On your phone or laptop, open the AI assistant app or website where you already have an account. No special "fraud tool" is needed — every general-purpose assistant can do this job. You'll see an empty chat box, usually labeled "Message", "Ask anything", or similar. If it looks different: every assistant has its own empty state, but the chat box is almost always in the middle of the screen on desktop or at the bottom on mobile. If you're stuck, start a "New chat". You'll know it worked when you see a blinking cursor and a send button, ready for you to type.

Step 3 — Paste in a "fraud detector" prompt first

Before pasting the suspicious message, give the AI a one-line role instruction. Type it into the chat box and hit send. The AI will respond with a short confirmation — usually something like "Sure, paste it in" — that primes it to focus on fraud patterns instead of generic advice. This single line is what turns a normal chat into a fraud check. Example: "Act as a fraud detector. Read the message I'm pasting next and list every red flag you can see, then tell me clearly whether it's safe to reply, click any link, or open any attachment. If you're not sure, say so." If it looks different: some assistants reply with a longer setup paragraph — that's fine; what matters is that it now knows to focus on scams. You'll know it worked when the AI answers something like "Got it — paste the message" or asks a clarifying question.

Step 4 — Paste the full suspicious message, then send

Highlight the entire message — sender name, subject line, body, and the link text — but leave out your own name, address, or any account numbers if they appear, since the AI doesn't need them to spot trouble. Paste it into the chat box and hit send. The AI now has the raw text and starts analyzing it. If it looks different: if the message is unusually long (a fake invoice, say), you can paste a trimmed version — but keep the opening sentence and any links exactly as written, because that's where the red flags usually live. Example of what to paste: "From: Amazon Support <> — Dear customer, your account has been suspended due to unusual activity. Click here to verify: hhtp://amaz0n-secure.example.com/login — please respond within 24 hours." You'll know it worked when the AI starts replying with a bulleted list of specific things it spotted — usually 4 to 8 items.

Step 5 — Read the red flags and trust the obvious ones

The AI's reply is usually a checklist: things like "the sender domain doesn't match Amazon," "the URL uses a zero instead of an 'o'," "the message creates urgency with a 24-hour deadline," or "the greeting is generic (no real name)." Read each one and ask yourself, "Does this match what I'm seeing on my screen?" You'll start to see a pattern you can recognize next time without any help. Example of an AI answer: "Red flags: sender domain is amaz0n-support.example.com, not amazon.com; the link is 'hhtp' (typo) and uses a zero in 'amaz0n'; urgency tactic ('within 24 hours'); generic greeting. Verdict: do NOT click, do NOT reply. Report and delete." If it looks different: the AI might use softer words ("likely a phishing attempt", "this looks suspicious"). That's the same meaning. If it ever says "I'm not certain", take that seriously and lean toward treating

the message as a scam. You'll know it worked when you can name at least three specific reasons the message is suspicious and explain them out loud.

Step 6 — Decide, then report and delete

Once the AI gives its verdict, decide based on the strength of the red flags. If even one major red flag is present — a sketchy sender domain, pressure to act NOW, or a request for passwords or payment — don't click and don't reply. Use the email or messaging app's built-in "Report phishing" or "Report junk" option if you can find it, then delete the message. The dangerous message is gone from your screen and flagged to the platform. 💬 Example of a good follow-up question to the AI: "Can you explain in one sentence what an 'urgency tactic' means, so I know what to look for next time?" If it looks different: Gmail, Outlook, Apple Mail, and most phone carriers all have a Report button, usually under a "⋯" three-dot menu or near "Mark as" / "Move to". If you can't find the report option, deleting is good enough — what matters is that you don't reply. You'll know it worked when the message is no longer in your inbox, you've flagged it to your email or carrier, and you haven't tapped the link.

Common mistakes

- Replying to the scammer just to ask "is this real?" — replying confirms your address is live, and more scam messages follow. Delete and report instead.
- Trusting the AI without reading the explanation — AI can miss things. If the message "feels" off but the AI says it's fine, trust your gut and don't click. The AI is a second opinion, not a final one.
- Pasting personal details along with the message — your name, address, and account numbers aren't needed to spot red flags; trim them out before pasting, because anything you send into a chat is processed by that service.

What this means for you

- **In everyday life:** A 30-second check before clicking unknown links saves you from stolen passwords, drained bank accounts, and long phone calls with your bank. It becomes a habit that protects you on autopilot.
- **For your work:** Make this your default first step before forwarding suspicious work emails to IT. You'll often catch obvious scams that IT would have flagged anyway — which saves everyone time and earns you a reputation as the careful one on the team.
- **If you're just getting started:** Try it on the very next weird message you receive. You don't need to be perfect — you only need to slow down for 30 seconds before you click.

Try it now

The next time an email or text feels off, copy it into your AI assistant with the prompt from Step 3, paste the full message from Step 4, and read the red flags from Step 5 before you tap anything. The whole check takes about half a minute.

How to Spot AI Scams, Deepfakes, and Fake Content

5 steps · Read online

> **In short:** AI-generated content is getting more convincing, but you can still protect yourself by verifying through independent channels and staying calm when messages create urgency.

By the end of this guide, you will know how to identify AI-generated images, synthetic voices, and deceptive messages before they trick you. This practical guide is designed to help anyone navigate the internet safely, without needing a degree in computer science. While these tips are powerful, remember that AI technology is always evolving. No single method is 100% foolproof, so staying updated and critically evaluating content is your best defense. This guide teaches you to be observant, but it cannot make decisions for you about the trustworthiness of a source.

Before you start

- **A device with internet access:** This could be your mobile phone, tablet, or computer.
- **About 10-15 minutes of your time:** Enough to read through the guide and try a few observations.
- **An observant and calm mindset:** AI changes quickly, so ongoing vigilance and a cool head are your best tools against deception.
- **No special apps or accounts needed:** We'll focus entirely on using your natural observation skills and everyday verification methods.

Step 1 — Spot visual "glitches" in photos

AI image generators are computer programs that create pictures based on text descriptions, much like a highly skilled digital painter who has studied millions of real photos. While they're incredibly talented, older versions of these programs often struggle with fine details, and examining photos closely can sometimes reveal inconsistencies.

When you see a photo that feels "off," your first action should be to zoom in and closely examine the specific details. You'll usually see the image take up more of your screen as you zoom. Once magnified, check areas like hands, feet, accessories, clothing, and backgrounds for subtle errors. If your device doesn't have a simple pinch-to-zoom feature, look for a magnifying glass icon or a dedicated zoom option in your photo viewer. You'll know it worked when you can clearly see distorted features, strange shadows, or illogical elements in the image.

💬 Example: You see a photo of a well-known person in an unusual situation. Zoom in on their hands and face. If you notice anything unusual — extra fingers, distorted features, or inconsistent lighting — investigate further before sharing or believing it.

Note: Modern AI generators have improved dramatically, so these artifacts are becoming less common. When in doubt, verify the image through a reverse image search or other means.

Step 2 — Identify "deepfake" video clues

A **deepfake** is a video where a person's face or body has been digitally altered or even completely swapped using smart computer systems. Think of it like a sophisticated digital mask that moves in real time, trying to perfectly mimic a real person. While detection is getting harder as the technology improves, there are still some things to watch for.

Start by playing the suspicious video, then pause it frequently to get a good look at specific frames. As you play the video, you'll see the frames move, and then when you pause, a still image will appear, allowing you to scrutinise the facial areas without the distraction of constant motion. If your video player doesn't have an obvious pause button, look for a large play/pause icon, often in the center or bottom of the screen. You'll know it worked when you spot any inconsistencies in how the face blends with the hair, unusual lighting, or areas that look slightly blurred or distorted.

💬 Example: You receive an urgent video message from your company's director asking for an immediate fund transfer. Pause the video and look closely at their jawline, neck, and hairline. If you notice anything unusual — a fuzzy outline, unusual shadows, or lip movements that don't quite match the audio — treat it as suspicious and verify through another channel before acting.

Note: High-quality deepfakes are increasingly difficult to spot with the naked eye. If a video requests money or sensitive information, always verify through a separate, trusted method.

Step 3 — Listen for voice cloning issues

Voice cloning is technology that takes a short recording of a real person's voice and then uses a computer to make them say entirely new things. Imagine it like a high-tech text-to-speech tool that can perfectly imitate your friend, child, or colleague. Modern voice cloning can be very convincing, so the best protection is not relying on voice alone.

When you receive a suspicious call or audio message, listen intently to the speaker's voice, paying close attention to their tone and speech patterns. You'll hear the words being spoken and can analyze how natural or unnatural they sound. If you are on a call, try to get to a quiet place or use headphones if available to reduce background noise. If anything feels slightly off — an unusual tone, unexpected phrasing, or a request that feels out of character — take that as a signal to verify through other means.

Example: You receive an urgent phone call from a friend asking for financial help. The voice sounds exactly like them, but the request is unusual or the tone feels strange in context. This is a red flag. Hang up immediately and call them back on their usual, trusted number to verify.

Step 4 — Verify suspicious messages using a second channel

Deceptive messages often leverage AI to write highly convincing, personalized emails or text messages, making them a potent form of automated scamming. The golden rule of online safety is simple: never use the contact details provided *within* a suspicious message to verify its authenticity, as these will likely lead you directly to the scammer.

Your action is to deliberately step away from the original message and initiate contact through a different, independently verified method. You'll see the original message on your screen (email, SMS, social media), which might contain urgent requests or "updated" details. By using an independent contact, you avoid giving the scammer direct access. If the request was via email, for instance, don't reply to that email. Instead, find a trusted phone number from a previous bill or the company's official website. You'll know it worked when the person on the other end confirms (or denies) the request, using a contact method you initiated and trust, separate from the potentially fake one.

Example: A supplier sends an email with "updated bank details" for a payment, urging you to use the new account. Do not reply to that email or call the phone number listed in its signature. Instead, find the phone number from your original paper contract or the supplier's official website and call them directly to confirm the change.

Step 5 — Look for generic or overly perfect text

AI-generated text can often be identified not by obvious errors, but by its *perfection* – or lack of human nuance. These **AI-written texts** might be overly formal, incredibly verbose, repetitive in their phrasing, or simply too generic to sound like genuine human communication. They lack the quirks, emotional inflections, and subtle imperfections that make human writing unique.

Your action is to read the message or article carefully, paying close attention to its style, flow, and word choice. As you read, you'll be evaluating the tone and identifying any unusual patterns. This applies whether you're looking at an email, a social media post, an online article, or even a review. If the text reads like it was written by a robot trying to sound human, it's a good clue. You'll know it worked when the writing feels unnaturally smooth, free of human quirks or slang, or uses overly formal language for a casual situation.

Example: You receive an email from a supposed friend, but it uses overly formal language like "Furthermore, I require immediate assistance" rather than their usual casual style, and every sentence is perfectly structured without any contractions or slang. This unnatural perfection suggests AI involvement.

Common mistakes

- **Trusting "official" branding:** Scammers can easily use AI to copy company logos, official-looking fonts, and professional language. The mistake is to assume legitimacy based on appearance. **The fix:** Always focus on the *action* they are asking you to take (like clicking a suspicious link or sending money), rather than how professional or authentic the message looks.
- **Panicking and acting too fast:** AI scams are designed to create urgency or fear, making you feel pressured to act without thinking. **The fix:** If a message demands immediate action, take a deep breath and pause. Time is always on your side when it comes to verifying facts and protecting yourself.
- **Missing the full context:** Sometimes a small snippet of an image, video, or audio can be shared out of its original context, making it seem suspicious when it's actually harmless. **The fix:** Always try to find the original source or the wider story if something looks questionable. Search online for the image or video to see where it first appeared and what the original context was.

What this means for you

- **In everyday life:** You can protect your family from distressing phone scams by realizing that a familiar voice on the phone might not always belong to the person you expect. Using these observation skills helps you verify crucial information before making hasty decisions, keeping your loved ones safe from emotional and financial exploitation.
- **For your business or work:** Safeguarding your company's assets and reputation becomes easier. You'll be better equipped to detect phishing emails, fraudulent invoices, or deepfake video calls that could lead to significant financial loss or data breaches. Protecting your team and clients from sophisticated AI-powered deception is vital for maintaining trust and security.
- **If you're just getting started:** The most important takeaway is to cultivate a healthy dose of skepticism about content you encounter online, especially if it elicits a strong emotional response or demands urgent action. Start by simply pausing and asking yourself, "Does this feel real?" before you click, reply, or share.

Try it now

Next time you're scrolling through social media or checking your emails, pause when you see an image, video, or message that looks a little too good to be true, or surprisingly out of character. Apply one of the quick checks from this guide – perhaps zoom in on a photo, or read a message aloud to detect robotic tones.

Conclusion

Staying informed and vigilant is your strongest defense against the evolving landscape of AI-driven scams and misinformation. By learning to recognize potential red flags in AI-generated content, you empower yourself and those around you to navigate the digital world with greater confidence and security. Always remember to question, verify, and trust your instincts when something feels off.

How to Spot AI-Generated Scam Messages Before You Click or Reply

5 steps · Read online

> **In short:** Modern scams often use AI to write surprisingly convincing messages, but you can paste the message into any AI assistant (ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek — whichever you already use) and ask it to flag the warning signs. This guide walks you through a simple five-step check before you click a link, reply, or send money.

Scammers used to write in clumsy English. Today, AI helps them write polished messages that look exactly like a real text from your bank, USPS, or a family member. The good news: the same AI assistants we use every day are also very good at spotting these tricks when you ask the right questions. In the next five minutes, you'll learn a repeatable check you can run on any suspicious message.

Before you start

- An account with at least one AI assistant you can chat with — ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, DeepSeek, or similar. The free tier is enough for this.
- A phone or computer where you can copy the suspicious message and paste it into the chat.
- The original message itself (screenshot or text), plus the sender's number, email or username if you have it.
- Roughly **5 minutes** for your first check; faster once you've done it two or three times.

> A quick honest note: an AI assistant can flag *likely* warning signs. It cannot guarantee a message is safe, and it cannot detect every new trick. When in doubt, treat the message as suspicious and contact the organization directly using a phone number from their official website.

Step 1 — Don't click anything yet; copy the message safely

The very first action is the most important: **do not tap any link, reply, or call any number inside the suspicious message.** Scam links can take you to convincing fake sites, and replying confirms your number is real. Instead, take a breath, then open the message and copy the full text. On most phones, press and hold the message bubble and look for a **Copy** option in the small menu that pops up. On a computer, highlight the text with your mouse and press **Ctrl+C** (or **Cmd+C** on a Mac). If the message includes an image (a screenshot of an invoice, for example), take a screenshot of that too — look for a button like **Share** or **Screenshot** on your device. Example: a message you might want to check — *"Hi Mom, I dropped my phone and I'm using a friend's number. Can you urgently transfer $450 to this account for a new one? I'll pay you back tonight."* You'll know it worked when the suspicious text is sitting in your phone or computer's clipboard, ready to paste, and you have not tapped any link inside it.

Step 2 — Open your AI assistant in a separate window

Open a fresh browser tab or your AI app — whichever you normally use — so you can see the suspicious message and the assistant side by side. The reason for a separate window: if the scam message came through email, replying inside your email app can be risky, and a fresh chat keeps things clean. Look for the chat box, which is usually a large text field at the bottom of the screen, sometimes with a placeholder like **"Message ChatGPT"** or **"Ask Gemini"**. If you don't see a chat box, look for a button like **New chat** or a **+** icon at the top-left. If your assistant is on a paid plan, this check works the same; free plans handle short messages easily. Example: open the assistant you already use, and start a brand new chat titled something like *"Check this message for scam signs"*. You'll know it worked when you can see an empty chat with a text field ready for typing, and the original message is still in your clipboard.

Step 3 — Paste the message and ask for a scam-spotting breakdown

Tap inside the chat field and paste the message with **Ctrl+V** or **Cmd+V** (or long-press and **Paste** on a phone). Then, on a new line under the pasted message, ask the assistant to look for scam warning signs. A good prompt is short and direct — you don't need any special wording. If the assistant offers extra features (like web search, voice mode, or "deep research"), you can leave them off for this task; the default chat is plenty. Example prompt to copy: *"I got this message and something feels off. Can you list any scam warning signs you notice? Please look at the tone, the request, any pressure to act fast, and anything that doesn't quite add up. Don't tell me what to do — I just want a list of red flags."* You'll know it worked when the assistant replies with a numbered or bulleted list pointing out specific things in the message (for example: *"1. Urgency — the word 'urgently' is a pressure tactic"*).

Step 4 — Ask the assistant to check the link and the sender

Now add the sender's details and any link, and ask the assistant to look at those too. The **From** field on an email, the phone number on a text, or the username on a DM can all reveal patterns — for example, a "PayPal" email sent from a Gmail address, or a "bank" link that doesn't start with the bank's real website. Paste any link in full, exactly as it appeared (don't tap it). Then ask the assistant three things: who the sender claims to be, whether the link looks like a real address for that organization, and whether the sender's address matches the company they say they represent. If a tool can browse the web, you can also ask it to summarize what the public link actually points to. Example: *"The sender claims to be from USPS. Their email is 'usps-delivery-support@gmail.com' and the link is 'usps-tracking.example.com/pay-now'. Does that look like a real USPS*

address? Just point out anything suspicious — don't visit the link for me." You'll know it worked when the assistant flags the mismatch — a Gmail address for a big company, or a link that almost spells the brand but is slightly off.

Step 5 — Read the summary with your own judgement

The assistant's list is a **starting point**, not a verdict. Take a moment to read its answer and ask yourself: do these warning signs match my gut feeling about the message? If the assistant flagged several red flags (urgency, money request, weird sender, off-brand link), treat the message as a scam — don't reply, don't click, and don't send any money. If you want to verify, contact the real organization using a phone number from their **official website** (type the address into your browser yourself; don't use any number from the suspicious message). Delete the message afterwards. If the assistant found no red flags but the message still feels wrong, trust your feeling — scams evolve faster than any checklist, and no AI can promise a 100% clean bill of health. 💬 Example: ask a follow-up like *"Based on the red flags you found, is this more likely a scam or a real message? Give me your best guess and a one-sentence reason."* You'll know it worked when you can clearly say to yourself either "I'm confident this is a scam, I'm deleting it" or "I'm confident enough to verify it through the official channel — and I still haven't clicked the link in the message."

Common mistakes

- **Tapping a link "just to look".** A scam site can install harmful software the moment it opens, or quietly capture data. Always check links in the assistant's chat, never in your browser.
- **Asking the assistant to "visit the link and tell me what's there".** Some assistants can browse, but pasting an unknown URL into a tool you don't fully understand adds risk. Ask for an analysis of the URL itself (the words and structure) instead.
- **Trusting the message because the writing sounds perfect.** Perfect grammar used to be a sign of a real company. With AI, polished writing is now a common sign of a scam. Polished ≠ safe.
- **Ignoring pressure words.** "Urgent", "today only", "act now", "or your account will be closed" are some of the strongest scam signals. If you see them, slow down — that is exactly what the scammer doesn't want you to do.
- **Replying to ask "is this real?".** Replying tells scammers your number is active. If you genuinely want to verify, contact the organization through a channel you found yourself, not through the message.

What this means for you

- **In everyday life:** treat any unexpected message asking for money, passwords, codes, or quick action as suspicious until you've checked it with your AI assistant. This covers texts from "your bank", "USPS", "IRS", "CouriersPlease", and even messages that look like they're from a family member using a new number.
- **For your business or work:** scam messages often target work emails (fake invoices, fake login pages, fake boss messages). Run the same five-step check on any work message that asks for payment changes, gift card purchases, or password resets. When in doubt, talk to a colleague in person before acting.
- **If you're just getting started:** keep one AI assistant open in a tab during the day. The moment a message feels off, paste it in and run through the checklist. After two or three times, the whole check takes under a minute — and you'll start spotting red flags on your own before you even paste.

Try it now

Find one message in your phone or inbox that you've been unsure about (not from the last hour — something older is fine). Paste it into your usual AI assistant with the prompt from Step 3, and read the first three red flags it lists. That's your first five-minute check.

How to Spot a Deepfake Video or Voice Clone Before You Share It

5 steps · Read online

> **In short:** Look at the eyes, mouth and face edges for small glitches; listen for flat, robotic speech; check who sent it and what it asks. If anything feels off, don't forward it — verify with the real person through a trusted channel first.

A short clip of a public figure saying something outrageous lands in your family group chat. A frantic voicemail from your "boss" asks for gift cards. None of these are obviously fake, and that's exactly the point. A **deepfake** is a video, image or audio clip made with AI to look or sound like a real person who never actually said or did the thing, and **voice cloning** (reproducing someone's voice convincingly from a short audio sample) is now easy and cheap. The good news: a few careful seconds of looking is usually enough to catch the most common giveaways. Even so, no single check is perfect, so your final safety net is always verifying with the real person before you share.

Before you start

- The clip itself (saved to your phone or laptop, or a link in a message)
- A web browser (for the reverse image check in Step 4)
- About 2–3 minutes per clip
- No account, no app, no paid plan needed

Step 1 — Watch the face for small glitches

Pause the video and study the face for about 10 seconds. Focus on the **eyes, the edges of the face, and the mouth**. Common giveaways: eyes that don't quite blink in time; glasses that seem to slide or warp when the head turns; teeth that blur into a single white block when the mouth moves; a face outline that looks soft or "painted" where it meets the hair and background; lighting on the face that doesn't match the rest of the scene. If you can't pause or zoom, slow playback to 0.5x and look again. On a phone, tap the screen to bring up the playback speed option — it's usually a small "1x" button near the timeline. You'll know it worked when you spot at least one thing that feels off — or, just as usefully, confirm the face looks natural.

Step 2 — Listen to the voice for flat delivery

Close your eyes for a moment and just listen. Real human speech has small imperfections: a quick breath, a half-stutter, a sentence that gets quieter when the speaker is unsure. AI-generated voices (TTS, short for text-to-speech — software that turns written words into spoken audio) tend to be too clean, too even in volume, and oddly flat in emotion. Listen for sentences with no breath at all, words that all sit on the same pitch, or a tone that doesn't match the situation (someone supposedly panicking, but the voice sounds calm and steady). On a phone, plug in headphones so background hiss doesn't mask the small tells. You'll know it worked when you can tell whether the voice feels human or **synthesized** (made by a computer, not a real person).

Step 3 — Read the message around the clip

Before you forward anything, stop and read the context. Who sent it? Is it really them, or just a forwarded link from someone you trust? What does the clip ask you to do — donate money, click a link, send a code, stay quiet about something? Most scams lean on **urgency, secrecy or big emotional claims**: "send this in the next hour", "don't tell anyone", "this is huge." Those are the same red flags as any scam, and they're a strong reason to slow down. If you can't answer "who, what, and why" from the message alone, that's already a signal. You'll know it worked when you have a clear answer to those three questions — or a clear reason to be suspicious.

Step 4 — Run a quick source check

Take a screenshot of the person's face in the video (or use the platform's "save image" option). Open a reverse image search (commonly available on **Google Images**, **TinEye** or **Yandex Images** — no account needed) and upload the picture. Look for earlier appearances of the same face. If the face is supposed to be a public figure but the search shows nothing, or if older photos don't match the person in the clip, treat it as a likely fake. If the clip is supposedly a private person, send a quick message to that person on a channel you already had before the clip — not the one the clip came from — and ask if they really sent it. 💬 Example: "Hey, I got a voice note from someone claiming to be you asking about gift cards. Did you actually send it? Just checking." You'll know it worked when you can either confirm the source or decide it isn't worth the risk.

Step 5 — Decide what to do next

If any of the checks above raised doubt, the safest move is simple: **don't share it**. Reposting a fake — even with a "this might be fake" caveat — still spreads the original. Tell the person who sent it, politely, that you couldn't verify it, and ask them where they got it. If the clip claims to be a real person you know, give them a heads-up that someone may be using their face or voice. If the clip is a scam attempt, report it to the platform where you saw it, and if money or account access was involved, contact your bank or the company being impersonated directly. You'll know it worked when you stop and think before tapping "forward."

Common mistakes

- **Trusting the "share if you agree" framing.** If a clip is trying to get you to forward it before you've checked, that's the scam working. Slow down — the share button can wait 30 seconds.
- **Assuming a friend's name means a friend sent it.** Accounts get hacked and contacts get **spoofed** (a message made to look like it came from someone you know). Verify through a separate channel before trusting.
- **Relying on a single "AI detection" app.** No consumer app is reliable enough to be your only check, and some detection tools themselves are scams. Your eyes, ears and a quick source check beat any single app.

What this means for you

- **In everyday life:** Treat any unexpected video or voice message that asks for money, codes, or secrecy as a scam until proven otherwise. A two-minute check protects your bank account and your reputation.
- **For your work:** If a "colleague" sends an unusual request by voice or video, especially one involving wire transfers or credentials, confirm it through a known channel like a phone call or in-person chat before acting.
- **If you're just getting started:** Save this page and use it as a checklist the next time a clip makes you feel strongly — good or bad. Strong feelings are exactly what fake clips are designed to trigger.

Try it now

Pick one short video you saw recently that made you feel a strong reaction. Spend two minutes running Steps 1, 2 and 4 on it. Even if it turns out to be real, the practice makes the next check faster and easier.

Money & admin

Build a Simple Personal Budget Spreadsheet Using Any AI Assistant

7 steps · Read online

By the end of this guide, you'll have a working monthly budget in Excel, Google Sheets, or Apple Numbers — even if you've never built a spreadsheet before. You'll simply ask an AI assistant (think of it as a chat-style helper that writes things for you) to design the structure, then copy the result into your favorite spreadsheet app. No formulas memorising, no finance jargon, no fuss.

Before you start

- An account with any AI assistant. ChatGPT, Gemini, Claude, Copilot, Le Chat, and DeepSeek all work for this task — pick whichever you already have. Most have a free tier.
- A free spreadsheet app: Microsoft Excel (often pre-installed on Windows), Apple Numbers (on Mac), or Google Sheets (free in any browser at sheets.new).
- About 15 minutes the first time. A bit less once you've done it.

Step 1 — Open your AI assistant

Visit the website of your chosen AI assistant in a browser tab, or open its app on your phone. Sign in if needed. You'll see a chat box — usually a large text field at the bottom of the screen with placeholder text like "Message ChatGPT" or "Ask Gemini" inside it. If you don't have an account yet, look for a "Sign up" or "Get started" button; sign-up usually takes under a minute with an email address. Once you're in, the chat sits quietly waiting — the AI won't reply until you send something in Step 2. If your assistant offers a model selector (a small drop-down at the top of the chat), the default option is fine for budget work; no need to change it.

You'll know it worked when you can see the empty chat box and your cursor is blinking inside it.

Step 2 — Send the budget prompt

Click inside the chat box and type (or paste) your request. Press Enter, or tap the small upward-pointing paper-plane icon ➤ next to the box, to send. Within a few seconds, the AI starts typing a reply, and a budget table usually appears — complete with column headers like "Category", "Planned", "Actual", and a few sample rows. Some assistants (ChatGPT, Gemini, Claude) also offer a downloadable spreadsheet file; look for a small spreadsheet icon or "Download" link above the reply. If you see a paperclip 📎 icon for attachments, ignore it; you don't need to upload anything.

💬 Example: "Design a simple monthly budget spreadsheet for a single person earning a regular salary. Include columns for Category (Rent, Groceries, Transport, Utilities, Entertainment, Savings, Other), Planned amount, Actual amount, and Difference. Add 2 example rows per category. Keep it under 20 rows total."

You'll know it worked when the AI replies with a clearly formatted table — or a downloadable file link.

Step 3 — Check what the AI produced

Read the reply slowly. You'll usually see one of three things: (a) a neat table with rows and columns inside the chat, (b) a downloadable spreadsheet file, or (c) plain text describing the layout. Look for sensible categories (housing, food, transport, savings), clear column headers, and at least one mention of a formula — typically a "Difference" or "Total" column that subtracts one number from another. If the response feels too complex (mentions macros, pivot tables, or technical jargon), just reply asking for a simpler version. If the AI gave you plain text instead of a table, ask it to "format that as a table" in the same chat — it usually complies in one reply. Don't worry if the numbers are placeholders; you'll swap them for your own in Step 6.

You'll know it worked when you can clearly see columns, rows, and at least one total or formula reference in the reply.

Step 4 — Open your spreadsheet app

Switch to (or open) your spreadsheet app: Excel on Windows, Numbers on Mac, or Google Sheets in any browser (sheets.new is the fastest way to start). You'll land on a blank grid — columns labeled A, B, C across the top, rows numbered 1, 2, 3 down the side. In Google Sheets, click the colorful "+" button to start a new blank spreadsheet. In Excel, choose "Blank workbook". If spreadsheets are new to you, know this: each little rectangle is called a "cell", and you click one to type into it. The cursor blinks in the top-left cell (A1), waiting for your first entry.

You'll know it worked when you see an empty grid of cells ready to receive your budget.

Step 5 — Copy the AI's table into your spreadsheet

If the AI gave you a downloadable file, click "Download" and open it in Excel or Google Sheets — it lands as a ready-to-edit spreadsheet, and you can skip straight to Step 6. If the AI gave you a table inside the chat, click the small "Copy" icon that appears when you hover over the table (or highlight the text manually, then press Ctrl+C on Windows or Cmd+C on Mac). Click cell A1 in your empty spreadsheet and paste with Ctrl+V or

Cmd+V. The categories and column headers appear neatly in row 1, with example data filling the rows below. If the paste comes out as one long paragraph instead of a grid, undo with Ctrl+Z and try Ctrl+Shift+V instead — most apps offer a "paste as plain text" option that cleans up the formatting.

You'll know it worked when the AI's budget rows are sitting in your spreadsheet, ready to edit.

Step 6 — Add your real numbers and a totals row

Click each cell in the "Actual" or "Amount" column and type your own real spending for the past month — rough estimates are fine if you don't have receipts handy. Then click an empty cell directly below your last row and type a label like "Total". In the cell next to it (one column to the right), type =SUM(B2:B10), replacing B2 and B10 with the first and last row numbers of your actual data. Press Enter, and that cell instantly shows the sum. If your spreadsheet puts amounts in column C or D instead, change the letter to match. If you see #REF! or #VALUE! instead of a number, double-check the formula and the range — usually one cell was empty or the letters don't line up with your data.

You'll know it worked when the Total cell shows a single dollar figure that updates automatically as you change any of the amounts above it.

Step 7 — Save the file and make it yours

Save your file: in Google Sheets, click the title at the top-left to rename it (it saves automatically); in Excel, press Ctrl+S. From here, you can add new categories by typing into new rows, duplicate the sheet for a new month (look for a small "+" at the bottom of the screen), or come back to the AI assistant for more formulas like =AVERAGE() to see your typical weekly grocery spend. Try one small change right now — edit a number and watch the Total update. That single moment is when spreadsheets start to feel less like a chore and more like a quiet little superpower.

You'll know it worked when you've saved a file with your name on it, full of your real numbers, and the totals update themselves whenever you change anything.

Common mistakes

- **Asking for "the best budget spreadsheet".** There's no universal best, and the AI may invent complex features you don't need. Stick to simple columns (Category, Planned, Actual, Difference) and add complexity later only if you genuinely want it.
- **Pasting a screenshot of the AI's table into your spreadsheet app.** Most apps treat images as pictures, not data — so the cells stay empty. Use copy-and-paste from the text table, or download the file the AI offers.
- **Trusting the AI's placeholder numbers as personal advice.** The example figures are made up to show structure, not to recommend how you should spend. Replace them with your own real spending, and treat any rule-of-thumb percentages the AI mentions as a starting point rather than gospel.

What this means for you

- **In everyday life:** A 10-minute weekly glance at your "Actual" column tells you where your money actually goes — usually a few small surprises (daily coffees, ride-shares, app subscriptions) you can adjust without much pain.
- **For your business or work:** The same template scales to a side-project budget, a freelance invoice tracker, or a small team's monthly expenses — just rename the categories.
- **If you're just getting started:** Open one AI assistant, paste the prompt from Step 2, and download whatever it gives you. Don't try to perfect it. A rough budget you actually use beats a perfect one gathering dust in a folder.

Try it now

Open your favorite AI assistant right now, paste the prompt from Step 2, and download the spreadsheet it produces. Spend two minutes renaming one column to something that matters to you — "Coffee", "Rent", "Dog walk" — and you're already further along than most people ever get.

Decode Any Confusing Phone, Internet, or Medical Bill in 10 Minutes

7 steps · Read online

> **In short:** You don't need an accounting degree to question a bill. Copy the line items into any AI assistant, ask it to flag anything unusual, and use that short list to decide what to push back on. No app, no signup, no jargon.

Bills use jargon to make you give up and pay. This guide shows a calm way to read any bill with an AI as your second pair of eyes, so you pay only what you actually owe. It works for phone, internet, and medical statements alike, and you don't have to finish it in one sitting.

Before you start

- **One bill you want to understand** — paper, PDF, or a photo on your phone.
- **Any AI assistant in a browser or app** — ChatGPT, Claude, Gemini, Copilot, or another you already use. The free tier is enough.
- **About 10 minutes**, in one go or split across two short breaks.
- **A pen or note app** to mark up the line items as you go.

Step 1 — Find a quiet 10 minutes and put the bill in front of you

Sit down with the bill, your phone or laptop, and a piece of paper. Don't try to read it on the bus or while cooking — line items blur into each other when you're distracted. If the bill is paper, snap a clear photo so you have a digital copy to work with. If it's already a PDF in your email, open it on the same device you'll use for the AI assistant, so you can copy text straight from it. Example: a small note on your paper — "Today's goal: figure out what each line means and which ones look off." You'll know it worked when the bill is on screen, your AI tab is open next to it, and you've got a clear few minutes ahead.

Step 2 — Strip out personal details before you share anything

Before any line goes near an AI, remove identifying information — your full name, address, account number, payment card digits, and any claim or policy number. Scribble them out on paper, or cover them with rectangles in a free phone app. The AI only needs the line items, amounts, and dates. Personal details do not help and create unnecessary risk if a screenshot ever ends up in the wrong place. Example: keep visible — "$84.50 — Base Plan — Sept 3 to Oct 2". Black out everything else. You'll know it worked when the version you're about to share looks like a generic sample, with no name or number pointing back to you.

Step 3 — Paste the bill and ask one focused question

Copy the cleaned-up line items and paste them into your AI assistant. Then ask one specific question, not five. The best request is short and direct: "Look at these charges. Which ones look unusual, and what should I ask the company about each one?" That single sentence turns the assistant into a checklist-maker instead of a general explainer. If your first try produces a wall of text, reply "shorten to a bulleted list under 200 words" and it will reset. Example prompt to copy: "Here are the line items from my phone bill. For each charge, tell me: (1) what it usually means, (2) whether it's normal or worth questioning, (3) one question I should ask the provider." You'll know it worked when the assistant gives you a list where each charge has a small label — like "normal," "worth asking," or "likely duplicate."

Step 4 — Read the response like a checklist, not a verdict

The AI's reply is a starting point, not a final answer. Treat each line like a Post-it note: green means "I understand this and it matches my plan," yellow means "I want a human to explain it," and red means "I want this removed or credited." Don't argue with the bill yet — your only job here is to mark each charge. If the assistant flags something it can't identify, write "?" and move on; the next step is where you check those yourself. Example: mark "monthly plan $40" green, "regulatory fee $3.20" yellow (you don't know if it's accurate), and "international roaming $27" red (you didn't travel). You'll know it worked when you can count the yellow and red marks on one hand and feel ready to look them up.

Step 5 — Double-check the suspicious lines with the company itself

For every yellow or red mark, search the provider's official help pages for the exact phrase. Type the term in quotes — for example, `"universal service fund"` or `"device installment"` — and read the company's own definition. This second pass protects you from a confident-but-wrong AI answer. Most bills also have a phone number, chat, or "contact us" page; note the one you'll use so it's ready when you need it. Example search: type `"administrative fee"` plus the company's name into your browser to see how they explain the charge in their own words. You'll know it worked when every red and yellow line has either a clear explanation in your own words or a one-sentence question written next to it.

Step 6 — Send one short, polite message with a single focused question

Now write to the company. Keep it short, polite, and focused on one line item per message — not the whole bill. Front-load the question: "I'm reviewing my October bill and noticed a $27 charge labeled 'international roaming' for dates I didn't travel. Can you confirm what triggered this and which dates it covers?" Include the bill date and exact amount. If they answer clearly, decide whether to accept it, ask for a credit, or escalate. Example message you can adapt: "Hi, on bill dated 3 October, line 'Equipment Fee — $12.00' appeared for the first time. I don't recall being

notified about a new charge. Can you explain what this covers and the start date? Thank you." You'll know it worked when you get a specific answer — not a copy-paste paragraph — that either makes sense or gives you grounds to escalate.

Step 7 — Decide your next move based on what they say

If the answer resolves it, mark it green and move to the next yellow. If it's a copy-paste non-answer, reply once more asking for a specific date or reference number. If the charge turns out to be wrong, ask politely for written confirmation of the credit and how it will appear on your next bill. For medical bills especially, many regions have a patient advocate or billing hotline whose job is to translate line items. Save your email thread in case you need it later. 💬 Example closing reply — "Thanks. Could you confirm the credit amount in writing and the date it will appear on my next statement?" You'll know it worked when every flagged line is now green, yellow (waiting on the company), or removed — and you have a paper trail for each.

Common mistakes

- **Pasting a bill with your name, address, or account number still visible.** The fix: redact first, always. The AI only needs the line items and amounts to do this job.
- **Asking the AI ten questions at once.** The fix: one focused thing, then a follow-up. Shorter prompts give shorter, more useful answers.
- **Treating the AI's answer as proof of a billing error.** The fix: check the company's own help pages or a previous bill before you push back. The assistant can misread a code as a charge.
- **Sending a long, angry email with every yellow line at once.** The fix: one focused question per email gets a real answer; a wall of complaints gets a template.

What this means for you

- **In everyday life:** any time a bill arrives — phone, internet, streaming, utilities, medical — you can run it through this 10-minute routine and decide what to question, instead of paying on autopilot.
- **For your business or work:** the same habit works for vendor invoices, software subscriptions, and freelancer contracts. The AI is a triage nurse, not a doctor; the human on the other end still confirms.
- **If you're just getting started:** start with the bill that's been bugging you the longest. Even if you only save a small amount, the next bill will feel easier to read.

Try it now

Grab the most recent bill sitting on your kitchen counter, snap a photo, scribble out your name and account number, and paste the line items into any AI assistant with the prompt from Step 3. Ten minutes from now, you'll have a checklist — and a clearer head.

How to read and summarize a long lease or insurance policy with Gemini

6 steps · Read online

> **In short:** You can upload a long lease or insurance PDF to Google Gemini and ask for a plain-English summary that highlights the clauses you actually need to read before signing. The free version handles most everyday documents, and you don't need any technical skill.

A 40-page lease or a 25-page insurance policy is the kind of thing most people skim, sign, and regret later. Gemini can read the whole document for you, pull out the parts that cost money or restrict what you can do, and explain the legal-sounding words in everyday language. This guide walks through that process step by step, whether you're moving into a new apartment, choosing a health plan, or reviewing a contract your employer sent over. Gemini can give you a strong first pass — but remember, it's not a lawyer, and for big decisions the original document still matters.

Before you start

- **A Google account.** If you use Gmail or YouTube, you already have one.
- **A browser or the Gemini app.** On a computer, open any web browser. On a phone, install the Gemini app from your app store if you don't have it yet.
- **The PDF file.** Save the lease or insurance document to your phone or computer so you can find it quickly. If it's a paper document, scan it to PDF first — most phone camera apps can do this.
- **About 10 to 15 minutes.** That's all it takes to go from a confusing PDF to a clear summary.
- **Free or paid plan.** The free version of Gemini handles most everyday documents well. Some advanced features — like reading very long files or uploading many files at once — may require a paid plan; Gemini will tell you if it hits a limit.

Step 1 — Open Gemini and start a fresh chat

In your browser, go to **gemini.google.com** and sign in with your Google account if prompted. If you're on a phone, open the **Gemini app** instead. You'll see a clean chat window with a message box near the bottom and your Google profile picture in the corner. Click or tap inside the message box so the cursor blinks — that's how you know you're ready to type. Starting a new chat, rather than reopening an old one, keeps the document summary focused and easier to find later.

*If you don't see a chat box, look for a button labeled **"New chat"** or a **pencil icon** somewhere on the page.*

You'll know it worked when you see an empty message box with a blinking cursor, and your Google account name or avatar visible somewhere on screen.

Step 2 — Upload your PDF

Next to the message box, look for an icon that looks like a **paperclip**, a **"+"**, or the words **"Upload"** or **"Files"**. Click or tap it. A file picker will open — the same window you use to attach files to an email. Navigate to wherever you saved the PDF (your Downloads folder, Files app, or Google Drive), select the lease or insurance document, and confirm. The file will start uploading, which usually takes a few seconds for a normal-sized PDF. When it's done, you'll see the file name appear in the chat as a small chip or attachment.

*If the icon names don't match, try the **"+" or plus button** — most chat interfaces tuck uploads there.*

You'll know it worked when the filename shows up in the conversation as an attachment and the message box is ready for your first question.

Step 3 — Ask for a plain-English summary

Now type your first message. Be specific about what you want — vague prompts give vague answers. Ask Gemini to summarize the document in everyday words, and tell it which topics matter to you (money, rules, deadlines, restrictions). Press **Enter** or tap **Send**. Gemini will read through the PDF and write back a structured response. Don't worry if it takes 10–20 seconds — that's normal for long documents. Read the summary carefully, but remember it's a starting point, not the final word on what the contract means.

Example: *"Please summarize this lease in plain English. I want to know what I have to pay each month, what I'm not allowed to do, what happens if I break the rules, and how I can end the lease."*

You'll know it worked when Gemini returns a summary that mentions specific dollar amounts, dates, or rules from your document — not generic advice about leases in general.

Step 4 — Ask Gemini to flag the high-stakes clauses

A summary is useful, but the real value is knowing which sections could cost you the most money or get you in trouble. Send a second message asking Gemini to point out the parts you should read closely. The clearer you are about what "high-stakes" means to you (money? termination?

liability? privacy?), the more useful the list will be. If the list feels too long, ask Gemini to rank the top three so you don't have to chase every detail.

Example: *"Which clauses in this document could cost me the most money or cause the biggest problems if something goes wrong? List them and quote the key sentence from each one if you can."*

You'll know it worked when you see specific clause names, dollar amounts, or quoted phrases from your actual document — not generic warnings about "all contracts".

Step 5 — Ask follow-up questions about any confusing term

If a clause still doesn't make sense, copy a sentence from it and ask Gemini to explain it like you're twelve years old. This is the part most people skip, and it's often the most valuable — a single confusing phrase ("indemnify and hold harmless," "actual cash value," "joint and several liability") can hide thousands of dollars in real-world consequences. If Gemini's explanation still feels too formal, ask again in simpler words. You can keep asking as many follow-ups as you want in the same chat.

Example: *"What does 'indemnify and hold harmless' mean in everyday language? And does this clause apply to me as the renter, or only to the landlord?"*

You'll know it worked when you can explain the term to a friend in one short sentence, in your own words.

Step 6 — Save the summary somewhere you can find it

Once you have a summary and a list of important clauses worth re-reading, copy the key parts into a notes app, an email to yourself, or a simple document. Don't rely on the Gemini chat history alone — those conversations can be hard to find weeks later, and uploads sometimes expire. A short note with the date, the document name, the three biggest risks, and any open questions you still have is usually enough. If you sign the document, attach or link your notes so future-you knows what you agreed to.

If you'd rather skip the copy-paste, take a screenshot of the summary and save it to your phone. It's not searchable later, but it's better than nothing.

You'll know it worked when you can open your notes app right now and see, at a glance, the three most important things you learned from the document.

Common mistakes

- **Asking "summarize this" with no direction.** Gemini will give you a generic overview that misses what you actually care about. Tell it which topics matter: money, deadlines, restrictions, termination, who pays for what.
- **Treating the summary as the final answer.** Gemini is a starting point, not a lawyer or insurance advisor. For big decisions — a house, a surgery, a long business contract — re-read the original document or talk to a professional before you sign.
- **Uploading a scanned image instead of a real PDF.** If your document is just photos of pages, Gemini may not be able to read the text inside. Open the file and try to select a sentence; if you can't highlight any words, it's an image scan, not a text PDF, and you'll need a higher-quality scan with selectable text.

What this means for you

- **In everyday life:** Before you sign a lease, an insurance renewal, or a phone plan, you can spend ten minutes with Gemini and walk away knowing what you're actually agreeing to — instead of finding out the hard way when a fee shows up.
- **For your business or work:** Vendor contracts, freelance agreements, and client terms of service are exactly the kind of long PDFs that benefit from a quick AI pass before you forward them to legal or a partner.
- **If you're just getting started:** Pick one document you already have sitting in your email — a recent subscription, an old lease, even a saved terms-of-service page — and run it through Steps 1 to 3 today. Ten minutes, one real document, no risk.

Try it now

Open gemini.google.com (or the Gemini app), upload any PDF you already have — even an old utility bill or a saved terms-of-service page — and ask for a plain-English summary in the next two minutes. You'll know the whole flow works when you see your first AI-written summary of a real document, and you can scroll back through the conversation to find it again.

How to Use ChatGPT to Explain a Confusing Bill or Invoice

6 steps · Read online

> **In short:** Open ChatGPT, attach a photo of the bill, and ask it to translate the jargon into plain English. ChatGPT can explain what each charge probably means — but it can't verify the bill itself is correct, so call the company if something looks off.

Every month, millions of people stare at a bill packed with terms like "regulatory cost recovery fee," "prorated charge," or "facility rental" and have no idea what they actually paid for. You don't have to guess anymore. In about five minutes, you can turn that wall of small print into a short, plain-English explanation you can actually use.

Before you start

- **A ChatGPT account.** A free account works; ChatGPT Plus and Team give you more image uploads per day and faster replies.
- **The app or the website.** On a phone, install the **ChatGPT** app (iOS or Android). On a computer, open **chat.openai.com** in any modern browser.
- **The bill itself.** Either a paper copy on a table, or a digital version (PDF, screenshot, or photo from email).
- **Roughly 5–10 minutes.** Most of that is reading the explanation, not typing.

Step 1 — Get a clear photo of the bill

Lay the paper flat on a table in good light, then hold your phone straight above it — directly down, not at an angle. Make sure all four corners of the page are inside the frame and the small print at the bottom is readable. If the bill is already digital (PDF in your email, screenshot in your camera roll), use that file directly instead of photographing the screen.

If it looks different: If you're working from a paper bill and the corners keep curving in the photo, place a heavy book on the edges. If you're screenshotting an email bill, scroll down first so the whole thing fits in one image.

Example: Take one photo straight down, then pinch to zoom and check the smallest line — if you can read it, ChatGPT probably can too.

You'll know it worked when you can read every line yourself when you zoom in. If the photo is blurry, cropped, or shadowy, ChatGPT will guess — and guesses about money are a bad idea.

Step 2 — Open ChatGPT and start a fresh chat

Open the ChatGPT app on your phone, or go to **chat.openai.com** in your browser. Look for a small **pencil icon** or a button labeled **"New chat"** — usually at the top-left corner of the screen. Tap it. A clean conversation opens, with just an empty message box at the bottom and no previous messages above it.

If it looks different: Some versions show a **speech bubble** with a plus sign, or a **"+"** button in the same spot. If you can't find it on mobile, tap the **menu icon (three lines)** in the top-left — your chats list opens, and you can tap "New chat" from there.

You'll know it worked when the screen shows a blank chat with only the message box at the bottom. Starting fresh matters because old messages can confuse the answer.

Step 3 — Attach the bill photo to your message

Inside the message box, look for a **paperclip**, a **plus (+)**, or an **image icon**. Tap it, then choose **"Photos"** or **"Camera"** on your phone, or **"Upload file"** on a computer. Pick the bill photo from your library. A small thumbnail of the image appears inside your message area, waiting to be sent.

If it looks different: On some versions of the app the image icon sits on the **toolbar just above the keyboard**, not inside the box itself. On a computer, you can also **drag and drop** the photo straight onto the chat window — no clicking needed.

You'll know it worked when a small preview of the bill is sitting in your message area, ready to go. If you see only text and no image, the upload didn't go through — try again.

Step 4 — Ask ChatGPT to explain it in plain English

Now, in the message box, type a clear request. Tell ChatGPT what kind of bill it is and ask it to walk you through each charge in everyday language. Tap **send** (the arrow or paper-plane icon). Within a few seconds, ChatGPT reads the image and starts typing a plain-English explanation line by line.

If it looks different: If the chat box appears empty after you tap send, check that your message contained **both** the image and your text — sometimes the image uploads but the words disappear. Resend them together.

Example: "This is my home internet bill. Please list every charge on it in plain English, and flag anything that looks unusual or higher than I'd expect for a typical plan."

You'll know it worked when ChatGPT replies with a numbered list or short paragraph that names each line item and says what it means, then asks if you want a closer look at any of them.

Step 5 — Ask follow-up questions if something is still unclear

If a term or charge doesn't make sense, just type another message in the same chat. ChatGPT remembers what you just talked about, so you can ask **"What does 'prorated' mean here?"** or **"Is this a normal charge for internet?"** without re-uploading the bill. The more specific your question, the more useful the answer tends to be.

If it looks different: If ChatGPT seems to have lost the thread or replies with something generic, scroll up and confirm the photo is still visible in the conversation. If it isn't, attach it again in your new message.

Example: "What does 'regulatory cost recovery fee' usually cover on these bills, and is it something I can ask the company to remove?"

You'll know it worked when you get a direct, plain answer to your specific question — not a generic reply and not a request to upload the bill again.

Step 6 — Save the explanation for later

Once ChatGPT has explained the bill, you'll probably want to keep that explanation. Look for a **copy icon** near the assistant's reply, or simply **highlight the text**, copy it, and paste it into Notes, an email, or a message to yourself. Some plans also offer a **"Share"** button near the top of the chat that creates a link to the whole conversation.

If it looks different: If you don't see any copy or share option, the simplest fallback is a **long-press on the text** (on a phone) or a **right-click → Copy** (on a computer), then paste wherever you keep notes.

You'll know it worked when you can open your Notes app (or wherever you pasted it) and read the explanation without going back to ChatGPT.

Common mistakes

- **Sending a blurry or cropped photo.** ChatGPT will fill in the gaps with guesses — and guesses about money are dangerous. Fix: retake the photo flat, in good light, with all four edges visible.
- **Asking "is this bill correct?"** ChatGPT can't compare the bill against the company's actual records, so it can only explain what the text says. If something still looks wrong after the explanation, call the company's billing line directly.
- **Uploading bills full of sensitive numbers.** Full account numbers, Social Security numbers, or detailed medical codes don't need to be on the photo — the charges do. Fix: cover sensitive numbers with a piece of paper or a finger before snapping, or scribble them out with a marker.

What this means for you

- **In everyday life:** Next time an unfamiliar charge shows up on a phone, utility, or medical bill, you don't have to guess. A two-minute chat can translate the jargon into a list you actually understand.
- **For your business or work:** The same trick works on supplier or contractor invoices — useful for spotting strange line items before you pay, and for explaining a charge to a teammate who wasn't on the original call.
- **If you're just getting started:** Pick one bill you've been avoiding opening, snap a photo, and run it through. That single experiment will teach you more than any tutorial.

Try it now

Find the most recent bill you've been ignoring — utility, internet, phone, or medical — and snap a clear photo of it. Open ChatGPT, attach the image, and paste this prompt: **"This is my [type of] bill. Please list every charge in plain English and flag anything that looks unusual or higher than I'd normally expect."** Read the reply before you pay anything.

How to Decode a Confusing Medical Bill with AI

5 steps · Read online

> **In short:** Paste an itemized medical bill or insurance statement into any AI assistant and ask it to translate the jargon, list every charge, and flag anything that looks off — in plain English. The AI explains; it cannot dispute a bill or make medical decisions.

A medical bill arrives in the mail, or an Explanation of Benefits (EOB) shows up in your insurer's portal, and within five lines you've hit "CPT 99213", "allowed amount", and a number that doesn't match what you expected to pay. This guide walks you through using an AI assistant to turn that wall of codes and figures into a plain-English breakdown you can actually read — no medical or insurance background needed.

Before you start

- **The document itself**: an itemized medical bill, an EOB from your insurer, or both. Paper, PDF, or a photo on your phone all work.
- **An AI assistant**: any major one will do — ChatGPT, Gemini, Claude, Copilot, Le Chat, or DeepSeek. Free versions are usually enough for a single bill; longer documents may need a paid plan with a bigger context window (the amount of text the AI can read at one time).
- **A second screen is helpful but optional**: phone for the bill, laptop for the AI, or two browser tabs side by side.
- **About 15 minutes** for your first bill. You'll get faster on the second one.
- **Important limit**: AI can explain what a charge means. It cannot file an appeal, call your insurer, or tell you whether a procedure was medically necessary. Treat its output as a translation, not a verdict.

Step 1 — Make the bill digital

If your bill is on paper, you need it as text the AI can read. Open any free scanning app on your phone — most phones have one built in (look in your camera or notes app for a "scan document" option). Lay the bill flat, line up the corners in the camera frame, and snap. The app saves it as a PDF with selectable text. Already have a PDF in email or your insurer's portal? Just download it to your phone or computer. **If it looks different**: if you don't see a built-in scanner, search your app store for "free document scanner" — any top result works fine. You can also type the charges out by hand, though that's slower. 💬 Example: "Scan the front and back of the bill into one PDF named 'bill-jan2026.pdf'." You'll know it worked when the bill opens on your screen and you can highlight the words with your cursor — or, on a phone, long-press a word and see it selected.

Step 2 — Strip out personal details

Before you paste anything into an AI, take 60 seconds to remove identifying numbers. On a paper bill, use a marker or your finger to cover your full Social Security Number, member ID, and any bank or card numbers. On a PDF, print and mark it up, or copy the body text into a notes app and delete those numbers by hand. **If it looks different**: a quick rule of thumb — leave your name and the provider's name (the AI needs them for context), but remove any number that looks like an account, policy, or government ID. Some bills also print diagnosis codes that are very specific to your health; if that makes you uneasy, swap them for a generic description like "knee injury" before pasting. 💬 Example: "Black out my SSN, member ID, and the last four digits of my card before pasting." You'll know it worked when a stranger could read the document and follow it, but no one could use it to impersonate you or pull up your records.

Step 3 — Paste it in and ask for a plain-English summary

Open your AI assistant in a browser or app and start a new conversation. Paste the cleaned-up bill or EOB into the message field near the bottom of the screen. The message field is often labeled something like "Message" or "Ask anything," though your app may use different words — look for the long text box with a blinking cursor, just above or beside a send button. Ask the AI to explain it, and keep your first prompt simple — one request, one paragraph. **If it looks different**: if the document is very long and the AI says it was cut off, split it into two messages — first half in one prompt, second half in the next — and tell it to combine the analysis. 💬 Example: "Here's my itemized bill from a clinic visit. In plain English, please list every charge, what each one was for, what my insurance paid, and what I still owe. Flag anything that looks unusual." You'll know it worked when the AI gives you a structured answer — usually a short list of charges in everyday language, followed by a total or a clear note that it couldn't calculate one.

Step 4 — Ask follow-up questions about specific terms

Now zoom in. Pick any line you don't understand — a procedure code, a phrase like "allowed amount" or "patient responsibility", a charge that surprises you — and ask the AI about it. This is where the chat format really helps: you don't have to ask everything at once. **If it looks different**: if the AI's reply is too generic, paste the exact line back and ask about that one item only. Still vague? Try rephrasing — instead of "explain this", try "in one sentence, what does a 99213 CPT code usually mean, and roughly how much does it cost without insurance?" And if every line on the bill actually looks clear to you already, still ask the AI to flag any codes or charges it thinks are worth questioning — a fresh pair of eyes often catches something you skimmed past. 💬 Example: "On line 4 it says 'CPT 99213, allowed amount $84.20, patient owes $25'. What is a 99213, what does 'allowed amount' mean, and why is the patient share $25?" You'll know it worked when the AI names the procedure in plain words, explains each money term, and tells you roughly how the numbers relate — usually with a note that exact prices vary by region and plan.

Step 5 — Sanity-check the totals and spot anything odd

Before you trust any of this, do your own quick math. The AI is good at explaining, but it can mis-add numbers or invent a percentage that wasn't in the original. Open your bill or EOB in another tab, grab a calculator (your phone has one), and add the charges yourself. Compare that to the AI's total. **If it looks different**: if the totals don't match, paste the raw lines back and ask it to re-add — most assistants catch their mistake on the second try. If the gap is large, treat the bill, not the AI, as the source of truth, and call the number on the bill to ask about it. 💬 Example: "Please re-add every line on this bill and confirm the total. I'm getting $342.10 when I add it by hand but your last answer said $298.50." You'll know it worked when either the AI corrects itself and the new total matches yours, or it points to a specific line it had missed — and you've decided what to do with the gap.

Common mistakes

- **Pasting a bill with sensitive numbers still in it.** Even well-known assistants log conversations to improve their models. Strip SSNs, member IDs, and card digits first — Step 2 takes a minute and prevents a much bigger headache later.
- **Treating the AI's "this looks normal" as proof.** AI assistants can sound confident even when they're guessing. Use their output as a starting point, then confirm any big charge by calling the provider's billing line — the number is almost always at the top of the statement.
- **Asking the AI whether a charge is fair.** That's a judgment call that depends on your plan, your region, and the specific service. The AI can explain what a code is; only your insurer can tell you whether it should have been covered. For disputes, start with the provider's billing department, and if you want formal help, contact your state or country's insurance ombudsman.
- **Uploading a blurry photo of the bill.** If the AI keeps saying "I can't read this", re-scan in good light with the page flat. Bad input always produces bad output.

What this means for you

- **In everyday life:** the next time a confusing statement shows up, you'll spend 15 minutes instead of an afternoon figuring out what you actually owe — and you'll know exactly which line to call about if something looks wrong.
- **For your business or work:** if you run a small team, you can use the same workflow to help employees decode their first bills after a new benefits plan kicks in. It won't replace HR, but it cuts down on "what does this mean?" emails.
- **If you're just getting started:** try it first on an old, already-paid bill you no longer need to act on. You'll learn how the assistant responds without any pressure — and you'll have a much easier time when a real, current bill arrives.

Try it now

Pick any old medical bill or EOB you can find — even one you already paid. Strip the sensitive numbers, paste the rest into your favorite AI assistant, and ask: "Explain every line in plain English and tell me what I might want to question." That single two-minute ask is the whole skill.

How to Read Through Insurance Policies and Bank Contracts with Claude

6 steps · Read online

> **In short:** Paste any insurance policy, loan agreement, or bank contract into Claude and ask for a plain-English summary. Claude will flag hidden fees, exclusions, and tricky clauses — but it cannot replace a lawyer, and very long documents may need to be split into parts.

Insurance policies and bank contracts are written in a style called *legalese* — language that sounds official but is designed to be hard to read. A 30-page home insurance PDF often hides the answer to "am I actually covered if my pipe bursts?" in paragraph 47. Claude can read that whole document in seconds and explain it like a friend would over coffee. Here's how to use it without losing the details that matter.

Before you start

- A free Claude account at **claude.com** (the free plan is enough for short documents; very long contracts may need a paid plan with a bigger file upload).
- The contract or policy you want to read — either as a PDF file on your device, or as text you can copy from your email.
- About 10–15 minutes for your first try.
- A quiet moment. You'll want to actually read Claude's summary, not just skim it.

Step 1 — Open a fresh chat with Claude

Go to **claude.com** and log in. Click **"New chat"** at the top-left of the screen. A blank conversation box opens. Starting fresh matters because Claude remembers the previous conversation — and you don't want yesterday's homework question mixed up with today's insurance policy. If you're on the phone app, tap the **pencil or "New chat" icon** instead; if you can't find it, look for a button like **"Start new conversation"**. Example: there is nothing to type yet — just open the chat. You'll know it worked when you see an empty message box with a **paperclip icon** (for attaching files) and a **send arrow**.

Step 2 — Attach the document (or paste the text)

Click the **paperclip icon** to the left of the message box, then choose your file (PDF, Word, or a photo of a printed page usually works). If your contract is already open in another window, you can also simply copy the text — **Ctrl+C** on Windows or ⌘**+C** on Mac — and paste it straight into the chat box with **Ctrl+V** or ⌘**+V**. After you attach or paste, you'll see a small file thumbnail or a chunk of text appear in the box. If you don't see a paperclip, look for a **"+" button** or an **upload icon** instead — apps change the icon from time to time. Example: a PDF named `home-insurance-2026.pdf` shows up as a thumbnail before you send. You'll know it worked when Claude replies that it has read the document and is ready for your question.

Step 3 — Ask for a plain-English summary first

Before asking tricky questions, ask Claude to give you the basics. A good first prompt tells Claude who you are and what kind of document it is. Type a short message in the box, then press **Enter** or tap the **send arrow**. Claude will write a multi-paragraph reply covering the main points. If the reply is too short or too long, follow up with "Can you go deeper on the parts a normal person would care about?" Example: "I'm a homeowner, not a lawyer. Please summarize this home insurance policy in plain English. Cover: what am I covered for, what am I NOT covered for, my deductible, and anything that could void my coverage. Keep it under 400 words." You'll know it worked when Claude's reply lists the key sections in everyday language, with no big unexplained words.

Step 4 — Ask Claude to flag the "gotcha" clauses

Now for the part that saves real money. Many contracts include clauses (sentences or whole sections) that limit what the company has to pay you — things like "we only cover damage from sudden events, not gradual wear." Type a follow-up question in the same chat. Claude uses what it already read, so you don't need to re-upload anything. If the reply feels vague, add "Give me at least five specific examples." Example: "Now read it again like a skeptical friend. List every clause that could limit my payout or surprise me. For each one, quote the original wording, then explain what it actually means in plain English." You'll know it worked when Claude gives you a bulleted list with both the original sentence and a plain translation — those are the spots to read twice in the real document.

Step 5 — Compare two offers side by side (optional but powerful)

If you're choosing between two insurance quotes or loan offers, open a second chat (back to **"New chat"**) and upload both documents — or paste both into the same chat one after the other. Then ask Claude to compare them on the points that matter to you. Claude reads both files in seconds and builds a comparison. Be specific about what you want to compare, because "which is better?" is too vague to answer well. Example: "I've pasted two home loan offers below. Compare them on: total interest paid over 5 years, early-repayment penalties, fixed vs variable rate, and any hidden fees. Show your work in a simple table." You'll know it worked when Claude produces a clean table or side-by-side list — and you can spot the better deal in about a minute.

Step 6 — Save the summary where you'll actually find it

Claude's replies disappear if you lose the browser tab. Before you close the chat, copy the most useful reply (the summary, the gotcha list, or the comparison) into a notes app, email it to yourself, or take a screenshot. On a computer, click and drag to highlight the reply, then **Ctrl+C** / **⌘+C**. On a phone, long-press the text and choose **"Copy"** or **"Share"**. Saving this now means you don't have to ask Claude again six months later when you actually need the info. 💬 Example: paste into a note titled `home-insurance-summary-2026` so future-you can find it. You'll know it worked when you can close Claude, reopen your notes app, and see the summary waiting for you.

Common mistakes

- **Pasting only the first page.** Insurance policies often hide the most important exclusions near the end. If the file is long, upload the whole PDF instead of copying a paragraph.
- **Trusting the summary without checking the originals.** Claude is helpful, but it can occasionally miss a clause or misread a number. Always open the real document and confirm any decision that affects money or coverage.
- **Asking vague questions like "is this a good deal?"** Claude needs specifics to compare well. Name the criteria (fees, interest rate, coverage limits, cancellation terms) so the answer is useful.

What this means for you

- **In everyday life:** Reading a car insurance renewal no longer means an hour of squinting. Ten minutes with Claude gets you the parts that matter before you sign.
- **For your business or work:** Comparing supplier contracts or client terms of service becomes a 15-minute task instead of an afternoon — useful for freelancers and small-business owners who don't have legal staff on call.
- **If you're just getting started:** Pick one document you've been meaning to read — a bank fee schedule, a phone plan, an old insurance renewal — and run it through Steps 2 and 3 today. The habit starts there.

Try it now

Find one PDF or email attachment from your bank, insurer, or phone company that you've been avoiding. Open a new Claude chat, upload it, and paste this prompt: *"I'm not a lawyer. Summarize this in plain English, then list the three things in here that could cost me money if I don't know about them."* Read the reply once, then decide whether the document deserves a closer look — that one minute of friction-free reading is the whole point.

Home & family

Build a Smart Packing List for Any Trip in Just a Few Minutes

7 steps · Read online

> **In short:** Tell any AI assistant where you're going, for how long, and what you'll be doing, and it drafts a weather-smart packing list in under five minutes. The list won't book your flight or check the weather for you, but it covers the basics most travelers forget.

You don't need a packing app, a spreadsheet, or a Sunday afternoon of planning. An AI assistant can turn a few plain sentences about your trip into a clear, organized list, and you can tweak it until it fits your real life. This guide walks you through the whole thing, step by step, even if you've never typed a **prompt** (the instruction you give to an AI — the message box is the same one you'd use to text a friend).

Before you start

- An AI assistant you can type to: ChatGPT, Gemini, Claude, Copilot, Le Chat, or any other you already use. A free plan is fine for this guide.
- A device with internet access — phone, tablet, or laptop all work.
- About 10 minutes total: roughly 3 minutes to gather your trip details and 5 to chat with the assistant and tweak the list.
- Optional but helpful: a rough idea of the weather at your destination and any planned activities (hiking, dinners, a wedding, the beach).

Step 1 — Gather your trip details in one place

Open any note app — the built-in Notes app on your phone works fine — and jot down four things: your destination, the dates (or length of stay), the rough weather (warm, cold, rainy, mixed), and the main activities you have planned. Keep it short, like a few text messages to a friend. You'll paste this paragraph into the assistant in the next step, so have it ready to copy. Example: "Lisbon, Portugal, 6 days in mid-October. Mild days around 20°C, cooler nights. Hiking one day, two nice dinners out, sightseeing the rest." If you're not sure about the weather, write "not sure" — the assistant can still suggest layering options and remind you to check the forecast closer to your trip. You'll know it worked when you have a single short paragraph (3–4 sentences is plenty) with all four details sitting in one place, ready to copy.

Step 2 — Open your AI assistant and paste your details

Open the AI assistant you usually use, either in a browser or in its app, and start a fresh chat. If you've never opened an AI assistant before, the easiest way to start is to go to ChatGPT.com in any browser, or open your phone's app store and search for "ChatGPT", "Gemini", or "Claude" — all of them have a free tier and a similar message box. Look for a clearly labeled "New chat" or "+" button (usually top-left or top-right), tap the message box (it often says "Message", "Ask anything", or shows a small paper-airplane icon), then paste the paragraph you just wrote. Example (paste this, then edit the city and dates to match your trip): "I'm going to Lisbon, Portugal for 6 days in mid-October. Days will be mild around 20°C and nights cooler. I'll hike one day, have two nice dinners out, and sightsee the rest." If the chat looks different from what you remember, look for any button that says "New chat", "New conversation", or a "+" sign — it's the way to start fresh and avoid an old conversation mixing in. You'll know it worked when the assistant replies with a starting list or a clarifying question — both are fine, because either way you've kicked off the planning.

Step 3 — Ask for a weather-smart packing list by category

Reply in the same chat with one short follow-up that asks for the list to be organized into clear categories and adjusted for the weather. The assistant usually gives a wall of bullets at first — asking for categories makes it scannable and easy to edit later. Type your request as a new message in the same conversation; the assistant already has your trip details from Step 2. Example: "Please turn that into a packing list organized by category (clothes, toiletries, electronics, documents, activity-specific). Focus on layers for mild days and cooler nights, and call out anything I'm likely to forget." If your assistant gives you one long flat list instead, just reply "Can you group this into categories like Clothes, Toiletries, Electronics, Documents, and Extras?" — it will reorganize without losing anything you already had. You'll know it worked when the response shows clear headers or bold category names and you can scan the whole list in under a minute.

Step 4 — Add your personal must-haves and remove what doesn't fit

The list still lives right here in the chat thread — you don't need to copy it anywhere yet. Just type a new reply to the assistant, the same way you'd text a friend, and tell it what to add and what to drop. Be specific: name items you actually own (your rain jacket, your prescription medications, a USB-C power adapter for the country you're visiting) and items you know you won't use (gym clothes if you won't actually go to the gym, formal shoes if there's no dressy event). Example: "Add a compact rain jacket, my prescription medications, and a USB-C power adapter. Remove the gym clothes and the formal shoes — I won't need them." If you realize mid-edit that you've forgotten an activity (a wedding, a beach day, a conference), mention it in the same message: "Also one beach afternoon and one semi-formal dinner — what should I add for those?" You'll know it worked when the assistant's next reply reflects every change you asked for, and nothing on the list feels like it belongs to a stranger.

Step 5 — Ask for a "things people forget" add-on

Every traveler forgets something — charger cables, a reusable water bottle, sunscreen, a photocopy of their passport. Ask the assistant, in the same chat thread, for a short list of commonly forgotten items based on your specific trip. This single step catches more oversights than the rest

combined, and it works best when you reference the trip details already in the conversation. 💬 Example: "Now add a short 'often forgotten' list tuned to a 6-day city-and-hike trip in mild autumn weather." If the assistant produces a generic list that doesn't quite fit, reply "Skip anything that doesn't apply to a city trip with one day of light hiking" — most assistants will trim it quickly into a tighter version. You'll know it worked when you spot at least two or three items you genuinely would have left at home.

Step 6 — Turn the final list into a copy-able checklist

Ask the assistant, still in the same chat, to convert the final version into a simple checkable list you can copy into your notes app, print, or share with a travel companion. Plain markdown checkboxes (the small `[ ]` boxes) usually paste cleanly into most note apps and stay editable when you tap or click them. Once you have the checklist, tap and hold the assistant's reply, choose "Copy" (or "Select all" then "Copy"), open your notes app, and paste. 💬 Example: "Please rewrite the full list as a checklist I can copy and paste — one item per line with empty checkboxes." If the assistant produces checkboxes that don't paste well into your notes app, just ask "Give me the same list as plain lines starting with a dash instead." Either format works for packing. You'll know it worked when you can paste the result into your notes app, tick items off as you pack, and the boxes actually behave like boxes.

Step 7 — Save the prompt for next time

This whole process works just as well for your next trip. Save your original message from Step 2 (the one with destination, dates, weather, and activities) in the Notes app, an email draft, or wherever you keep travel info — so the next time you travel you only have to swap in the new destination and dates. The bones of the prompt stay the same. 💬 Example to save: "I'm going to [city] for [X days] in [month]. Weather: [warm/cold/rainy/unknown]. Activities: [list]. Please make a packing list by category for mild days and cooler nights, then add an 'often forgotten' section, then format as a checklist." If your trip details change later, you don't have to start over — just paste the new details into the same chat and ask the assistant to update the list, or open a new chat and reuse the saved paragraph. You'll know it worked when you open the saved prompt, change three lines, and get a useful first draft in under a minute.

Common mistakes

- **Skipping the weather detail.** Without it, the assistant defaults to a generic list and you'll end up overpacking or missing a rain jacket. Fix: even a rough guess ("probably mild, maybe rainy") is enough to get sensible layers.
- **Asking for everything in one mega-message.** You'll get a generic list that doesn't sound like your trip. Fix: add your personal details in Step 4 — that's where the list actually becomes yours.
- **Trying to copy and edit the list yourself.** The assistant is faster and catches things you'd miss. Fix: stay in the chat and ask it to make edits — just type your changes as a new message.
- **Trusting the list blindly.** The assistant doesn't know you already own a hairdryer or that you hate wearing jeans. Fix: do one quick pass to add or remove items before you copy the final version.
- **Forgetting to check airline rules.** The assistant can't see your ticket. Fix: double-check luggage size, weight limits, and any banned items (large lithium batteries, for example) on your airline's website before you zip the suitcase.

What this means for you

- **In everyday life:** A packing list built in five minutes beats the Sunday-afternoon panic pack. You'll stop arriving somewhere and discovering you left the charger or the rain jacket on the kitchen counter.
- **For your business or work:** The same prompt pattern works for work trips, conferences, and team offsites. Tweak "activities" to include "client meetings" or "booth setup" and the list adapts in seconds.
- **If you're just getting started:** The easiest first move is to take a trip you already took — even last month's — and run it through the prompt. You'll see what the assistant catches and what it misses, so your real next trip goes even smoother.

Try it now

Open your AI assistant right now and paste this — fill in the brackets — and you've got a working packing list before your coffee gets cold: *"I'm going to [city] for [X days] in [month]. Weather: [warm/cold/rainy/mixed]. Activities: [sightseeing, hiking, dinners, beach, whatever applies]. Make me a packing list by category, then add an 'often forgotten' section, then format it as a checklist I can copy."*

Plan a Week's Meals from What's Already in Your Fridge

7 steps · Read online

> **In short:** You don't need a recipe app or a shopping list to plan a week of meals. Open the fridge, sort what you have into "use first" and "use later", then build a 7-day plan that uses the most perishable items first. About 30 minutes of work, no special tool required.

You stare into the fridge, see half a bag of carrots, two eggs, and some yoghurt that expires tomorrow, and feel that low-key dread of "what do I even make?" This guide turns that moment into a 30-minute plan. It's the same approach cooks and busy families use to cut waste and save money, and you can do it tonight with a pen and a piece of paper.

A quick note on date labels before we start. **Use-by** dates are about food safety — food past its use-by date may not be safe to eat, so use it or freeze it by that day. **Best-before** dates are about quality, not safety — food past its best-before date is often still perfectly fine to eat, just maybe not at its peak. When you read a label later in this guide, that's the difference.

Before you start

- **Time:** about 30 minutes for the first plan, less once you get the hang of it.
- **Tools:** a pen and paper, or the notes app on your phone. Optional: your phone camera for a quick fridge photo.
- **Pantry check:** a quick look in the fridge, freezer, and pantry (the cupboard where you keep rice, pasta, tinned food, oils, and spices).
- **No app or account needed.** This is a hands-on method, not a software tutorial. If you'd like AI help later, any assistant (ChatGPT, Gemini, Claude, Copilot, Le Chat) can take your list and suggest recipes, but that's a separate step.

Step 1 — Empty the fridge onto the counter

Take everything out of the fridge and place it on the kitchen counter, grouped by type: vegetables together, dairy together, meat or fish together, condiments and jars together. The point is to see what you actually own, not just what's visible through the door. What you'll see is a surprisingly large pile, often more than you expected. This is normal — most people underestimate their fridge by about a third. If it looks different on your counter (very small fridge, or a deep freezer), just work top-shelf down so nothing hides. Example: pull out the cheese drawer, the veg crisper, and the little door shelves too. A 5-minute job. You'll know it worked when every shelf is empty and you can see the back wall of the fridge.

Step 2 — Sort into "use first" and "use later"

Make two piles on the counter. **Use first** is anything that will go off in the next 3 days: soft vegetables, opened milk, fresh meat or fish, ripe fruit, leftover containers. **Use later** is hardier stuff: carrots, potatoes, onions, whole eggs, unopened packs, jars, anything from the freezer. Don't throw anything away yet — sorting first helps you decide what to cook, not what to bin. If a piece of meat smells off or vegetables are clearly slimy, set those aside in a small "compost" pile, but be honest: most "iffy" food is still fine to cook today. Remember the date-label note: use-by dates are the safety line, best-before is only about quality. Example: a half-empty milk carton with a 2-day use-by date goes in "use first". A sealed jar of pickles goes in "use later". You'll know it worked when each item on the counter sits in exactly one pile, and the "use first" pile is the smaller one.

Step 3 — Check the pantry and freezer too

Spend 2 minutes looking at the dry pantry and the freezer. Write down — or take a phone photo of — anything you might forget: rice, pasta, tinned tomatoes, lentils, frozen peas, frozen bread, leftover roast chicken. Think of this as your "supporting cast". The fridge gives you the stars; the pantry and freezer give you the carb, the protein backup, and the speed. If your freezer is packed with mystery bags, label them right now with a sticky note ("beef, March 2026") so future-you isn't guessing. Example: "We have basmati rice, two tins of tomatoes, a bag of frozen peas, and half a loaf of bread in the freezer." You'll know it worked when you have a short list of staples you can mix with anything from the fridge.

Step 4 — Pick one "anchor" meal per day

For the next 7 days, give each day one anchor meal that uses up the most perishable ingredients. Spread the fresh meat, fish, and soft vegetables across the week so nothing sits past its prime. The other meals can be simple repeats: breakfasts are usually the same 2–3 things, lunches are often leftovers from dinner, and one or two nights a week can be a "pantry meal" (pasta, omelette, fried rice) that doesn't need any fresh shopping. Example week: Monday stir-fry (uses the soft capsicum and half the chicken), Wednesday frittata (uses the eggs and any leftover veg), Friday curry (uses the sad carrots and a tin of lentils), with toast, yoghurt, and leftovers filling the gaps. You'll know it worked when you can glance at the plan and see that nothing in the "use first" pile is left unassigned by Day 3.

Step 5 — Write the plan in a simple grid

On a piece of paper, draw 7 columns (one per day) and 3 rows (breakfast, lunch, dinner). Fill in each box with one short phrase — not a recipe, just a name. "Eggs on toast", "leftover curry", "pasta with peas and parmesan". The whole plan should fit on one side of an A4 page. If a box is hard to fill, swap two days around so the most perishable food gets cooked earliest. If it looks different on your phone (you prefer digital notes), use a notes app with a simple table — the shape is the same. Example: "Tues lunch: last night's chicken + salad leaves. Sat dinner: pantry fried rice

(rice, egg, frozen peas, soy sauce)." You'll know it worked when you can read the plan in under 30 seconds and know exactly what to take out of the fridge each morning.

Step 6 — Shop only for the gaps

Now look at your plan and make a short shopping list — only for things you genuinely need to make it work. Common gaps: fresh bread, milk, one or two vegetables, maybe a protein. Most weeks you'll find you need almost nothing. Stick to the list; this is the part that actually saves money. If you spot a gap you can't fill (you planned fish tacos but have no tortillas), change the plan, not the shopping list — swap the meal for something that uses what you already have. 💬 Example list: "1 loaf bread, 1 L milk, 1 lettuce, 6 tortillas." That's it. Estimated cost: about $12. You'll know it worked when your shopping basket is noticeably lighter than usual and you haven't bought anything random.

Step 7 — Cook once, eat twice

Pick one evening — usually Sunday or your day off — to cook a slightly larger portion, and box the leftovers for a lunch or a quick dinner later in the week. This is the single biggest time-saver in weekly cooking, and it turns "leftovers" from a chore into the plan. Cooked rice, cooked pasta, roast vegetables, and soups all keep 3–4 days in the fridge and freeze well. If it looks different on your schedule (you work long hours), don't skip this step — instead, cook a double batch on any evening you already have the oven or stove on. 💬 Example: while the curry is simmering, boil an extra cup of rice. Tomorrow's lunch is sorted in 30 seconds. You'll know it worked when at least 2 lunches and 1 dinner this week come straight from the fridge without any new cooking.

Common mistakes

- **Mistake:** Throwing away "iffy" food before checking. **Fix:** Smell it, look at it, and remember that best-before is a quality date, not a safety date for most things. Use-by is the safety date — respect it, but food past best-before is usually still fine if it looks and smells normal. Soft vegetables and slightly tired herbs are usually fine if cooked today.
- **Mistake:** Planning 7 ambitious recipes. **Fix:** Cap it at 2 or 3 new dishes per week. The rest should be simple repeats, leftovers, or pantry meals. A plan you don't cook is not a plan.
- **Mistake:** Shopping before planning. **Fix:** Always plan from the fridge first, then write the shopping list. If you shop first, you'll overbuy and the plan stops mattering.
- **Mistake:** Forgetting breakfasts and lunches. **Fix:** Most plans fail because dinner gets all the attention. Give breakfast one default (eggs, oats, toast) and let lunch be leftovers twice a week.

What this means for you

- **In everyday life:** food waste is a real cost. American government and waste-reduction reports (such as those from Foodbank and Fight Food Waste the US) regularly find that many households throw out a meaningful share of the food they buy — often around one-fifth of the weekly shop. A fridge-first plan usually shrinks that quickly, often by half or more within a month of starting. The exact dollar saving depends on your household size and shopping habits, but reducing waste is the most reliable way to put money back in the food budget.
- **For your business or work:** the same "use what you have first" logic works for office lunches, shared kitchen snacks, and even team lunches. Cook once on Sunday, eat well all week, and avoid the 4pm takeaway trap.
- **If you're just getting started:** do this just for dinners this week. Skip breakfast and lunch planning entirely. Once dinners feel easy, add the rest.

Try it now

Pick one shelf of your fridge — just one — and plan three dinners from what's on it. Not the whole fridge, not the whole week. Just one shelf, three nights, shopping list of nothing. That single small win is how fridge-first planning actually becomes a habit.

Plan a Family Weekend Getaway with DeepSeek: A Step-by-Step Guide

7 steps · Read online

> **In short:** DeepSeek can draft a full weekend itinerary for your family — kid-friendly activities, meal stops, rough travel times — in under five minutes. It does the brainstorming and structuring; you still book the hotels, tickets and restaurants yourself.

Picture this: it's a Wednesday evening, the kids are finally asleep, and you're trying to figure out where to go this Saturday and Sunday. DeepSeek — the free AI chatbot from the Chinese AI lab DeepSeek — can act like a brainstorming travel buddy. Tell it who is coming, where you're starting from, and how much you want to spend, and it returns a tidy day-by-day plan you can actually use.

Before you start

- A free DeepSeek account (sign up at deepseek.com with an email address — no payment needed for the chat feature).
- The DeepSeek mobile app, or any web browser on a phone or computer.
- About 15–20 minutes, mostly waiting for the AI to type.
- Optional but useful: a notes app on your phone (Apple Notes, Google Keep, etc.) so you can save the final plan.

Step 1 — Open DeepSeek and start a new chat

Open the DeepSeek app or go to **deepseek.com** in your browser. Sign in with the account you just created. You'll see a clean chat screen with a text box at the bottom — it's intentionally simple, looking similar to most messaging apps. Tap that text box once, and your keyboard slides up, ready for typing. If you've used DeepSeek before, look for a small **"New Chat"** icon (often a pencil or plus sign) at the top of the screen so you start with a blank conversation. Starting fresh matters — the AI will remember earlier messages in the same chat and may get confused by leftovers from old trips.

Example: just type the word **"Hi"** for now to make sure the box works.

You'll know it worked when a response starts streaming in within a few seconds.

Step 2 — Tell DeepSeek the basics of your trip

In the chat box, write a short paragraph that names the essentials: who's going, where you're starting from, the weekend dates, and what kind of trip you want (relaxing, active, food-focused, somewhere new). The more concrete you are here, the better the plan will be. Press **Enter** or tap the send arrow. DeepSeek usually replies in a few seconds — it does not browse the live web for this kind of question, so it won't know current hotel prices, but it will use its general knowledge of cities, neighborhoods and family activities. Think of this message like giving a friend the brief before they make suggestions.

Example: *"We're a family of four — two adults, a 7-year-old and a 10-year-old — based in Chicago. We want to drive somewhere within 3 hours for the weekend of 14–15 June. We love outdoorsy stuff and kid-friendly museums. Budget is medium."*

You'll know it worked when DeepSeek writes back with one or two clarifying questions or an early outline of ideas.

Step 3 — Ask for a day-by-day itinerary

Now ask for the structured plan you actually came for. Be specific that you want a **day-by-day** schedule, **kid-friendly** activities, and **rough travel times** between stops. Press send and let it write. You'll see the answer appear paragraph by paragraph — the response can be quite long for this kind of request, so give it a moment to finish. If the answer is heading the wrong direction (too many hikes, only adult restaurants), you don't need to start over — just tell it what to change in the next step. If you ever want a shorter summary, type "summarize that" and it will condense the plan into a few lines.

Example: *"Please make a day-by-day itinerary for Saturday and Sunday. For each day, give me morning, midday and afternoon activities that work for kids aged 7 and 10. Add rough drive times between stops."*

You'll know it worked when you see a clearly labeled "Saturday" and "Sunday" section with specific times or blocks.

Step 4 — Add meals, budget notes and a backup plan

The first plan is usually solid but a bit bare on food and money. Send a follow-up message in the same chat asking for restaurant suggestions that welcome kids, an estimated cost for the day, and a simple rainy-day backup. DeepSeek keeps the whole conversation in mind, so it builds on what it just suggested. Watch for the reply to come in the same chat window — no need to refresh or start a new thread. If it suggests a fancy tasting-menu place, just say "make meals more casual, kid menus only" and it will revise. This step turns a generic plan into one that fits your actual family instead of an imagined one.

Example: *"Add lunch and dinner suggestions for each day — nothing too fancy, kids need to see familiar food on the menu. Estimate the total daily cost for our family of four, and give me a backup indoor plan in case it rains."*

You'll know it worked when the new reply includes named restaurants or food types, a rough dollar figure per day, and an alternative activity for each day.

Step 5 — Refine the plan with focused follow-ups

This is where DeepSeek really earns its keep: you can keep tweaking the same plan without starting over. Ask narrow questions like "is there a shorter drive option?" or "swap the museum for something more outdoorsy." Press send after each question and read the reply before sending the next one — the chat history stays short, so you don't have to scroll far. If the assistant suggests something you don't recognize (a town name, a small park, a museum), quickly search it on your phone to confirm it's real and open on your dates. AI sometimes invents place names that sound convincing but don't actually exist, especially for lesser-known spots.

Example: *"Swap the Sunday morning museum for an outdoor activity — the kids get restless indoors. Also, can you suggest a town within 2 hours of Chicago instead of 3?"*

You'll know it worked when the revised plan reflects every change you asked for, in the same chat thread, with no contradictions from earlier suggestions.

Step 6 — Save the plan outside the chat

Once the plan looks good, get it out of the chat window — old chats can disappear after app updates, and you don't want to lose your weekend to a glitch. Tap and hold inside the answer (on mobile) or highlight it (on desktop) and choose **Copy**. Then paste the whole thing into a notes app, a Google Doc, or your family's group chat. If the plan is very long, copy it in two parts to avoid truncation. Some versions of DeepSeek also show a small **share icon** near a message — that creates a link you can send to your partner. Either way, the goal is simple: by the end of this step, the plan lives somewhere you can read it on Saturday morning without opening the AI again.

Example: copy the final plan, then paste it into a new note titled **"Family weekend — 14/15 June"**.

You'll know it worked when you can open your notes app and read the full itinerary without going back to DeepSeek.

Step 7 — Check the key details before you go

DeepSeek can't book anything, see live prices, or know if a museum is closed for renovations — so the last step is yours. Spend ten minutes verifying the big things: opening hours for any ticketed attraction, the restaurant is still serving on those dates, drive times on Google Maps, and the weather forecast. If something is closed or sold out, ask DeepSeek in the same chat for a swap ("the aquarium is closed Sunday, what else works for kids that age?"). This quick check is the difference between a smooth weekend and showing up at a locked gate with tired kids.

Example: type *"Quick check — can you remind me of the backup plan for Sunday if it rains? I want it in one short list."*

You'll know it worked when you have a confirmed plan in your notes, real opening hours in your calendar, and a fallback ready in case anything changes on the day.

Common mistakes

- **Asking a question that's too vague.** "Plan me a trip" gets a generic answer that doesn't fit your family. Always include the family size, ages of kids, starting city, dates and budget in the first message.
- **Trusting every place name without checking.** DeepSeek knows a lot, but it can invent small museums or restaurants that sound real. Google each venue before you drive there, especially for places you've never heard of.
- **Forgetting to mention budget.** Without a budget signal, the plan may lean toward expensive tasting menus and pricey hotels. Say "low", "medium" or "high" up front, and add a number if you have one in mind — it changes the suggestions a lot.

What this means for you

- **In everyday life:** the same prompt style works for any weekend — birthday parties, day trips, visiting friends in another city. Once you have the rhythm, planning takes minutes instead of evenings.
- **For your business or work:** if you ever run team off-sites or client visits, you can paste the same brief and ask for a more grown-up version — swap "kids" for "teammates", "budget" for "company card", and the structure stays useful.
- **If you're just getting started:** the easiest first move is to ask DeepSeek for a one-day plan for *this* coming weekend, even if you don't take it. The practice of writing a clear brief is the actual skill — the trip is a bonus.

Try it now

Open DeepSeek, paste the exact prompt from Step 2 with your own city and dates, and see what it comes back with in under a minute. That's the whole skill: one good paragraph of context, and a useful plan follows.

Create a soothing bedtime story routine with Alexa

7 steps · Read online

Imagine tucking your child into bed and simply saying, “Alexa, start bedtime story.” In just a few minutes, you can set up a gentle routine that Alexa will read aloud each night, complete with a soothing voice and soft background sounds. Once it’s saved, you can use it every evening with a single voice command, creating a lovely ritual for your family.

Before you start

- An **Amazon account** (free).
- The **Alexa app** installed on your phone or tablet (free).
- A **device with Alexa** (like an Echo Dot, Echo Show, or similar).
- Access to a **bedtime story skill or content** (many free options are available; some specific premium content may require a paid subscription like Audible or Amazon Music Unlimited).
- Roughly **10-15 minutes** to set up.

Step 1 — Open the Alexa app and find Routines

First, open the **Alexa app** on your phone or tablet. When it loads, you’ll typically see the home screen with various cards and information. Look at the bottom of the screen for a navigation bar with several icons. Tap the icon that looks like **three horizontal lines** or a **square 'More' button** in the bottom-right corner. This will open a menu. On the next screen, find and tap the word **Routines**. Routines are like your personal scripts for Alexa, telling her to perform a series of actions based on a trigger.

If it looks different: If you don't see "More" or three lines, look for a 'Settings' gear icon, or a menu icon that looks like nine dots or a hamburger menu. The goal is to find the main menu where 'Routines' is listed.

Example: Tap `More` → `Routines`.

You'll know it worked when you see a screen titled "Routines" with a large "Create Routine" or "Add a Routine" button.

Step 2 — Create and name your routine

Now you're on the Routines screen, tap the **+ Create Routine** or **Add a Routine** button. The app will then present you with options to build your new routine. The first thing to do is give it a simple, memorable name. Look for a field labeled "Enter routine name" or "Routine name." This name is just for you to keep your routines organized.

If it looks different: The button might be a plus (+) icon in the top-right corner, or simply "New Routine." Tap it to begin.

Example: Type `Bedtime Story` into the name field.

You'll know it worked when your chosen name appears in the designated field, and you see options to set "When this happens" and "Add action."

Step 3 — Set the voice command (trigger)

Next, you need to tell Alexa what phrase you'll say to start this routine. Under the section labeled **When this happens**, tap the **+** or **Add** button. From the list of trigger options, tap **Voice**. A new screen will appear where you can type your preferred phrase. This phrase is your "voice trigger" – the specific command Alexa will listen for to activate your routine. Keep it simple and easy to remember for everyone who might use it.

If it looks different: Instead of "Voice," you might see an option like "Phrase" or "Custom command." Select that to type your phrase.

Example: Type `start bedtime story`.

You'll know it worked when your chosen phrase is displayed under "When this happens" on the routine creation screen.

Step 4 — Add the bedtime story action

Now it’s time to add the actual story. Under the **Add action** section, tap the **+** or **Add** button. You’ll see a list of actions Alexa can perform. Since we want a story, we’ll choose a content-playing option. Tap on **Skills** or **Music & Audio**. If you choose Skills, you can search for a specific bedtime story skill (many are free). If you choose Music & Audio, you can type "bedtime story" and Alexa will find suitable content from your linked services or free options. Select your preferred story source.

If it looks different: Look for categories like "Play," "Entertainment," or "Listen" if "Skills" or "Music & Audio" aren't immediately visible.

Example: Tap `Skills` → search for `bedtime stories` → select a skill like `Kids Bedtime Stories` and tap `Add`.

You'll know it worked when the selected story skill or music command appears in the list under "Add action."

Step 5 — Customize Alexa's voice for the routine

You can give your bedtime story a calm, consistent voice. While still on the routine setup screen, scroll down and look for an option called **Alexa's Voice** (or sometimes "Speaking Style" or just "Voice"). Tap on this to see the available voices. Choose one that sounds soothing, like "Matthew" or "Emma," or another voice if your app offers more options. This setting ensures Alexa uses this particular voice for all the speaking actions within this routine.

If it looks different: The option might be labeled "Choose Alexa's voice" or "Voice Settings" within the routine.

Example: Tap `Alexa's Voice` and choose `Matthew`.

You'll know it worked when the name of the voice you selected is displayed next to the "Alexa's Voice" option in the routine.

Step 6 — Add a calming background sound (optional)

To make the storytime even cosier, you can add a gentle background sound that plays after the story or simultaneously. Tap **Add action** again. This time, navigate to **Music & Audio** (or sometimes "Sleep Sounds" or "Nature Sounds"). Search for a calming sound like "rain sounds," "ocean waves," or "soft piano music." Select your desired sound and confirm. Remember that this will typically play *after* the story action in the routine unless you adjust the order.

If it looks different: Some background sounds might be listed directly under "Sleep Sounds" as a separate action type.

Example: Tap `Music & Audio` → search for `rain sounds` → select a track and tap `Add`.

You'll know it worked when the background sound action appears below your story action in the routine's list of actions.

Step 7 — Save and test your routine

Once you've added all your desired actions and settings, it's time to save your routine. Look for the **Save** button, which is often found in the top-right corner of the screen. Tap it to finalise your new routine. After saving, the routine will appear in your list of active Alexa Routines. To test it, go to any Alexa-enabled device and speak the voice command you set up earlier.

If it looks different: The "Save" button might be a tick mark (✓) or a "Done" button. If Alexa doesn't respond when you test, double-check your trigger phrase for spelling or pronunciation in the Alexa app.

Example: Say: `Alexa, start bedtime story`.

You'll know it worked when Alexa begins speaking the bedtime story in the voice you chose, potentially followed by the selected background sound.

Common mistakes

- **Mistake:** Alexa doesn't respond when you say the trigger phrase. **Fix:** Go back into the Alexa app, find your "Bedtime Story" routine, and check the spelling of your trigger phrase under "When this happens." Ensure it's exactly what you're saying, and try speaking it slowly and clearly to your Alexa device.
- **Mistake:** Alexa plays the wrong story or just says "playing bedtime story" instead of reading one. **Fix:** The "Alexa Says" action is for Alexa to speak text, not to play content. Go back to your routine, delete the incorrect action, and add a new one by tapping "Add action" → "Skills" or "Music & Audio" to select a proper story skill or music.
- **Mistake:** The background music is too loud or not playing at all. **Fix:** If the sound isn't playing, ensure it's added as a separate "Music & Audio" action *after* your story. If it's too loud, you can add another action "Device Settings" → "Volume" to set a specific volume for your Alexa device as part of the routine.

What this means for you

- **In everyday life:** You can now start a calming storytime routine for your child with a single voice command, freeing up your hands and giving you a few extra minutes to get ready for bed yourself, or simply relax while your child enjoys a custom story.
- **For your business or work:** While this specific routine is personal, the ability to create custom, multi-step actions with a single command can be adapted for work tasks, like a "start my workday" routine that plays news, adjusts lights, and reads your calendar.
- **If you're just getting started:** This is an excellent way to turn your Alexa device into a helpful, customized assistant, making a daily task smoother without needing any complex technical skills.

Try it now

Open the Alexa app, go to Routines, and create a new routine named "Bedtime Story." Set the trigger phrase to "start bedtime story," add a "Skills" or "Music & Audio" action to play a bedtime story, choose a calm Alexa voice for the routine, and then save it. Say the phrase aloud to any Alexa device to test it out.

How to Plan a Kids' Birthday Party Without Losing Your Mind

7 steps · Read online

> **In short:** A simple 7-step plan — date, theme, budget, invites, food and a day-of run sheet — turns a kids' birthday party into a calm afternoon instead of a last-minute scramble. Set the budget early and most of the stress disappears.

You can throw a party your kid remembers without throwing yourself into chaos. This guide gives you a calm, step-by-step plan you can start today — even if the birthday is weeks away. It won't stop every sugar crash or shy-kid meltdown, and it shouldn't try to. The job of a good plan is to make the day easier for *you*, so you can be present for the moments that matter.

Before you start

- A calendar — phone or paper, either works.
- A rough idea of how many kids your child actually wants there (not "everyone in class").
- About 90 minutes total, spread over a week — most of this work happens in 10-minute chunks.
- No special app needed. A notes app, a napkin, or a piece of paper is enough.

Step 1 — Lock in the date, time and guest count

Open your calendar and pick a date plus a 2-hour window, like **Saturday 2:00–4:00pm**. Then ask your child which 5–10 friends they really want there — the honest list, not the polite "everyone in class" answer. Write the number down. These three things — date, time window, and headcount — are the foundation everything else sits on. If your calendar app looks different, a notebook works just as well; the tool doesn't matter, the decision does. Example: "Saturday, June 14, 2:00–4:00pm. Guests: 8 kids + 2 siblings." You'll know it worked when you can say the date, time and number of kids out loud without checking.

Step 2 — Pick a theme and set a budget

Ask your child one question: *"What do you want this party to feel like?"* Pirate ship, dance party, dinosaur dig — anything works. Write the theme in one short phrase. Then set a budget: a single number for the whole party, split into three buckets — food, decorations/activities, and a small "surprises" cushion. Most stress at kids' parties comes from spending that wasn't planned, not from the kids being picky. Example: "Theme: dinosaur dig. Budget total $200. Food $80, decor + activities $90, surprises $30." You'll know it worked when you have one clear theme phrase and one number for the total spend.

Step 3 — Build a simple countdown checklist

Open a notes app, a piece of paper, or a spreadsheet — whatever you already use. Create four buckets: **4 weeks out, 2 weeks out, 1 week out, day before**. Now drop the obvious tasks into each: book the venue (if not home), send invites, buy decorations, plan food, buy or bake the cake, prep goody bags. Don't try to make it pretty. A messy list you'll actually use beats a tidy one you won't. If a spreadsheet feels like overkill, four sticky notes on the fridge work fine. Example: "4 weeks: send invites. 2 weeks: buy decor + cake ingredients. 1 week: confirm RSVPs. Day before: set up." You'll know it worked when you can glance at the list and see exactly what to do today.

Step 4 — Send the invites and track the RSVPs

Send invites at least 2–3 weeks ahead — paper ones in mailboxes, a text to the parent group, or a free template from a basic invite app. Each invite should include the date, time, address, theme (so parents know what to expect), and an RSVP date about a week before the party. As replies come in, write each name next to yes, no, or "didn't reply yet." This list is your real headcount — the one you actually use to buy food and a cake. Example: "Hi! Maya's turning 6, dinosaur-dig party, Saturday June 14, 2–4pm at our place. Please RSVP by June 7." You'll know it worked when you have a clear yes-list 5 days before the party, not the morning of.

Step 5 — Plan food, cake and one anchor activity

Keep the food simple: one main, one snack, one drink, and the cake. Kids don't need variety; they need predictability. The **anchor activity** is the one thing kids will remember — a treasure hunt, a piñata, a dance freeze game — pick one, plan the supplies, and stop there. Two extra "just in case" activities (a coloring sheet, a movie queued up) cover you if the anchor ends early. If you can't decide between two activities, flip a coin — indecision takes more energy than a small mistake. Example: "Food: mini sandwiches + grapes + juice + cake. Anchor: backyard dinosaur dig with toy fossils." You'll know it worked when you can list the food in one breath and name the main activity without looking.

Step 6 — Write a day-of run sheet (the calm-maker)

The night before, write a single page with five time blocks: **setup** (1 hour before), **guest arrival**, **food/cake**, **activity**, and **goodbye**. Under each block, write the 2–3 things that MUST happen — not a wish list, only the must-haves. Example: "2:00 guests arrive. 2:15 snack. 2:45 dinosaur dig. 3:30 cake. 3:50 goody bags. 4:00 pickup." Tape it to the fridge. When something goes sideways (it will), the run sheet is your reset

button — you read the next block and keep moving. 💬 Example: "1:00 setup. 2:00 arrive. 2:15 snack. 2:45 dig. 3:30 cake. 3:50 bags. 4:00 done." You'll know it worked when you can read it in 30 seconds and know what comes next.

Step 7 — Day-of: do less, breathe more

On the day, give yourself one rule: **do less than planned**. Skip one decoration, shorten the activity by 10 minutes, and stop trying to photograph everything. Kids don't measure the party by how perfect it looked; they measure it by how relaxed you were. If something flops, no kid has ever gone home saying "the piñata was too short." After the last guest leaves, write down one thing that worked and one you'd change — that's gold for next year. You'll know it worked when the kids are smiling, the cake is eaten, and you actually enjoyed it.

Common mistakes

- **Skipping the budget.** Picking a theme and adding up the receipts later usually ends in a $50 surprise at the register. Fix: pick the total number in Step 2 and split it into three buckets *before* you shop.
- **Inviting "the whole class."** 25+ kids triples food cost, triples activity planning, and triples your stress. Fix: let your child pick 6–10 real friends, then add 2–3 siblings if needed.
- **Planning too many activities.** Three scheduled games means none of them actually get finished, and you spend the party refereeing transitions. Fix: one anchor activity + one backup, and let the rest be free play.
- **No run sheet.** Without one, the day turns into a blur of "what now?" moments and a hangry child at 3pm. Fix: spend 10 minutes the night before writing the five time blocks from Step 6.

What this means for you

- **In everyday life:** The shape of this plan — date, scope, budget, list, day-of script — works for playdates, holidays, and family trips too. Same bones, different costume.
- **For your work or small business:** Project planning uses the same skeleton: scope, budget, timeline, run sheet. The party is a low-stakes way to practice the muscle you'll use on bigger projects.
- **If you're just getting started:** Do Steps 1, 2, and 3 on a napkin tonight. Don't open a spreadsheet, don't download a template, don't overthink it. The messy first draft is the one that actually saves you on the day.

Try it now

Open your phone's notes app, write your child's next birthday date at the top, and list the 5–10 names of friends they'd actually want there. That's your foundation — everything else slots in on top of it, one step at a time.

How to use Qwen Chat to plan a weekend road trip in the US without getting lost

7 steps · Read online

By the end of this guide you'll have a complete weekend road-trip plan—routes, stops, times, and even backup ideas—built with Qwen Chat in about 20 minutes. No prior travel-planning experience needed.

Before you start

- A free Qwen Chat account (sign-up takes 2 minutes at chat.qwen.ai)
- Any device with a web browser (phone, tablet or computer)
- Rough idea of where you want to go (e.g., "Pacific Coast Highway" or "Blue Ridge Mountains")
- Total time: about 20 minutes

Step 1 — Open Qwen Chat and start a fresh chat

Open your browser and go to chat.qwen.ai. You'll see a clean white page with a big "Ask Qwen anything" box at the bottom. Tap or click inside that box to begin typing. A new chat window opens automatically. You'll know it worked when you see the blinking cursor ready for your first message. Example: `Hi`

Step 2 — Give Qwen the big picture of your trip

In the same message box type a short summary of what you want: where you're starting, where you're headed, how many days, and who's coming. This helps Qwen tailor every detail to your exact needs. You'll see Qwen reply with a friendly welcome and a quick question to confirm the basics. If the reply looks different on your screen, look for a small "Welcome" banner at the top of the chat. Example: `I want a weekend road trip from Portland to the Pacific Coast Highway with my partner, leaving Friday after work and returning Sunday evening. Can you plan every stop, time, and backup idea so we don't get lost?`

You'll know it worked when Qwen replies with a clear plan outline and asks for any must-see spots you already have in mind.

Step 3 — Ask Qwen to build the full day-by-day itinerary

Copy the itinerary request Qwen just gave you and paste it into the same chat, then add any personal must-sees. Qwen will now create a detailed hour-by-hour plan with driving times, lunch spots, scenic lookouts, and even fuel stops. On screen you'll see Qwen typing for about 10–15 seconds, then a long list appears with every stop numbered and timed. If the list looks different (some Qwen versions show cards instead of text), look for a "Download" or "Copy" button near the top right of the reply. Example: `Here's what I'd like to keep: the Twelve Apostles at sunset, a short rainforest walk at Maits Rest, and a seafood dinner in Apollo Bay. Please add those and any other must-sees, plus backup ideas if a stop is closed or crowded.`

You'll know it worked when you see a numbered list with Friday evening, Saturday full day, and Sunday morning—each with times, addresses, and backup notes.

Step 4 — Request turn-by-turn directions for each leg

Ask Qwen to turn every driving leg into simple turn-by-turn directions you can read aloud to a passenger or save to your phone. Qwen will list each leg separately with street names, kilometers, and estimated minutes. On your screen you'll see a new bullet list under the itinerary. If the directions look too short, look for a "Show full route" link or button inside the reply. If it's missing, just ask Qwen again. Example: `Please give me turn-by-turn directions for every driving leg, including street names and kilometres, so my partner can navigate while I drive.`

You'll know it worked when each leg shows something like "Drive south on Princes Hwy for 12 km, turn left onto Marengo Rd, continue 3 km to Apollo Bay Visitor Centre."

Step 5 — Ask for backup ideas in case something changes

Tell Qwen you want backup stops for every main attraction in case the weather turns or a venue closes. Qwen will add a second set of options with shorter walks, indoor museums, or different viewpoints. On screen you'll see the original list followed by a new "Backup ideas" section. If the section is missing, look for a "More suggestions" button or simply ask Qwen again. Example: `Add a backup plan for every main stop in case it's raining or too crowded. Include indoor alternatives and shorter walks.`

You'll know it worked when you see a clear "Backup ideas" heading with 2–3 options for each main stop.

Step 6 — Request a printable or shareable summary

Ask Qwen to create a single clean summary you can save as a PDF, print, or send to your travel buddy. Qwen will combine the itinerary, directions, and backups into one easy-to-read page. On screen you'll see a new message with a heading like "Your Weekend Road-Trip Plan – Ready to Save." If the text is too long, look for a "Copy all" or "Download PDF" button at the bottom of the reply. 💬 Example: `Please give me one clean summary I can save or print, with every stop, time, address, and backup idea in one place.`

You'll know it worked when you see a single block of text with a clear title and everything in order.

Step 7 — Double-check the plan with one final question

Ask Qwen to review the whole plan for anything you might have missed—toll roads, low-fuel areas, or time buffers. Qwen will point out any gaps and suggest tweaks. On screen you'll see a short checklist at the end of the chat. If the checklist is missing, simply ask Qwen again. 💬 Example: `Please review the entire plan for anything I might have missed—toll roads, fuel stops, or time buffers—and suggest any tweaks.`

You'll know it worked when you see a green checkmark list with no red flags.

Common mistakes

- **Mistake:** Giving Qwen a vague request like "Plan a trip." **Fix:** Always include start, end, days, and who's coming so Qwen can tailor every detail.
- **Mistake:** Skipping the backup step and assuming everything will be open. **Fix:** Ask Qwen for backup ideas right after the main itinerary—rain or closure can ruin a perfect day.
- **Mistake:** Trusting the plan without checking the turn-by-turn directions. **Fix:** Always ask Qwen for street-by-street directions and save them to your phone before you leave.

What this means for you

- **In everyday life:** Next time you're dreaming of a spontaneous weekend escape, open Qwen Chat, paste your rough idea, and in 20 minutes you'll have a complete, stress-free plan you can trust—no travel-agent fees, no last-minute scrambling.
- **For your business or work:** If you organize team off-sites or client retreats, use Qwen Chat to draft itineraries, directions, and backup plans overnight, so you can focus on the event itself instead of logistics.
- **If you're just getting started:** This is the easiest first move—open Qwen Chat, type a simple request, and watch Qwen build a full weekend plan for you. You'll feel instantly capable and ready for your next adventure.

Try it now

Open chat.qwen.ai, paste the example prompt from Step 2, and hit send. In under a minute you'll have a complete weekend road-trip plan you can trust.

Turn messy phone photos into a family album with Google Photos AI

7 steps · Read online

Imagine opening your phone's gallery and seeing every photo of your child's birthday, your last holiday, or your dog neatly grouped together — no endless scrolling, no guessing. With Google Photos' free built-in AI, your phone's photos can be automatically sorted by the people in them, so you can find what you need in seconds.

Before you start

- A phone or tablet with the **Google Photos app** already installed (Android or iPhone).
- At least a few dozen family photos already saved on that device.
- A free Google account (the one you use for Gmail or YouTube).
- Rough total time: **7–10 minutes**.

Step 1 — Open Google Photos and sign in

Unlock your phone and find the **Google Photos** app — it looks like a pinwheel made of red, yellow, green, and blue. Tap it once to open. If you see a prompt asking you to sign in, enter the email and password you use for your Google account. If you're already signed in, you'll see your username in the top-right corner. You'll know it worked when you see your photos appear in a grid.

Example: Open the Google Photos app → tap **Sign in** → type `your.name@gmail.com` and your password → tap **Next**.

Step 2 — Turn on the People group feature

At the bottom of the screen you'll see five icons: Photos, Search, Albums, Utilities, and Library. Tap **Search**. On the Search screen, look for a heading like **People & pets** or **Groups**. If you don't see it yet, scroll down a little. Tap the heading once. Google Photos will show a short animation and then list any people it already recognizes. You'll know it worked when you see a list of names or faces under **People**.

Example: Tap **Search** → scroll to **People & pets** → tap the heading → wait for the list to appear.

Step 3 — Wait for Google Photos to finish scanning

If you're doing this for the first time, Google Photos needs a few minutes to scan all your photos and find faces. While it works, you'll see a small spinning circle in the top-right corner of the screen. You can close the app and come back later; the scan continues in the background. You'll know it worked when the spinning circle disappears and the list under **People** is fully loaded.

Example: Leave the app open or close it — either way, come back in 2–3 minutes.

Step 4 — Name the groups you see

Under **People**, you'll see headings like **Unknown 1**, **Unknown 2**, and so on. Tap one of these headings to see the photos inside. To give it a real name, tap the heading once, then tap **Name** or the pencil icon. Type the person's name (for example, "Mom", "Liam", "Sarah") and tap **Save**. You'll know it worked when the heading changes from "Unknown 1" to the name you typed.

Example: Tap **Unknown 1** → tap the pencil icon → type `Dad` → tap **Save**.

Step 5 — Create a family album from the groups

Once all the people are named, tap **Library** at the bottom of the screen, then **Utilities**. Under **Create new**, tap **Album**. Give the album a name like "Family 2025" and tap **Create**. Google Photos will open a new empty album. Tap **Add photos** → choose **From your groups** → select the people you want (for example, "Mom", "Dad", "Liam"). Tap **Done**. You'll know it worked when you see all the photos of your family together in one album.

Example: Tap **Library** → **Utilities** → **Album** → name it `Family Summer 2025` → tap **Add photos** → choose **From your groups** → select "Mom", "Dad", "Liam" → tap **Done**.

Step 6 — Share the album with family in one tap

With the album still open, tap **Share** in the top-right corner. Choose how you want to send it: via text message, email, or any other app you like. The album link will open instantly for anyone who clicks it, and they'll see the same neatly grouped photos. You'll know it worked when you see the share sheet appear and the album link is ready to send.

Example: Tap **Share** → choose **Messages** → pick "Mom" from your contacts → tap the send arrow.

Step 7 — Let Google Photos keep the album updated

To make sure new photos are automatically added to the right groups, open the Google Photos app every few days and tap **Search** → **People**. If you see new "Unknown" groups, tap and name them as in Step 4. Google Photos will continue to group new photos by person without any extra effort from you. You'll know it worked when new photos of the same person appear under their correct name in the **People** list.

Example: Open Google Photos → tap **Search** → tap **People** → name any new "Unknown" groups → close the app.

Common mistakes

- **Mistake:** You don't see the **People & pets** heading after tapping Search. **Fix:** Make sure you have at least a few dozen photos on the device. If you still don't see it, tap the three-dot menu in the top-right corner and choose **Refresh**.
- **Mistake:** The names you give people disappear after you close the app. **Fix:** Always tap **Save** after typing a name. Your changes are saved to your Google account, so they'll be there next time you open the app.
- **Mistake:** The app runs slowly on an older phone. **Fix:** Close other apps while Google Photos is scanning. If it's still slow, try scanning smaller folders first (for example, just the last six months of photos).

What this means for you

- **In everyday life:** No more scrolling through hundreds of photos to find pictures of your child's birthday or your last holiday. One tap opens an album with every photo of your family together, ready to share with relatives.
- **For your business or work:** Use the same People groups to organize project photos, team events, or client meetings by person. Share the album link instead of messy email attachments.
- **If you're just getting started:** Start with a small folder of your favorite family photos. Once you see how easy it is, you can let Google Photos scan your entire collection.

Try it now

Open the Google Photos app on your phone, tap **Search**, and look for **People & pets**. If you see any "Unknown" groups, tap and name them right away. You'll have your first family album in under two minutes.

Organize your favorite recipes with Perplexity Q&A

7 steps · Read online

Imagine having every recipe you love—from Grandma's scones to your go-to curry—stored in one place you can search by ingredient, diet, or mood. With Perplexity's Q&A feature, you can turn your scattered notes into a living cookbook that answers your questions instantly. No apps, no logins, just one chat.

Before you start

- A **free Perplexity account** (sign up at perplexity.ai).
- A **device with a web browser** (phone, tablet, or computer).
- A list of 3–5 recipes you want to save (handwritten, photos, or links).
- **10 minutes** of focused time.

Step 1 — Create a new Perplexity space for your recipes

Open your browser and go to perplexity.ai. Sign in with your free account. On the left sidebar, look for a button like **"+ New Space"** (it might be under a "Spaces" heading or a plus icon). Tap it. A new blank space will open with a chat box at the bottom. This space will act like a private recipe folder. You'll know it worked when you see a clean, empty chat box and the space name at the top is editable.

Example: Type: `My Recipe Book`

Step 2 — Paste your first recipe into the chat

In the chat box, paste the full text of one recipe. Include the title, ingredients, and instructions. If your recipe is in a photo, describe it clearly in words. After you press Enter, Perplexity will show a summary card of the recipe at the top of the chat. It may also ask a follow-up question like "Would you like me to organize this?" You'll know it worked when you see the recipe summary appear above the chat and the assistant asks a helpful follow-up.

Example:

```
Pumpkin Soup
Ingredients: 1 kg pumpkin, 1 onion, 2 carrots, 400ml coconut milk, 1 tsp cumin
Instructions: Peel and chop veggies. Sauté onion. Add pumpkin, spices, and water. Simmer 20 mins. Blend. Stir in coconut milk.
```

Step 3 — Ask Perplexity to store the recipe in your space

Below the recipe summary, type: "Save this as a recipe in My Recipe Book." Perplexity will confirm by showing a message like "Added to your space: Pumpkin Soup." Your space now contains a growing list of recipes. If you don't see the confirmation, look for a small "Saved" badge or check the space sidebar. You'll know it worked when the recipe appears in the space sidebar or you see a "Saved" confirmation.

Example: Save this as a recipe in My Recipe Book.

Step 4 — Ask Perplexity to find recipes by ingredient

In the same space, type a question like: "Show me all recipes with coconut milk." Perplexity will scan your saved recipes and return a list with the matching ones highlighted. If you don't see results, check that you used the exact ingredient name (e.g., "coconut milk" not "coconutcream"). You'll know it worked when you see a clean list of recipe titles with the matching ingredient in bold.

Example: Show me all recipes with coconut milk.

Step 5 — Ask Perplexity to suggest meals based on what you have

Type: "What can I cook tonight with pumpkin and coconut milk?" Perplexity will analyze your saved recipes and suggest the Pumpkin Soup you added earlier. If the answer is off, rephrase with more detail: "Only recipes I've saved in My Recipe Book." You'll know it worked when you see a relevant recipe suggestion with a clear reason why it fits.

Example: What can I cook tonight with pumpkin and coconut milk?

Step 6 — Share your recipe space with a friend

At the top of your space, look for a button like **“Share”** or an icon that looks like a box with an arrow. Tap it and copy the share link. Send the link to a friend via message or email. They’ll open it in their browser and see your recipes without needing an account. If you don’t see a share button, look under a menu like “…” or “More.” You'll know it worked when your friend confirms they can view your recipes in their browser.

Example: https://perplexity.ai/space/12345-my-recipe-book

Step 7 — Ask Perplexity to scale up your collection

Now that your space is set up, add more recipes the same way: paste the text, ask Perplexity to save it, and tag it with keywords like “vegan”, “quick”, or “weeknight”. Over time, your space becomes a smart cookbook that answers questions like “What’s a quick vegan dinner?” or “What uses leftover pumpkin?” You'll know it worked when you can ask natural questions and get accurate recipe suggestions from your growing collection.

Example: What’s a quick vegan dinner I can make?

Common mistakes

- **Pasting messy text**: If the recipe summary looks wrong, clean up the pasted text and try again. Perplexity works best with clear, plain language.
- **Forgetting to save**: Always say “Save this as a recipe in My Recipe Book” after pasting. Without this step, your recipes won’t be searchable.
- **Using vague names**: Name your space something clear like “My Recipe Book” instead of “Stuff”. It helps Perplexity organize your recipes better.

What this means for you

- **In everyday life**: Keep your family’s favorite recipes in one place you can search by ingredient or diet. No more scrolling through old notebooks or photos—just ask Perplexity and get answers instantly.
- **For your business or work**: If you run a small cooking class or food blog, use a shared Perplexity space to organize recipes by theme or season. Clients can ask for meal ideas and get instant, accurate answers.
- **If you're just getting started**: Start with three recipes. Once you see how easy it is to search and share, add more over time. The key is consistency—small steps build a powerful cookbook.

Try it now

Open Perplexity in your browser, create a space called “My Recipe Book”, paste one recipe, and ask Perplexity to save it. You’ll have a living cookbook in under two minutes.

Create a personalized weekly meal plan based on your diet with Gemini

6 steps · Read online

Ever feel stuck trying to figure out what to cook for the week, especially when you have specific dietary needs or a few ingredients you want to use up? This guide will show you how to use Gemini, an AI assistant, to effortlessly create custom weekly meal plans. By the end, you'll be able to generate meal ideas, shopping lists, and even quick recipes tailored just for you.

Before you start

- **A Google Account**: You'll need one to access Gemini. It's free to create if you don't already have one.
- **Access to Gemini**: You can use Gemini through its website (gemini.google.com) on a computer, tablet, or smartphone.
- **Rough total time**: Allow about 10-15 minutes to generate and refine your first personalized meal plan.

Step 1 — Open Gemini and start a new chat

To begin, you need to navigate to the Gemini website. Open your preferred web browser and type `gemini.google.com` into the address bar, then press Enter. Once the page loads, you'll typically see the Gemini home screen, which features a prominent chat input field at the bottom and a "New chat" button, often found in the top-left corner or similar navigation area.

WHAT YOU SEE: The Gemini welcome screen or your most recent conversation. **WHAT HAPPENS NEXT:** The page loads, ready for you to begin interacting with the AI assistant. **IF IT LOOKS DIFFERENT:** If you're on a mobile device, the layout might be more compact; look for a clear text input box and a button or icon (like a plus sign `+`) to start a new conversation. If it automatically loads a previous chat, that's fine for now, we'll start a fresh one in the next step. **You'll know it worked when:** You see the Gemini interface and a noticeable text input area where you can type.

Step 2 — Initiate a fresh conversation

Now that you're in Gemini, it's a good idea to start with a blank slate for your meal planning request. Look for a "New chat" button or a similar option to begin a fresh conversation. This button is usually found in the sidebar on a desktop computer, or sometimes as a plus sign `+` icon on mobile devices. Clicking this will clear any previous conversation history, giving you a clean space for your meal plan.

WHAT YOU SEE: The Gemini interface with a button clearly labeled "New chat" or a similar icon. **WHAT HAPPENS NEXT:** The chat area will clear, removing any past messages and preparing for your new input. This helps Gemini focus solely on your current request without being influenced by previous topics. **IF IT LOOKS DIFFERENT:** On some versions or devices, it might be a button like "Start new conversation" or a simple `+` icon. Just tap whatever initiates a fresh interaction. **You'll know it worked when:** The main chat window becomes completely empty, presenting a clean message box where you can type your first instruction to Gemini.

Step 3 — Describe your meal plan needs to Gemini

This is where you tell Gemini exactly what you're looking for in your meal plan. In the message input box at the bottom of the screen, type out your request, including all your preferences. Think about dietary restrictions (like vegetarian, gluten-free), preferred ingredients you have on hand, how much time you have for cooking, and for how many days you need the plan. The more details you give, the better and more tailored Gemini's response will be.

WHAT YOU SEE: An empty message box ready for your input. **WHAT HAPPENS NEXT:** As you type, your message appears in the box. Once you send it (usually by pressing Enter or clicking a send icon), Gemini, your AI assistant (a powerful computer program designed to understand and generate human-like text, much like talking to a very smart, helpful person), will process your request and begin to generate a response. **IF IT LOOKS DIFFERENT:** If Gemini takes a moment to respond, that's normal; it's working to create a detailed plan based on your input. Just be patient. **Example:** `I need a 7-day vegetarian meal plan for one person. I prefer easy recipes that take less than 30 minutes to prepare each night. I have lots of lentils, spinach, and sweet potatoes that I want to use.` **You'll know it worked when:** You see Gemini's reply appearing in the chat window, starting to suggest meal ideas based on the information you provided.

Step 4 — Refine the meal plan with specific feedback

After Gemini gives you its first meal plan suggestion, take a moment to read through it. Chances are, it's a good start, but you might want to make a few tweaks. Don't hesitate to ask for changes or additions in the same chat. For example, you can swap out meals, add more side dishes, or ask for specific types of cuisine. This back-and-forth is how you get the most out of an AI assistant, refining the output until it's just right.

WHAT YOU SEE: Gemini's initial meal plan, usually presented as a list of meals for each day. **WHAT HAPPENS NEXT:** Gemini will process your new instruction and adjust its previous response. It's like having a conversation with a chef who's happy to modify the menu to your taste. The AI will often regenerate the entire plan or highlight the changes it has made. **IF IT LOOKS DIFFERENT:** If the initial plan is already perfect for you, you can skip this step and move on to asking for more details, like a shopping list. **Example:** `This looks great, but could you swap out the dinner on Tuesday for something with chickpeas instead of lentils? Also, can you add a simple snack idea for each day, like a`

`piece of fruit or some nuts?` **You'll know it worked when:** Gemini presents an updated meal plan that clearly incorporates all your new requests, showing the changes you asked for.

Step 5 — Request practical details like a shopping list

Once you're happy with the meal plan itself, you can ask Gemini to help with the practical side of preparing these meals. A very useful next step is to ask for a shopping list. This saves you the time of going through each recipe and jotting down ingredients. You can even ask for the list to be organized in a specific way, like by grocery aisle, to make your shopping trip more efficient.

WHAT YOU SEE: Your refined, personalized weekly meal plan in the chat window. **WHAT HAPPENS NEXT:** Gemini will quickly generate a comprehensive shopping list, pulling all the necessary ingredients from the meals it has suggested. It will typically present this as a clear, easy-to-read list directly in the chat. **IF IT LOOKS DIFFERENT:** If you prefer, you could also ask for estimated nutritional information or a rough breakdown of cooking times for each meal. **Example:** `Can you create a shopping list for this entire week's meal plan? Please group the items by type, such as 'Produce,' 'Pantry,' 'Dairy,' and 'Frozen.'` **You'll know it worked when:** A clearly organized shopping list, categorised as you requested, appears directly below your meal plan in the chat.

Step 6 — Get quick recipe overviews or variations

To make your meal plan even more useful, you can ask Gemini for quick recipe overviews for any specific meal you're unsure about, or request variations if you want to try something a little different. This is especially helpful if you're new to a particular dish or want to ensure you have all the basic steps before you start cooking. Gemini won't give you a full cookbook, but it can provide enough detail to get you started.

WHAT YOU SEE: Your complete meal plan and shopping list. **WHAT HAPPENS NEXT:** Gemini will provide a brief summary of how to prepare the requested dish, or offer ingredient substitutions and alternative cooking methods right there in the chat. It acts as a quick reference guide, saving you from searching online. **IF IT LOOKS DIFFERENT:** If you're confident with cooking and familiar with the suggested meals, you might not need to ask for this extra detail. **Example:** `For the 'Sweet Potato and Lentil Curry' on Thursday, could you give me a very quick overview of the main steps? Also, what's a good alternative if I don't have fresh ginger on hand?` **You'll know it worked when:** You see a short, easy-to-follow explanation of the recipe or practical alternative ingredient ideas for your chosen meal appearing in the chat.

Common mistakes

- **Being too vague in your initial request**: If you just say "Give me a meal plan," Gemini might provide a very generic one. **Fix:** Be specific from the start. Mention dietary needs (e.g., "vegan," "low-carb"), time constraints (e.g., "quick 20-minute meals"), and ingredients you want to use.
- **Forgetting to specify dietary needs**: You might get suggestions that don't fit your diet if you don't mention them upfront or early in the conversation. **Fix:** Always include your core dietary requirements in your first prompt, such as "I need a gluten-free plan" or "I am trying to eat more plant-based meals."
- **Not asking for enough detail**: Gemini can do more than just list meal names. If you only get a list, you might miss out on other helpful features. **Fix:** Once you have the basic plan, follow up with requests for a shopping list, approximate cooking times, or even quick recipe outlines for specific dishes.

What this means for you

- **In everyday life:** This skill means you can save a lot of time and mental effort each week. No more staring into the fridge wondering what to cook! It helps reduce food waste by planning around ingredients you already have, and makes sticking to dietary goals much easier and more enjoyable.
- **For your business or work:** While meal planning isn't a business task, the ability to prompt an AI for structured, personalized plans is highly transferable. You could apply this same approach to generate content calendars, project outlines, daily task schedules, or even event planning checklists, making your work more efficient and organized.
- **If you're just getting started:** Don't try to create a complex 7-day plan straight away. Instead, go to Gemini and ask for a simple 3-day meal plan based on one specific ingredient you love and already have at home, like "Give me a 3-day dinner plan using chicken and broccoli."

Try it now

Head over to Gemini (gemini.google.com) and ask for a 5-day meal plan tailored to your favorite vegetable, specifying if you have any dietary preferences like "gluten-free" or "vegetarian." You'll be surprised how quickly you get a useful response!

How to write heartfelt birthday cards, wedding toasts and thank-you notes with Copilot

6 steps · Read online

> **In short:** You can ask Microsoft Copilot to draft a birthday card, wedding toast or thank-you note from a few details about the person and the moment, then rewrite the draft in your own voice. Copilot gives you a starting point so you are not staring at a blank page — but it cannot know the person for you, so your edits are what make the message feel real.

Anyone can freeze up when a card or toast is due. The right words hide somewhere between your feelings and the page, and the clock is ticking. Microsoft Copilot is an AI assistant that lives in your browser, your phone, and inside apps like Word and Outlook. You give it a short description of who the message is for and what you want to say, and it writes a first draft in seconds. Your job is to read it, keep the parts that sound like you, and cut the parts that do not. The result is a message that feels personal without taking all evening to write.

Before you start

- A free Microsoft account (the same one you use for Outlook, OneDrive, or Xbox). If you do not have one, you can create one in a couple of minutes at the official Microsoft sign-in page.
- A way to reach Copilot: a web browser on a computer, the Copilot app on your phone, or Microsoft Word with Copilot enabled. Any of these works.
- A few honest details about the person you are writing to — their name, what you appreciate about them, a memory you share, the tone you want (warm, funny, formal, brief).
- About 10 to 15 minutes for the whole thing: 2 minutes to set up, 3 to 5 minutes to prompt and edit, and a few more to copy the final version into your card or print it out.

Step 1 — Open Copilot and start a new chat

In a browser, go to the official Copilot website and sign in with your Microsoft account. On a phone, open the Copilot app and tap the sign-in button. If you are in Microsoft Word, look for the **Copilot** option in the top ribbon. Once you are signed in, you land on a clean screen with a large text box near the bottom — that is the prompt box, the place where you type what you want. A new chat starts every time you load the page, so old messages will not confuse your new request. If you have used Copilot before and want a clean start, look for a button like a **new chat** or **pencil icon** at the top of the list of past conversations.

Example: type `Help me write a birthday card for my sister.`

You'll know it worked when the screen shows your words in the prompt box and Copilot is ready to respond.

Step 2 — Tell Copilot who the message is for and what you want to say

Click inside the prompt box and type a short description. A prompt is just the instruction you give to an AI. The more specific you are about the person and your feelings, the warmer the draft will be. Mention the recipient's name, your relationship, one thing you genuinely appreciate about them, and the tone you want. You do not need to write a lot — three or four sentences is plenty. If you are writing a wedding toast, also mention how long the toast should run when spoken out loud, usually one to two minutes. If you are writing a thank-you note, mention the specific gift or favor you are thankful for, because generic thanks can feel hollow.

Example: `Write a warm birthday card for my dad, who is turning 70. He's a quiet, kind man who taught me to fish and always fixes things around the house. I want it to feel grateful and a little funny, not too long. About 80 words.`

You'll know it worked when Copilot replies with a short block of text that sounds like a real card — a greeting, a personal line, and a closing.

Step 3 — Ask for a second version if the first one is not quite right

Read the draft out loud. If it sounds stiff, too formal, or just not like you, do not start over. Type a short follow-up asking for a change. Saying "make it warmer", "sound more like me talking to a close friend", or "add a small joke about our fishing trips" tells Copilot exactly how to adjust. This back-and-forth is normal — even short messages often need a second pass. Keep your edits as instructions, not as the full rewrite, so the AI does the heavy lifting and you stay in control of the feeling. If you would like a few options side by side, ask for two or three versions in the same message and pick the one you like best.

Example: `That is good, but can you make it sound less formal? Add a line about the time he fell asleep in his fishing chair and the fish got away. Keep it under 100 words.`

You'll know it worked when the new draft has at least one line that genuinely sounds like something you would actually say.

Step 4 — Add one true detail the AI cannot know

This is the step that turns a generic draft into a personal message. Think of one small true thing about the person or your relationship that you did not put in the prompt — a shared joke, a recent moment, a small habit you love. Type that detail into the chat and ask Copilot to weave it in. It might be the way your dad hums while he works, the inside joke from your sister's wedding, or the specific way a friend helped you last month. The AI cannot invent this on its own, which is exactly why your input is the most important part of the process. Once the detail is in, the message stops sounding like it could be for anyone and starts sounding like it is for one person.

💬 Example: `Can you add a line about how he still hums the same tune from the radio while he works in the garage? It's something I will always remember.`

You'll know it worked when the draft contains the small true detail in a way that feels natural, not forced.

Step 5 — Edit the draft into your own voice

Read the latest version one more time and put on your editing hat. Cross out any phrase that does not sound like you, even if it is well written. Swap big words for small ones. Shorten sentences that run too long. Add a word you actually use, like "honestly" or "you know". If you are writing a toast, read it out loud once — your ear will catch anything that sounds like a press release instead of a real person talking. Copilot is your first-draft helper, not your ghostwriter, so the final version should pass the "would I actually say this?" test. A good rule of thumb is that at least half the words should feel familiar to you.

💬 Example: change *"I am eternally grateful for your unwavering guidance"* to *"Thanks for always being there, Dad. I notice it more every year."*

You'll know it worked when you can read the whole message without cringing at any single line.

Step 6 — Copy the final version into your card, email or printout

Once the message sounds right, click the **copy** button that usually appears under Copilot's reply, or highlight the text and copy it yourself. Paste it into your card, your email, or a Word document you want to print. If the handwriting on the card matters to you, write the message by hand onto the card, using the draft as your script. If you used the Copilot add-on inside Word, the text is already in your document, so you can format and print straight away. For toasts, practice reading it once or twice before the event so the words feel natural in your mouth, not just on the page.

💬 Example: paste the draft into a blank card, print it on nice paper, or copy it into the notes app on your phone for the big day.

You'll know it worked when the message is sitting where it needs to be and you feel ready to send, write, or raise a glass.

Common mistakes

- **Forgetting to add a real detail.** A clean draft with no personal touch reads like a greeting card from a store. Fix: always add at least one small true thing only you would know, even if it is just a phrase you and the recipient share.
- **Sending the AI's first draft untouched.** The first reply is rarely the best one. Fix: treat it like a rough sketch and spend two minutes editing so it sounds like you, not like a polite robot.
- **Asking for a toast that is way too long.** A two-minute spoken toast is about 200 to 250 words, shorter than most people think. Fix: tell Copilot the time limit and ask for a version that fits, then read it out loud to check the pace.
- **Using the wrong tone for the moment.** A formal thank-you note to your boss and a funny birthday card for your best friend need very different voices. Fix: name the tone you want in the prompt, and if the draft misses, say so and ask for a rewrite.

What this means for you

- **In everyday life:** The next time a birthday, wedding, baby shower, or thank-you slips up on you, you can go from blank page to a draft in under five minutes, and finish the personal part in another five. That is roughly a ten-minute job instead of an evening of dread.
- **For your business or work:** The same prompt-and-edit habit works for client thank-you notes, retirement messages for a colleague, or a short team toast at a work event. You bring the real detail about the person, Copilot helps you shape the words.
- **If you are just getting started:** Pick the next card or note you actually need to write this week. Open Copilot, paste the example prompt from Step 2 with your own details swapped in, and treat the reply as a starting point, not a final answer.

Try it now

Open Copilot on your phone or computer, and write a real birthday card for someone whose birthday is coming up in the next month. Use the Step 2 example as your template, swap in one true memory of your own, and spend two minutes editing the draft until it sounds like you. Send the card, and notice how much lighter a small task feels when you do not have to start from zero.

Plan a 3-day weekend road trip with Meta AI in 15 minutes

6 steps · Read online

> **In short:** Meta AI (the free AI assistant built into Facebook, Instagram, and WhatsApp, or on meta.ai) can draft a complete 3-day road-trip itinerary in under a minute. You provide the starting city and what you want — lakes, good food, nothing touristy — and it fills in the rest. It cannot book hotels or reservations for you.

You've got a free weekend, a car, and a vague urge to get out of town — but the thought of hours of Googling stops, routes, and restaurant hours makes your eyes glaze over. Here's the good news: Meta AI can turn a single sentence like "I'm starting in Denver, I want mountains, good tacos, and scenic drives" into a full three-day schedule with driving times and meal ideas. It won't buy your hotel or reserve your table, but it will do the brainstorming heavy lifting so you're not staring at a blank spreadsheet at 10 pm on a Thursday. Here's how to use it step by step.

Before you start

- **Access to Meta AI:** You need a Facebook, Instagram, or WhatsApp account (all free), or you can go directly to **meta.ai** in any browser — no app download needed.
- **No paid plan required:** Meta AI is completely free to use.
- **Time needed:** About 15 minutes total — a few minutes to type your request, a minute to read the result, and a few more to ask for changes.
- **What you'll need in your head:** Your starting city, roughly how many days, and a one-sentence feeling about what you want (peaceful? food-forward? scenic drives?).

Step 1 — Open Meta AI

The easiest way is to go to **meta.ai** in any web browser on your phone, tablet, or computer. The page opens straight to a chat box — no account creation needed if you're already signed in through Facebook or Instagram in that browser.

If you prefer your phone, open **Facebook** or **Instagram**, tap the search bar at the top, and look for a small **Meta AI icon** (a blue circle with sparkles, often near the search field). Tap it and the chat opens inside the app.

On **WhatsApp**, open any chat, tap the **+ (attachment) button** next to the text field, then tap **Meta AI** from the menu that slides up. A chat thread with Meta AI opens in the same app.

If you can't find a Meta AI icon in the search bar, look inside the **Chats** section for a dedicated "Meta AI" conversation, or check for an update in your app store — Meta AI's placement changes as the apps are updated.

You'll know it worked when a chat window appears with a blue-and-white Meta AI logo at the top, ready for you to type.

Step 2 — Write your first prompt

Think of a prompt as a short email to a very well-read travel-buddy who plans fast. You don't need to be detailed — Meta AI can work with surprisingly little. Include your starting city, roughly how many days you have, and one or two words about the vibe you're after. The vaguer you are, the more general the result — so a few specifics help.

Keep your prompt to one or two sentences. You can always ask for changes after you see the first draft. Don't worry about being polite or formal — Meta AI doesn't need a greeting; it just reads the instruction and starts planning.

Example: *"I'm starting in Asheville, North Carolina, and I have three days. I want scenic mountain drives, good local food — nothing chain restaurants — and mostly quiet, not too touristy spots. Can you give me a day-by-day itinerary with approximate driving times between stops?"*

You'll know it worked when a detailed reply appears below your message — usually a paragraph or a formatted list with Day 1, Day 2, and Day 3 headings, each with stops and timing.

Step 3 — Read the first draft and check it

Meta AI will usually respond with a formatted plan. It tends to include a suggested route order, a few specific stops (restaurant names, viewpoints, small towns), and estimated driving time between each point. Look it over and ask yourself: does it mention my starting city? Are the days roughly the right length? Are there any stops that sound wrong or too touristy?

Here's a real thing to know: Meta AI can occasionally suggest a restaurant or attraction that sounds plausible but has closed or moved. Treat the names it gives you as a starting point, then double-check any specific place you're excited about with a quick Google search — it takes 20 seconds and saves frustration.

If the plan looks roughly right but a day feels too packed or too empty, move on to Step 4.

You'll know it worked when you've scanned the whole itinerary, made a quick mental note of one or two places you want to verify later, and picked the one thing you'd most like to change.

Step 4 — Ask for one specific change

Most people need at least one adjustment. Maybe the driving times feel too tight, or you realize you don't want to backtrack the way Meta AI suggested, or a day sounds great but you wish it included a specific type of stop (a farmers' market, a scenic overlook, a coffee shop). That's fine — just type a follow-up request as a normal sentence.

Meta AI can adjust one section at a time, so you don't need to re-explain everything. Just ask for the one thing that bothered you most. If the first change worked, you can ask for the next one in a new message — that keeps each request small and easy for the assistant to handle.

Example: *"Day 2 looks good but I don't want to drive more than 2 hours between stops — can you tighten it up?"*

Or: *"Can you swap out the afternoon on Day 1 for something outdoorsy instead of the museum?"*

You'll know it worked when the revised plan appears and the specific issue you mentioned has been addressed — for example, the driving times are shorter, or a new stop has replaced the one you flagged.

Step 5 — Save or screenshot the final plan

Meta AI doesn't have a built-in "save my trip" button, so a screenshot is the quickest way to keep the plan. On a phone, take a screenshot (usually power + volume-up at the same time). On a computer, you can screenshot the chat window or copy the text and paste it into notes, email, or a Google Doc.

Some people copy the text into a notes app and then add links to specific places they want to verify — that becomes your trip plan that you can actually use on the road.

If your phone or browser doesn't take screenshots easily, try the "share" button inside the chat (look for an arrow or three dots) — many versions now let you copy the full conversation in one tap, which is even faster than screenshotting piece by piece.

You'll know it worked when you have the itinerary saved somewhere you can open during your drive (notes app, email to yourself, or a screenshot you can pinch-to-zoom).

Step 6 — Check one or two key stops before you go

Before you leave, open Google Maps or Google Search and type in the name of any restaurant or attraction Meta AI suggested. Look for two things: is it still open, and are the hours listed on a recent review? This takes about two minutes per stop and means you're not pulling into a closed parking lot on Saturday morning.

Meta AI cannot look up real-time hours, read current reviews, or confirm that a place hasn't moved — it draws on information up to a certain point, not today's conditions. That's why this quick check matters.

Example: search *"Blue Ridge Bakery Asheville hours"* to see if it opens early enough for your planned breakfast stop.

You'll know it worked when you can confirm at least your top two or three stops are open and reachable on the days you planned.

Common mistakes

- **Vague prompts that give you a generic result.** "I want a fun road trip" might get you a list of famous national parks across the country — exciting, but not a 3-day weekend plan. Add your starting city and a rough direction (e.g., "starting from Austin, within a 4-hour drive") to keep it grounded.
- **Accepting the first draft without checking.** Meta AI sometimes suggests restaurants or attractions that sound plausible but have closed. A 20-second Google search on the names you care about most prevents a wasted detour.
- **Asking for too much at once.** If you pile on six changes in one message ("make it cheaper, add hiking, remove the museum, find a dog-friendly hotel"), Meta AI may lose the thread and miss one or two. Ask for one change, read the result, then ask for the next — it's faster in the end.

What this means for you

- **In everyday life:** Instead of spending an evening doom-scrolling travel blogs, you can spend 15 minutes with Meta AI drafting a real plan, then spend 5 minutes confirming the stops you care about. That's an evening back in your pocket.
- **For your business or work:** If you occasionally need to plan a short work retreat, a team offsite, or a client site visit across a few days, the same workflow works — just swap the vibe description for "not too formal" or "close to the airport."
- **If you're just getting started:** Your very first prompt can be as simple as "I'm starting in [your city]. Give me a 3-day road trip itinerary. I like [one or two things]." That's it. See what comes back, then ask for one change. You'll learn how the tool thinks by watching what it

gives you.

Try it now

Open **meta.ai** in your browser right now, or tap Meta AI inside Facebook, Instagram, or WhatsApp. Type in your starting city, say you have a free weekend, and describe one thing you want from the trip — a type of scenery, a type of food, a pace (busy or relaxed). Hit send and watch what appears. It takes less than a minute, and you'll have a real plan to build from.

Translate a foreign-language recipe or instruction manual with Hugging Face's free translator

6 steps · Read online

> **In short:** Hugging Face runs free, open-source translation models in your browser through its "Translation" page — paste a foreign recipe or manual on the left, pick your language pair, and an English version appears on the right. No signup, no payment, but it isn't perfect for every language and very long texts need to be split into smaller chunks.

Picking up a friend's grandmother's paella recipe — written only in Spanish — used to mean a lot of squinting at a phone dictionary. Hugging Face, the open-source AI platform, has a quieter but surprisingly handy tool: a free translator that lives right on its website. No app to install, no account to make, no payment. You paste the text, pick a model, and read the English on the other side. Here's how to use it in about three minutes.

Before you start

- A web browser (Chrome, Firefox, Safari, or Edge — anything modern).
- A computer, tablet, or phone with internet.
- The text you want translated — copy it from a website, a PDF, or type it out.
- About 3–5 minutes total.

Step 1 — Open the Hugging Face translation page

In your browser's address bar, type `huggingface.co/tasks/translation` and press Enter. A page loads with the title "Translation" and a short explanation of what the tool does, plus a demo box in the middle of the screen. The demo has two text areas — one on the left for the text you want translated, and one on the right where the English (or your chosen language) will appear.

If that exact URL doesn't work on your device, search "Hugging Face translation" in Google, or go to huggingface.co first, then look for a "Tasks" link in the top menu and choose "Translation" from the list. Sites get reorganized now and then, so a small detour is normal.

Example: type `huggingface.co/tasks/translation` into the address bar.

You'll know it worked when you see a page titled "Translation" with two text boxes sitting side by side.

Step 2 — Choose your language pair

Look just above or below the text boxes for two dropdown menus — one labeled something like "Source language" (the language your text is written in) and one labeled "Target language" (the language you want). The default is often English-to-something, so you'll likely need to change the source.

Click the source-language dropdown and scroll until you find the language of your text — Spanish, French, Japanese, German, Korean, and many more are listed. Then make sure the target language is set to English. If you don't see your language listed, that's the cue for Step 3.

If your screen shows only one dropdown, or the language names appear as codes like "es" or "fr" instead of full names, that's fine — the codes follow the international standard: es is Spanish, fr is French, de is German, ja is Japanese, zh is Chinese, ko is Korean, pt is Portuguese, it is Italian, ru is Russian.

Example: set the source dropdown to "Spanish" and the target dropdown to "English."

You'll know it worked when both dropdowns show the pair you want — for example, "Spanish → English" — and the text boxes are still empty and ready for you.

Step 3 — Pick a different model if your language is missing

Below the language dropdowns, you may see another selector labeled "Model" with a default choice already filled in (often something starting with "Helsinki-NLP" or "facebook/nllb"). The default usually works well for common European languages. If your language wasn't in the dropdowns, or the default gives you poor results, click this dropdown and pick a different one.

The most beginner-friendly choice for rare or non-European languages is "facebook/nllb-200-distilled-600M" — it's been trained on around 200 languages, so it usually covers whatever you throw at it. For common pairs like Spanish-English or French-English, the lighter "Helsinki-NLP/opus-mt" models tend to be quick and clear.

If you only see one dropdown and no model selector at all, skip this step — the page has already picked a sensible default for the language pair you chose in Step 2.

Example: if your dropdown only had ten languages and you need Thai or Vietnamese, switch the model to "facebook/nllb-200-distilled-600M" and the language list grows.

You'll know it worked when a model name appears in a small label like "Model:" or in a gray bar above the boxes, and both text boxes are still visible.

Step 4 — Paste your foreign-language text

Click inside the left text box — the one labeled something like "Input" or sitting under the source language. The cursor blinks. Now grab the text you want translated. If it's on a website, select it with your mouse, right-click and choose Copy, then come back here and right-click inside the box and choose Paste. If it's in a PDF or on paper, you may need to type it out, or use your phone's camera to capture it first.

For best results, paste a reasonable chunk at a time — about one paragraph or one recipe step. Most of these models work best on one to five sentences; longer blocks can confuse them and slow the page down. If your text is very long, break it into pieces and translate one piece at a time.

If the translation feels slow or never appears, it usually means the server is busy — wait ten seconds and try again, or refresh the page and paste a shorter piece.

Example: paste the Spanish line "Mezcle el arroz con el caldo y deje cocer a fuego lento durante 20 minutos."

You'll know it worked when the pasted text shows up in the left box, with accents and special characters preserved exactly as in the original.

Step 5 — Click translate (or watch it auto-translate)

Some pages have a button labeled "Translate," "Compute," or a small arrow icon between the boxes. Click it once. Other pages translate automatically the moment you stop typing or pasting — there is no button at all. Either way, give it a few seconds.

The right box — usually labeled "Output" or sitting under the target language — fills in with the English version. Below or beside it, you may see a small note like "Generated in 1.2 seconds" or the name of the model that did the work.

If you see an error message instead of a translation, the most common cause is a language pair the current model doesn't support — go back to Step 3 and pick a different model, or check that the source language matches what you actually pasted (don't paste Spanish if the dropdown still says French).

Example: after clicking Translate, the right box should show something like "Mix the rice with the broth and let it simmer for 20 minutes."

You'll know it worked when the right box shows readable English that matches the meaning of what you pasted — even if a word or two come out a little off.

Step 6 — Copy the result and use it

Click inside the right box and drag your mouse across the English text to highlight it all. Press Ctrl+C on Windows or Cmd+C on Mac to copy. Now open a Notes app, an email, or a recipe site and paste (Ctrl+V / Cmd+V) the translation where you need it.

For a long recipe or manual, repeat Steps 4 and 5 for each section. Keep the original and the English side by side in your notes — that way you can compare and spot anything that looks wrong.

If a sentence comes out awkward, try rephrasing the original slightly before pasting it back. Small wording changes in the source language often give the model a clearer path to a good English translation.

Example: paste the English result into a note titled "Paella — translated" alongside the original Spanish so you can double-check the measurements.

You'll know it worked when you have a clean English version of your recipe or manual in your notes, ready to read.

Common mistakes

- **Pasting the whole article at once.** Most of these models handle one to five sentences at a time well. Split longer text into paragraphs, recipe steps, or bullet points before pasting.
- **Trusting one translation blindly.** Models can mistranslate measurements, ingredients, or safety warnings. Cross-check anything important — oven temperature, cooking time, chemical dosage, electrical warnings — against a second source.
- **Forgetting to switch the source language.** If you paste Spanish while the source dropdown still says French, the output will be nonsense. Always glance at the dropdown before clicking Translate.

What this means for you

- **In everyday life:** Translate a recipe card from a trip, a sewing pattern, a manual for a secondhand appliance bought abroad, or a letter from a relative overseas — in seconds, without paying for an app.

- **For your business or work:** Translate short snippets of customer emails, product reviews, or instructions from international suppliers. For full contracts or official documents, you still want a professional human translator — but for quick understanding, this tool is plenty.
- **If you're just getting started:** Pick one short recipe or instruction sheet you've been curious about, paste a single paragraph in, and see what comes back. Two minutes, no signup, and you'll know exactly what the tool can and can't do for you.

Try it now

Open `huggingface.co/tasks/translation` in your browser right now. Find a recipe or a sentence in any language you don't speak — try a Spanish cooking site or a Japanese craft blog — paste one paragraph, and read the English on the other side. That is the whole tool, and you just used it.

Build a Simple Weekly Meal Plan With Grok Based on What's in Your Fridge

6 steps · Read online

> **In short:** Open Grok, list what's actually in your fridge, and ask for a 7-day dinner plan with a shopping list. You'll get dinner ideas matched to your real ingredients in one chat — no paid plan or cooking skill required.

By the end of this guide, you'll have a real 7-day dinner plan built around what's already sitting in your fridge, plus a shopping list for what you're missing. This is for anyone who keeps opening the fridge at 6pm and feeling stuck — no fancy equipment, just a free Grok account and about ten minutes.

One honest note up front: Grok can draft the plan and the list, but it can't go shopping for you, and it can't tell if your chicken is actually still good. That part is still on you.

Before you start

- A free Grok account (sign up through the Grok website or the X app — both work)
- Three minutes to peek inside your fridge, freezer, and pantry
- Optional: a notes app open so you can paste the final plan into it
- Total time: about 10 minutes, mostly for the food peek and saving

Step 1 — Open Grok and start a fresh chat

Open Grok on your phone or computer. On a computer, go to the Grok website and sign in. On a phone, open the Grok app — or open X and look for a small ⚡ icon somewhere in the menu. Once you're signed in, you'll see a chat box at the bottom of the screen with a placeholder like "Ask Grok anything." A fresh chat keeps the assistant focused on your meal plan instead of mixing it up with old conversations.

If you don't see a chat box right away, look for a button labeled "New chat" or a + icon in a corner. The Grok entry point inside X moves around as the app updates, so try the menu or the search bar if it's not where you expect.

💬 Example: type "Hi" to test the connection, or jump straight to Step 2 with your list ready.

You'll know it worked when you see an empty chat box and a blinking cursor, waiting for your first message.

Step 2 — Take quick stock of your fridge and pantry

Before opening Grok, spend two or three minutes looking through your fridge, freezer, and pantry. Write down what you actually have — not what you wish you had. Include basics like rice, pasta, canned beans, and any half-used vegetables. Note proteins too: chicken, eggs, tofu, ground meat, or fish. Spices and condiments matter as well, because they change what meals are possible.

Group items roughly by type — vegetables, proteins, grains, dairy, pantry staples — so the list is easy to scan. Also jot down how many people you're feeding. Grok uses the serving size to scale recipes up or down.

💬 Example: "Fridge: 2 chicken breasts, half a head of broccoli, 4 eggs, a wedge of cheddar. Pantry: rice, canned chickpeas, soy sauce, olive oil, garlic. Freezer: frozen peas. Serving 2 adults."

You'll know it worked when you have a written list of 8–20 ingredients plus a serving size, ready to paste into the chat.

Step 3 — Paste your list and ask for a 7-day plan

Paste your ingredient list into the Grok chat box and add a clear request. Be specific about the timeframe, the number of meals, and any restrictions. Grok answers best when you spell out exactly what you want — for example, "dinners for 7 nights, no seafood, weeknight meals under 30 minutes, vegetarian on Fridays." More context means a closer first answer.

After you hit send, wait a few seconds. Grok usually replies with a day-by-day breakdown: meal name, short description, and a rough cooking method. The first answer is rarely perfect — that's normal, not a failure.

If you don't see a typed reply within 15 seconds, check your connection and try again. The chat often shows a small "thinking" indicator before the answer appears.

💬 Example: "Here's what I have: [paste your list]. Build me a 7-day dinner plan for 2 adults. No seafood. Keep weeknight meals under 30 minutes. Use what I already have first, then suggest what to buy."

You'll know it worked when Grok returns a numbered list of seven meals, each tied to ingredients from your fridge.

Step 4 — Adjust the plan to your taste and schedule

Look over Grok's first answer. Anything you don't like? Speak up — that's the easiest part of using a chatbot. Reply in plain language and tell it what to swap. Maybe Tuesday's curry doesn't fit your weeknight energy. Maybe Saturday feels too involved for a relaxed weekend. Ask for replacements, and Grok will rewrite just those days.

This is also the moment to mention things you forgot: a picky eater, a slow cooker you want to use, or leftovers for lunch the next day. The more you iterate, the more the plan feels like yours rather than a generic template.

If a meal keeps coming back that you really don't want, say so directly. The chat remembers your earlier messages in the same conversation, so a clear "no mushrooms anywhere" once is usually enough.

💬 Example: "Swap Tuesday's curry for something quicker — under 20 minutes. Add a slow-cooker recipe for Sunday. And please skip mushrooms entirely; nobody here eats them."

You'll know it worked when the updated plan reflects your changes and you can imagine actually cooking each meal.

Step 5 — Ask Grok to build a shopping list of missing items

Once the plan feels right, ask Grok to compile the shopping list. Specifically, request only the items you don't already have — otherwise you'll end up buying things already sitting in your pantry. A good prompt names the servings and asks Grok to group the list by aisle or category, which makes the actual shopping trip faster.

Grok will usually return a checklist sorted into sections like produce, dairy, pantry, and frozen. Some answers even estimate quantities. If the list looks too long, ask it to shrink by reusing ingredients across two or three meals — that's both cheaper and faster to cook.

If you want a printable version, ask Grok to format it as a simple numbered list or a table — then copy and paste it into your notes app.

💬 Example: "Great. Now give me a shopping list of only what I'm missing, grouped by aisle (produce, dairy, pantry, frozen, meat). Quantities for 2 adults. Keep it under 25 items if you can."

You'll know it worked when you have a clean, grouped shopping list of just the items you don't already have at home.

Step 6 — Save the plan somewhere easy to find

Now save your meal plan and shopping list before the chat scrolls away. Open your phone's notes app, email, or any list tool you already use, and paste both lists in. Give the note a clear name like "Week of June 8 — Meal Plan" so you can find it next week.

If you're on a computer, you can print the shopping list straight from the browser. Most grocery apps also let you email yourself a list — copy from Grok, paste into the app, and check items off as you walk the aisles.

One more thing worth doing now: skim the plan once more for anything you actually dislike — easier to swap a meal now than at 6pm on a Wednesday. If you want to come back to this conversation later, look for a way to bookmark or pin it in the Grok interface (if available), so you can ask for a new plan in seconds next week.

💬 Example: "OK, I'm ready to save this. Can you format the final 7-day plan and the shopping list as plain text I can copy into my notes app?"

You'll know it worked when both lists are saved in your notes app.

Common mistakes

- **Asking Grok to invent ingredients you don't have.** If you say "what can I make?" without listing what's actually in your fridge, you'll get random recipes — not a plan that uses your food. Fix: always paste your real inventory first, then ask for ideas.
- **Accepting the first plan without iterating.** Grok's first answer is a draft, not a final menu. If a meal feels wrong, just say so. Fix: reply with swaps ("Tuesday's too spicy — give me something milder") until the plan matches your week.
- **Forgetting to mention dietary needs or picky eaters.** The assistant can only work with what you tell it. Fix: spell out allergies, vegetarians in the household, and "no fish" rules at the start, so the plan fits your family from the first reply.
- **Saving the chat but losing the shopping list.** It's easy to skim past the second reply. Fix: ask for the shopping list as a separate, grouped message and paste it into your notes app or grocery list right away.

What this means for you

- **In everyday life:** You spend less time staring at the fridge and less money on food that quietly goes bad. A good plan turns "what's for dinner?" from a daily decision into a quick glance at your phone.
- **For your work:** If you cook for a team, a book club, or a shared household, the same workflow scales — just bump the servings in your prompt and ask for batch-cooking tips. Grok can suggest what to make ahead on Sunday for the rest of the week.
- **If you're just getting started:** Don't aim for perfection on your first try. Make a list, send a quick prompt, and accept a rough plan you can fix later. Each week gets easier, and your fridge will start feeling a lot less random.

Try it now

Open your fridge right now, take a quick mental photo of what's inside, and paste this into Grok:

"Here's what's in my fridge and pantry: [your list]. Make me a 3-day dinner plan for [number] people. Use what I have first. No [food you dislike]."

See what comes back. If it sparks an idea, you've just made your week easier in under two minutes.

Plan a Budget-Friendly City Vacation in One Afternoon

6 steps · Read online

> **In short:** You can use any AI assistant — ChatGPT, Gemini, Claude, Microsoft Copilot, or Le Chat — to draft a realistic day-by-day city plan that fits a budget you set, in roughly an hour. The AI can't book flights, can't see today's prices, and can't know last-minute closures, so you'll confirm the practical details yourself before you go.

You don't need a travel agent, a giant spreadsheet, or three weekends of research. One afternoon with an AI assistant, a notebook (digital or paper), and a rough idea of where you want to go is enough to walk away with a plan you can actually use. This guide walks you through it step by step, with prompts you can copy and paste as you go.

Before you start

- **An AI assistant** — any one will do: ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or another you already use. A free plan works fine; a paid plan just gives longer answers. **Rough total time:** about 60–90 minutes, including the back-and-forth.
- **Two rough numbers** — your total budget (all-in, including flights) and the number of days you have. Ballpark is fine; you'll tighten them as you go.
- **One open notes app** — Google Docs, Notes, or a notebook, so you can save the good parts of each answer as you draft.

Step 1 — Tell the AI where, when, and how much

Open your AI assistant in a browser or app and start a fresh chat. You'll see a message box at the bottom of the screen, usually with a placeholder like "Message ChatGPT" or "Ask Gemini". If you don't see a blank chat, look for a **"New chat"** or + button first — you want a clean conversation so the AI doesn't mix your trip with old questions. In that box, type a clear one-paragraph brief: destination, number of days, month or season, who is traveling, and your total budget. Be specific about the kind of trip you want — food, walking, museums, neighborhoods, kid-friendly, quiet — because the first draft depends on it.

Example: *"I want to spend 4 days in Lisbon in mid-October. Just me, no car. My total budget including flights from the US East Coast is $1,800. I like walking, local food, neighborhoods, and one or two paid museums. I don't love big crowds. Give me a first-pass plan."*

You'll know it worked when you get back a clear day-by-day outline with morning, afternoon, and evening blocks, plus a short budget breakdown.

Step 2 — Ask for a money breakdown you can actually trust

Generic answers like "you can do Lisbon for $50 a day" are often too optimistic. In the same chat, reply and ask the AI to break the budget into categories: lodging, food, local transport, attractions, and a small "buffer" for surprises. Then ask which two categories usually eat most of the budget, and which is the easiest to cut without ruining the trip. If the assistant forgets the buffer, push it to add one — that single line is what keeps a budget plan from collapsing on day two.

Example: *"Thanks. Now break the $1,800 into lodging, food, transport within the city, attractions, and a buffer. Which two categories usually eat most of the budget for a 4-day solo trip, and which is the easiest to cut without ruining the trip?"*

You'll know it worked when the AI gives you specific categories with rough shares (for example: lodging 40%, food 30%) and names a clear "save here" lever — like staying just outside the center, eating lunch as the main meal, or using a city tourist card.

Step 3 — Build the day-by-day plan around free and cheap anchors

Now ask the AI to rebuild each day using mostly free or cheap anchors: public parks, walking neighborhoods, viewpoint walks, free museum days, markets, churches, street art, and self-guided routes. Tell it to mix in one paid "anchor" per day (a museum, a boat ride, a cooking class) and to group nearby stops so you don't waste hours on transport. If the first reply is just a list of sights, push back and ask for a route, not a list — that's the difference between a plan and a brochure.

Example: *"Rebuild the 4 days using mostly free or cheap anchors — parks, walking streets, viewpoints, free museum hours, markets. Put one paid anchor per day, and group nearby things together so I'm not crossing the city for every stop. Flag anything that's usually very crowded."*

You'll know it worked when each day reads like a real route, not a list. You should be able to picture yourself walking from one stop to the next, with a coffee break and a meal built in.

Step 4 — Pressure-test the plan against your real life

A pretty plan on paper is not the same as a plan you'll actually do. In the same chat, copy the day-by-day plan back in and ask the assistant to challenge it. Specifically: how much walking is that, do I need to pre-book anything, and what should I do if it rains? This is where the AI is most

useful as a calm second opinion, not as a source of fresh facts. If the assistant hedges ("it depends..."), ask for a worst-case number — that gets you a real answer.

Example: *"Here's the Day 2 plan you gave me: Alfama walking tour → Sé cathedral → lunch in Graça → National Tile Museum → sunset at Miradouro da Senhora do Monte → dinner in Bairro Alto. Push back on this: how much walking is that, do I need to book the museum, and what's a good rainy-day swap?"*

You'll know it worked when you get honest pushback — for example, "this is 8–10 km on hills, you'll be tired by evening" or "yes, book the Tile Museum, lines can be long." Treat those notes as a starting point, then sanity-check the critical ones with a quick search.

Step 5 — Lock in the food and the small print

The difference between a budget trip and a stressful one usually lives in two places: food, and the small daily decisions. Ask the AI for a short list of typical local dishes at different price points (a cheap street-food version and a sit-down version for each), and for a "don't do this" list — tourist-trap areas to avoid, common scams, and transport passes worth buying versus ones to skip. Keep the food list short and memorable; you only need 5–8 dishes to feel like you ate the city, not 50.

Example: *"Give me a short list of Lisbon dishes with a 'cheap-eats version' and a 'sit-down version' for each. Then give me 5 'don't do this' warnings — the kind of thing a first-timer always falls for, like taxi scams, tourist-menu restaurants, or transport passes that don't pay off for 4 days."*

You'll know it worked when you have a small, memorable food list and 4–5 specific "watch out for" notes you can actually remember. These two things alone save most budget travelers the most money.

Step 6 — Save it as one clean page and share it

Finally, ask the AI to compress the whole plan into a single plain-text page you can copy into your notes app or print: the budget breakdown, the day-by-day plan, the food list, the warnings, and a short packing note. Keep it under one screen if you can. Then, if you're traveling with someone, share it as a doc so they can mark what they actually want to do. The skill here is editing — almost no first draft survives being copied into a real trip without one pass of cutting.

Example: *"Now compress everything we discussed into one tight page: budget shares, 4-day plan, food list, top 5 warnings, and a 5-line packing note for October weather. Keep it tight — I want to fit this on one screen on my phone."*

You'll know it worked when you can scroll the whole plan in under a minute, and a friend reading it cold understands your trip without you having to explain it.

Common mistakes

- **Trusting the first plan as final.** The first draft is a starting point, not the answer. Always do one round of "push back on this" — it catches the unrealistic walking days and the missing reservations.
- **Skipping the buffer line.** A budget without a 10–15% buffer will not survive contact with reality. Add it before you start booking, not after the first surprise.
- **Asking the AI about current prices or openings.** AI assistants don't reliably know today's hotel rates, flight prices, or whether a small museum is closed for renovation. Confirm anything money- or time-sensitive on the official site or a recent review.

What this means for you

- **In everyday life:** The same one-afternoon method works for a long weekend, a family visit, or any trip where you have a budget and a vague idea and need a workable plan by tonight.
- **For your business or work:** If you ever plan client visits, team offsites, or conference travel on a per-diem, the "brief → money breakdown → day plan → pressure-test → compress" pattern is how a good assistant plans anything, and you can reuse it for non-travel projects too.
- **If you're just getting started:** The single biggest upgrade is Step 4 — asking the AI to challenge its own plan. Most beginners stop at Step 1 and assume a pretty plan is a good plan. It usually isn't.

Try it now

Pick one city you've been curious about, open any AI assistant, and paste this: *"I want to spend 3 days in [city] in [month]. My budget is [$X] all-in. I like [2 things you like]. Build me a first-pass day-by-day plan with mostly free or cheap activities."* Spend the next 15 minutes reacting to what you get — that back-and-forth is the whole skill.

Plan Your Weekly Meals and Shopping List with Any AI Assistant

6 steps · Read online

By the end of this guide, you'll have a five-day dinner plan and a ready-to-shop shopping list — written in your words, for your family. This works with any AI assistant (think ChatGPT, Gemini, Claude, Copilot, Le Chat, or DeepSeek), so use whichever one you already have an account with.

Before you start

- **An AI assistant** you can type into — a free account with ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, DeepSeek, or a similar tool is fine. Both the website and the mobile app work.
- **Roughly 10–15 minutes** to answer the AI's follow-up questions and tweak the plan.
- **A pen and paper, or your phone's notes app**, to copy the final list.
- **No paid plan is required** for this task on any of the assistants listed above. If you want to save the chat history, you may need a free sign-in.
- One honest heads-up: the AI can **draft** a plan and a list, but it can't taste your food, check your fridge, or place a grocery order for you. You'll do the cooking and the shopping yourself — the AI just makes the thinking part quicker.

Step 1 — Open your AI assistant and start a new chat

Open the website or app of the AI you use, and start a **fresh, empty chat**. Think of a chat like a blank page in a notebook — the AI only remembers what you write in *this* conversation, so starting clean keeps things tidy. You'll see a text box at the bottom (or middle) of the screen with a prompt like *"Message ChatGPT"* or *"Ask Gemini"*. If you don't see that box, look for a **+** or **New chat** button — usually at the top of a sidebar on the left, or behind a pencil icon.

Example: don't type anything yet — just make sure you can see the empty text box.

You'll know it worked when the screen shows a blank conversation with your name at the top and a blinking cursor in the message box.

Step 2 — Tell the AI about your household in one message

In that empty text box, **type one short message describing your family**. The more concrete you are, the better the plan you'll get back. Mention the number of people eating, any allergies or foods you avoid (gluten, nuts, dairy, pork, shellfish — whatever applies), and roughly how much you'd like to spend per week. If you have a slow cooker, an air fryer, or only one working hob, mention that too — it changes the suggestions.

Example: *"Family of four in Chicago: two adults, a hungry 14-year-old, and a 7-year-old who won't eat spicy food. No nuts or shellfish. One of us is vegetarian most nights. Budget around $120 for five dinners. We have an oven and a stovetop, no microwave."*

You'll know it worked when the AI replies by **summarizing back what you told it** and asking one or two clarifying questions. If it just says "Sure, what would you like?" — that's a sign your message was too short, so add a few more details and send it again.

Step 3 — Ask for the five-day dinner plan

Now **send a follow-up asking specifically for a five-night dinner plan**. Use the word "dinner" rather than "meal" so the AI doesn't accidentally add breakfast or lunch. Ask it to keep each recipe under 30 minutes, to repeat ingredients across nights where it makes sense (so you cook once and re-use), and to flag any night that needs you to **soak, marinate, or defrost something in advance**. If it looks different on your device and the AI offers bullet points, that's fine — just read them in order.

Example: *"Great, thanks. Please give me a five-night dinner plan based on that. Keep each dinner under 30 minutes of active cooking, repeat ingredients across nights where it makes sense (saves money and waste), and call out any night that needs overnight prep like defrosting or marinating."*

You'll know it worked when the AI responds with a numbered list of five dinners, each with a short ingredient list and a rough time estimate. If a recipe uses something obscure or expensive, the AI usually notes it — read that note before moving on.

Step 4 — Ask for the matching shopping list

Reply in the same chat and **ask the AI to turn the plan into a single shopping list, grouped by aisle**. Grouping by aisle (produce, meat, pantry, dairy, frozen) saves you real time in the supermarket because you walk through the shop once instead of zig-zagging. Also ask it to **include rough amounts** (e.g. "500g chicken mince", "1 brown onion") so you can compare prices, and to **note anything you probably already have** — olive oil, salt, pepper, plain flour. If your AI's reply comes back as a wall of text, just reply "Could you put that as a bullet list grouped by aisle?" and it will reformat.

Example: *"Nice plan. Can you turn that into one shopping list grouped by aisle (produce, meat/fridge, pantry, dairy, frozen) with rough amounts? And please mark staples I likely already have at home — oil, salt, pepper, basic spices — so I don't have to buy them."*

You'll know it worked when you see a clean, aisle-grouped list with amounts beside each item, and a short "you probably already have" note at the end. If anything is unclear (like "some herbs" instead of "a small bunch of parsley"), reply and ask for specifics — the AI is happy to be more precise.

Step 5 — Personalize and stress-test the plan

Read through the list with a **"will my family actually eat this?"** eye. This is the step most people skip, and it's where a plan either succeeds or ends up half-eaten. Reply with any swaps you want — *"swap the salmon for chicken thighs"*, *"the kids won't eat couscous, can we have rice instead?"* — and ask the AI to **re-output the updated shopping list**. This is also a good moment to ask it to check for **duplicate ingredients** across the week (buying one bunch of coriander and using it three ways is cheaper than buying three).

Example: *"Looks good but please swap the salmon for chicken thighs (cheaper this week), change the couscous to jasmine rice, and double-check I'm not buying the same herb or spice twice across the five nights. Then re-print the shopping list."*

You'll know it worked when the AI gives you a revised plan *and* a revised list, and you can see that the changes actually flowed through. If the old salmon is still on the list, point that out and ask for another pass.

Step 6 — Save the list where you'll see it on shopping day

Finally, **copy the shopping list out of the chat and into a place you'll actually look at it**. The AI can't push the list to your phone's Reminders app for you, so this last step is on you. Most people either screenshot the list, paste it into their phone's Notes app, or write it on the back of an envelope. While you're at it, copy the five dinners into a separate note titled "This week's dinners" and stick it on the fridge — it stops the 5pm *"what's for dinner?"* panic.

Example: a quick message you can send to yourself or a family member: *"Here are this week's dinners: Mon – chicken stir-fry, Tue – veggie pasta, Wed – sheet-pan sausages and veg, Thu – rice bowl with leftover chicken, Fri – omelette night. Shopping list attached."*

You'll know it worked when the list is in your pocket, your fridge, or your phone — anywhere you'll actually see it before you walk into the shop. If the AI's formatting was messy, just tidy it as you copy.

Common mistakes

- **Asking for "a meal plan" without saying "dinner".** The AI often defaults to three meals a day, which is overwhelming. Always specify *"dinner"* (or breakfast, or lunches) so the list stays useful.
- **Forgetting to mention one dietary need.** A single missed allergy can ruin a whole week. Before you press send, re-read your message and check: did I mention all allergies, all the fussy eaters, and all the equipment I actually have?
- **Trusting the first list without stress-testing it.** The AI doesn't know your family's taste buds. If a recipe sounds like something they'd refuse, swap it now in Step 5 — not at 5:30pm when you're starving and the kids are hungry.

What this means for you

- **In everyday life:** You skip the Sunday-afternoon "what should we eat?" spiral, walk into the supermarket knowing exactly what to grab, and throw away less food because the plan repeats ingredients on purpose.
- **For your business or work:** The same prompt-writing habit transfers anywhere — drafting meeting agendas, planning a project, or organizing a trip. Once you've practiced telling an AI *who you are, what you need, and what to avoid*, you can use that pattern everywhere.
- **If you're just getting started:** Don't worry about getting the prompt perfect on the first try. The AI is patient — if its first answer isn't quite right, just say *"close, but please change X"* and it'll try again. That's the whole trick.

Try it now

Open your AI assistant right now and paste this single message to get started:

> *"I'm a family of [number] in [your city]. We need a five-night dinner plan on a budget of $[X]. We don't eat [allergies or foods to avoid]. Please keep each dinner under 30 minutes, repeat ingredients across nights, and tell me what to defrost the night before."*

Reply to whatever it gives you, ask for the shopping list, and you're done. Ten minutes tonight, five relaxed dinners next week.

Plan a Weekend Trip From a One-Line Idea With Cohere

6 steps · Read online

> **In short:** A single sentence like "weekend in Portland in October" is enough for Cohere to draft a day-by-day itinerary — but it will invent place names if you let it, so prime it with a short trip-planner instruction first, then tighten the plan one change at a time.

You already know how to send a message in Cohere Chat. This guide teaches a workflow you'll reuse every time you plan a trip: a one-time setup step, a structured first ask, and a series of small follow-ups that tighten the plan into something you can actually use on the day. By the end, you'll have a printable weekend outline and a reusable prompt you can paste into any future trip chat.

Before you start

- A Cohere account and the chat interface (web or app) you're already using.
- One rough idea of where you want to go and when — even "fall weekend, somewhere coastal" is enough.
- About 15 minutes the first time, around five minutes for future trips once you have the reusable prompt saved.

Step 1 — Open a new chat and set a trip-planner preamble

Start a fresh conversation so old context doesn't bleed in. Before you type your actual request, look at the top of the chat — there's usually a small field or a collapsible panel for **custom instructions** or a **system message** (sometimes labeled "Persona," or hidden behind a settings icon). Paste a short, plain-English brief: who you are, how detailed you want the answer, and any hard limits. This one-time setup makes every future prompt land better because the model already knows the shape you want. Example: *"You are a practical trip planner. I travel with one other adult, no kids. We like walking, local food, and one museum or landmark per trip max. Avoid luxury recommendations. Keep each day to three time blocks: morning, afternoon, evening. Never invent specific restaurant names — describe the cuisine and neighborhood instead."* You'll know it worked when the next reply already feels structured — it begins with a clear "Day 1 / Day 2" layout rather than a wall of suggestions. If it looks different on your device: if there's no obvious system-prompt field, paste the brief as the very first line of your message instead, prefixed with "Instructions for this chat:" — the effect is similar, just less convenient for reuse.

Step 2 — Send your one-line trip idea as a structured ask

Now type your vague idea, but wrap it in a request for a specific shape. A bare "weekend in Portland in October" is too open — Cohere will pick a structure for you, and you may not like it. So tell it what shape you want: how many days, what time blocks, what each block should contain. This is the difference between "give me ideas" and "give me a Saturday morning plan with a coffee spot." Example: *"Weekend trip idea: Portland, Oregon, second weekend of October, two adults, arriving Friday evening, leaving Sunday afternoon. Build a day-by-day plan in three time blocks per day (morning, afternoon, evening). For each block, give: one main activity, one place to eat (describe the cuisine and neighborhood, don't invent names), and a backup if it rains."* You'll know it worked when the reply comes back split into clearly labeled days and blocks — not a single paragraph of suggestions. If it looks different on your device: some Cohere views collapse long replies; if you only see the first half, look for a "Show more" or down-arrow at the bottom of the reply.

Step 3 — Push the answer toward your actual style

The first draft is rarely the one you'd actually use — it's a starting point. Read it once, then ask for one specific change. The most useful follow-ups are about *style*, not content: shorter walks, later mornings, more indoor options, less touristy. Asking for one thing at a time keeps the model from rewriting the whole plan and accidentally undoing the parts you liked. Example: *"Looks good overall, but tighten this: replace any block that requires a car with a public-transit option, and move any suggestion that starts before 9 a.m. to 10 a.m. or later. Keep everything else the same."* You'll know it worked when the reply keeps the same structure but only the blocks you asked about actually changed — not the whole plan. If it looks different on your device: if Cohere offers a thumbs-up/down or a "regenerate" button next to the reply, don't use regenerate here — that gives you a brand-new plan. Instead, send your tweak as a new message in the same chat so the previous version stays visible above it.

Step 4 — Add a rain backup for one specific slot

Weekend plans live or die by their backup. Pick the *single* most weather-dependent block — usually an outdoor afternoon — and ask for a swap. Naming one slot (not "the whole weekend") keeps the rest of your plan intact and gives you a clean A/B choice for the morning of the trip. Example: *"For Sunday afternoon only, give me a backup plan that works if it rains: an indoor activity, a nearby warm place for lunch, and a coffee stop I can walk to between them. Keep Saturday exactly as it is."* You'll know it worked when you see one new afternoon block with three labeled items, and Saturday's plan is unchanged above it. If it looks different on your device: if Cohere rewrites the whole plan instead of just the one slot, your message probably sounded too general ("give me rain backups") — rephrase to name the day and time block, like "For Sunday afternoon only…"

Step 5 — Ask for a one-screen summary you can actually use on your phone

Long replies are great for thinking; they're bad for using on the go. Ask Cohere to compress the whole plan into a single tight screen — short bullet lines, no descriptions, no explanations. This becomes the version you screenshot or paste into your notes app. Example: *"Now compress the whole plan — both days plus the Sunday rain backup — into a one-screen version I can read on my phone. Each item is one short line. No descriptions, no flowery language. Order: Friday evening, Saturday morning/afternoon/evening, Sunday morning/afternoon, Sunday rain backup."* You'll know it worked when the reply fits in roughly 12–18 short lines and reads like a checklist, not an essay. If it looks different on your device: if the summary comes back too long, send a follow-up like "Cut that in half — keep only the place and the time."

Step 6 — Move the plan out of the chat

A trip plan in a chat window is a plan you'll forget the moment you close the laptop. Send the one-screen summary out into a tool you actually open every day — your phone's notes app, a shared doc, or a calendar entry. This step takes 30 seconds and is the difference between a plan you read once and a plan you actually follow. Example: *"Now write me a copy-paste version of the one-screen summary that I can paste into Apple Notes or Google Keep. Use plain text, no markdown bold, one item per line, with a blank line between days."* You'll know it worked when the reply is plain text you can select, copy, and paste into any notes app without reformatting. If it looks different on your device: if Cohere still outputs markdown (bold headers, dashes), paste it into your notes app first, then strip the symbols manually — usually a 20-second job.

Common mistakes

- **Asking for "the best" anything.** Cohere will confidently invent names. Always ask it to *describe* the cuisine and neighborhood instead — then look up real places yourself on a maps app or review site.
- **Regenerating when you meant to tweak.** The regenerate button (or any equivalent) throws away your working plan. Use it only when the whole direction is wrong, never to fix one line.
- **Trying to make one prompt do everything.** A single huge message gets a generic reply. Better: one structured ask, then a series of small follow-ups. Each follow-up is a small commitment you can keep or throw away.

What this means for you

- **In everyday life:** any time you face a vague-to-specific task — a birthday plan, a moving-day checklist, a guest's weekend visit — the same workflow works: short preamble, one structured ask, one tweak at a time, then a one-screen summary.
- **For your business or work:** the pattern translates directly to any project with a known shape and a few preferences — meeting agendas, conference prep, onboarding checklists. Set the preamble once for the project, then iterate.
- **If you want to go further:** next time, before sending your trip idea, paste in a small block of *your own* notes — a list of places a friend recommended, or three things you already booked — and ask Cohere to fit its suggestions around them. It gets dramatically more useful when you give it something real to anchor against.

Try it now

Open Cohere and start a new chat. Type one sentence — your real, vague idea for a weekend, even if it's just "somewhere warm in February." Then paste the Step 1 preamble above it, followed by the Step 2 prompt structure with your destination filled in. Two minutes, and you'll have a draft you can react to.

Use ChatGPT to Plan a Road Trip with Stops, Fuel Breaks, and Drive Time

6 steps · Read online

> **In short:** Give ChatGPT your start, destination, dates, and a daily driving cap, and it drafts a day-by-day plan with stops, fuel breaks, and sights. Double-check distances and closures yourself — it has no live traffic or gas data, and sometimes invents places that don't exist.

You know that moment when you're staring at a map, scrolling between two pins, wondering where to stop for lunch? That's exactly the kind of question ChatGPT handles well. You give it the boring facts and it gives you back a plan you can edit and actually drive.

Before you start

- A free ChatGPT account works. Paid plans unlock the newest, fastest models — check the official pricing page for current tiers.
- A phone or computer with the ChatGPT app, or any browser pointed at the official site.
- About 20–30 minutes, plus a notes app on your phone for the final plan (Apple Notes, Google Keep, or Notion all work).
- Know four things before you begin: starting city, final destination, travel dates, and roughly how many hours you want to drive each day.

Step 1 — Open ChatGPT and describe your trip in one paragraph

On a computer, open chatgpt.com in your browser and sign in. On a phone, open the ChatGPT app. If this is your first time, you land on a blank screen with a message field at the very bottom of the page — that's where you type. If you see an old conversation instead, look for a **New Chat** button (a pencil or a square with a plus sign) at the top of the screen and tap it to start fresh. Type one short paragraph covering who's traveling, where you're starting, where you're going, and your dates. A few sentences is plenty, because ChatGPT works better with concrete details than vague ones. Example: "My partner and I are driving from Portland, Oregon to San Francisco, California on June 12, leaving around 8am. We like small towns, good coffee, and short hikes. We want to drive no more than 6 hours a day." You'll know it worked when your paragraph appears in the chat and a "typing..." indicator shows up below it.

Step 2 — Ask for a day-by-day itinerary with drive times

Once ChatGPT replies, type a follow-up asking it to break the trip into days, with the exact driving time for each leg and the suggested overnight stop. A new chat bubble appears below your question, and within a few seconds ChatGPT starts typing a numbered day-by-day plan. Each day should show roughly how many hours you'll spend driving, where you'll sleep that night, and a one-line reason for the stop. Example: "Great. Now please give me a day-by-day plan. For each day, tell me the start city, the end city, total driving hours, and one sentence on why that stop is worth it." If the first reply is too general — say, just listing famous cities without hours — reply with "Please add driving hours for each day and the rough total mileage." ChatGPT remembers the conversation, so it rewrites the plan with that detail included. You'll know it worked when each day in the plan has a clear start, end, and a number of hours in the car.

Step 3 — Ask for fuel stops and meal breaks

With the skeleton in place, ask ChatGPT to mark fuel stops and rest breaks along each driving day. Type a short message asking for fuel stations roughly every 3–4 hours, plus one real meal stop per day. The chat updates in place, so the fuel plan lives right below the day-by-day driving plan you already have. Example: "Add a fuel stop about every 3 hours of driving, and one proper meal stop per day. Just give me the town or exit name and roughly when in the day I'd hit it." If your car is electric or your range is shorter, say so in a follow-up — for example, "My EV gets about 250 miles per charge, so plan charging stops around that." ChatGPT will adjust the plan in place. You'll know it worked when each driving day now lists one or two fuel stops and a meal break, with rough timing.

Step 4 — Ask for one or two things to see at each stop

Type a short follow-up asking for one or two small, specific things to see or do at each overnight stop — nothing fancy, just enough to make the drive feel like a trip. ChatGPT adds a "things to do" line under each day in the same chat, mixing one outdoor stop, one food pick, and one small museum or viewpoint where it makes sense. Example: "For each overnight stop, suggest one outdoor thing, one local food place, and one quick museum or viewpoint. Keep it short — we're not trying to cram the whole trip in." If a suggestion feels off — maybe it's closed, far from your route, or you've already been there — just reply "Skip the lighthouse one, give me a different outdoor stop" and ChatGPT will swap it out. You don't need to start over. You'll know it worked when each day has a small "things to do" line that fits the time you actually have.

Step 5 — Ask ChatGPT to spot problems with your plan

This is the step most people skip, and it's the one that saves the most headaches. Type a short message asking ChatGPT to act like a skeptical road-trip reviewer and poke holes in your plan. ChatGPT rewrites its previous answer with a "Things to double-check" section at the bottom, flagging things like long stretches with no fuel, late-evening arrivals in quiet towns, or driving legs that are clearly too long. Treat this as a free second opinion. Example: "Now read the plan back to me like a skeptical friend. What looks too rushed? Which day has the longest stretch between fuel stops? Anything I should double-check before I leave?" If it doesn't flag anything, the plan isn't necessarily perfect — it just means nothing

obvious jumped out. Always verify highway exit numbers, current fuel prices, and road closures with your map app before you drive. You'll know it worked when the plan ends with a short "Things to double-check" section you didn't have before.

Step 6 — Save the final plan to your phone or notes

Type one last short message asking ChatGPT to combine the whole trip into a clean, numbered list you can paste into your notes app. ChatGPT responds with a tidy plain-text version: Day 1, Day 2, each day on its own short block, fuel stops inline, things to do in a few words. Copy it with the **Copy** button at the top of the message, then paste into Apple Notes, Google Keep, Notion, or even a plain email to yourself. Now the plan lives somewhere you can read at a gas station without an internet connection. 💬 Example: "Now condense the whole trip into a clean numbered list I can paste into my notes app. Keep fuel stops inline with each day. Don't add anything new." If your plan changes mid-trip — a road closes, you want to stay an extra night — come back to this same chat and ask ChatGPT to rebuild just that day. It still has the full conversation in memory. You'll know it worked when you have a clean, copyable list on your phone that doesn't depend on ChatGPT being open.

Common mistakes

- **Trusting every name without checking.** ChatGPT sometimes invents a town, café, or gas station that doesn't exist. Before you drive anywhere new, paste the name into a search engine or your map app and confirm it's real and still open.
- **Forgetting that ChatGPT has no live data.** It can't see today's gas prices, current road closures, or weather. Use it for the shape of the trip, then use a real-time map and weather app for the day-of details.
- **Asking for too much in one giant prompt.** A 600-word question usually gets a vague answer. Break the planning into the six steps above — the conversation builds a better plan than any single mega-message could.

What this means for you

- **In everyday life:** The same pattern works for any multi-stop day — a city break, a weekend of errands across town, a move between two apartments. Tell it the start, the end, the time you have, and ask for one chunk at a time.
- **For your business or work:** If you drive for work — delivery, sales calls, field service — paste tomorrow's address list into ChatGPT and ask it to group stops into the most efficient route, with fuel and meal breaks folded in. It's a quick first draft before you load the addresses into a real route planner.
- **If you're just getting started:** The easiest first move is to plan a half-day drive you already know well, just to see how ChatGPT structures a route. You'll learn what to ask for, what to ignore, and how to push back when an answer feels off.

Try it now

Open ChatGPT and paste this: "I'm driving from [your city] to [a town 2–3 hours away] this Saturday, leaving at 9am. Give me a one-day plan with total driving hours, one fuel stop, one lunch stop, and one thing to see along the way." Read the reply, change one detail ("actually make the lunch stop a town with a bakery"), and notice how fast it rewrites. That's the muscle you want before planning the big trip.

Creating a Travel Packing List for Any Trip with AI Guidance

6 steps · Read online

By the end of this guide you'll have a packing checklist that's actually tailored to your trip — not a generic list from the internet — and a reusable prompt you can re-use for every future holiday. This is for anyone who has ever stood in front of an empty suitcase wondering if they've forgotten something obvious.

Before you start

- **An AI assistant with a free account** — ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat or DeepSeek all work for this. Any of them will do. Paid plans work too, but the free tier is plenty.
- **A phone or laptop with internet** — the mobile app or the website both work.
- **A few trip details handy** — destination, travel dates (or rough length), accommodation type, and what you plan to do. Two minutes of thinking here saves ten minutes of back-and-forth later.
- **Roughly 5–10 minutes** for the whole process.

Step 1 — Pick an AI assistant and open a fresh chat

Open any AI assistant you already use, or grab a new one from your phone's app store. Look for a text box on the home screen, usually with placeholder text like "Ask anything..." or "Message [assistant name]". Start a **brand-new chat** rather than continuing an old one — leftover context from a previous conversation can muddy the answer. On a phone, tap the AI's app; on a laptop, head to the assistant's website.

If you don't have one yet, any of these work well: ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or DeepSeek. They all handle this kind of task. If you can't see a new-chat button straight away, look for a small **plus icon**, a **pencil icon**, or a **"New chat"** label, often in the top corner.

Example: open ChatGPT, Gemini, Claude, Copilot, Le Chat or DeepSeek — whichever you already have an account with, or download a free one from your app store.

You'll know it worked when you see a blank conversation screen with a text box ready to type into.

Step 2 — Give the trip basics in one short paragraph

In the text box, type a short paragraph covering where you're going, for how long, and roughly what kind of trip it is (city break, beach holiday, hiking, business). Think of this as a **prompt** — a prompt is just the instruction you give to an AI. The more useful detail you give up front, the less back-and-forth you'll need later. A few clear sentences beat a single vague one.

If the assistant has a **microphone icon** near the text box, you can speak your prompt instead of typing — useful on a phone. Some assistants support this; if you don't see a microphone, just type instead. Look for a small **paperclip or attach icon** if you want to upload a screenshot of your itinerary, a weather forecast, or a booking confirmation; many assistants will read it, but not all support this.

Example: "I'm going to Tokyo for 5 nights in early December with my partner. We'll stay in a hotel in Shinjuku, do a mix of sightseeing, two day trips, and 3–4 nicer dinners. One of us has dodgy knees so we don't want to walk more than about 10 km a day."

You'll know it worked when the assistant replies with a list, or asks one or two quick follow-up questions to clarify.

Step 3 — Add the details that change the list

Reply with the specifics that make or break a packing list: expected weather, planned activities, accommodation type, and any special needs. Don't worry about being too detailed — this is the moment to mention kids, medical items, electronics, or anything unusual like formal events or outdoor gear. Think of it like briefing a friend who has never met you, so they can picture your whole trip.

If you'd rather not type it all out, and if your chosen assistant can read links, paste a link to the weather forecast or your booking email and ask the assistant to read it. If your assistant can't read links, just copy and paste the text instead. The goal is the same: give it enough information to make smart choices for your specific trip.

Example: "Weather forecast says 8–14°C with a chance of rain on two days. We're planning one full-day Disneyland trip, one day in Hakone with onsen, and lots of walking in Tokyo. I take a daily medication and one of us wears glasses. We each have a phone, one small camera, and one universal power adapter."

You'll know it worked when the assistant's reply clearly reflects your specific situation — not a generic list anyone could have written.

Step 4 — Ask for a checklist you can actually use

A wall of text is hard to pack from. Ask the assistant to turn the suggestions into a clean **checklist** (a list with a little box next to each item that you can tick off as you pack), grouped by category — clothes, toiletries, electronics, documents, and so on — and ask it to flag anything easy to

forget. If you want to take it shopping or print it, ask for a format that copies cleanly, like a markdown list with tick boxes.

If the assistant's first reply looks generic, push back: reply with "Please tailor this to my specific trip — I want a custom version, not a template." If it dumps everything into one long paragraph, ask: "Can you put this in a bulleted checklist grouped by category?" Most assistants will happily redo the formatting.

Example: "Turn this into a printable checklist grouped by category (clothes, toiletries, electronics, documents, medications, miscellaneous), and add a short 'easy to forget' section at the end. Use markdown checkboxes so I can copy it into my notes app."

You'll know it worked when the reply is a tidy grouped list with tick boxes you can copy or screenshot in one go.

Step 5 — Edit the list to match real life

No AI knows everything about you. Skim the list, cross off anything that doesn't apply, add anything it missed (your lucky undies, that one specific cable, the kids' teddy), and tweak quantities if they look off. Treat the AI's list as a strong first draft, not gospel. If you spot the same item listed twice, delete the duplicate.

If you want a shopping list of things to actually buy, ask: "From this list, give me a separate short 'what to buy' list of items I probably don't already own at home." That second list is often the most useful part. If the assistant can't tell you local prices or shops, it won't pretend to — and that's fine, that's where your own knowledge comes in.

Example: "Now make a separate 'to buy' list with just the items I probably don't already have, like travel-size toiletries or any specific gear for the onsen."

You'll know it worked when the shopping list is short, specific and only contains things you actually need to purchase.

Step 6 — Save it and reuse the prompt for your next trip

Once the list feels right, copy it into your notes app, email it to yourself, or screenshot it so it's with you when you actually pack. Save the original prompt too — it's a template you can re-use by swapping in the new destination, dates and activities. Next trip, paste the old prompt, change the details, and you'll get a fresh tailored list in under a minute.

One honest note: an AI can build you a great checklist, but it can't pack your bag, check the live weather on the morning you leave, or remind you to charge your phone. Treat the list as a strong starting point, not a guarantee.

Example: "Save this whole conversation as 'Packing prompt template'. Next time: paste it, change the destination and dates, delete anything that no longer applies, and re-send."

You'll know it worked when you have a saved packing template you can re-use in under a minute for any future trip.

Common mistakes

- **Vague prompts give vague lists.** "Help me pack for a holiday" usually returns a generic template. Always include destination, dates, and trip style up front.
- **Forgetting the unusual stuff.** AI lists are great for clothes and toiletries, but easy to miss on medications, glasses, kids' items, or visa documents. Add a sentence about anything special.
- **Treating the first reply as final.** If the list looks generic, ask the assistant to redo it specifically for your trip. One more prompt usually fixes it.

What this means for you

- **In everyday life:** less pre-trip stress, fewer forgotten items, less last-minute shopping at inflated tourist prices.
- **For your business or work:** a reusable template for every work trip — same structure, just swap in the city, dates and meeting types.
- **If you're just getting started:** copy the example prompt from Step 2, paste it into any AI assistant, and edit three details to match a real trip on your calendar.

Try it now

Open any AI assistant you already have, paste the example prompt from Step 2, and change three things — destination, number of nights, and one planned activity — to match a real trip on your calendar. You'll have a tailored packing list before your coffee gets cold.

How to scale a family recipe up or down for any number of guests

6 steps · Read online

> **In short:** Multiplying ingredients by a single number is the easy part — but cooking time, pan size, and a few bold flavors (like salt and baking powder) don't scale the same way. Here's how to scale a recipe without ruining it.

Every family recipe was written for one number of people, but most of us cook for crowds that refuse to stay that size. The good news: the ingredient math is the smallest part of the job. The bigger part is knowing *what doesn't' scale and adjusting for it on purpose.*

Before you start

- A printed or saved copy of the recipe you want to scale (look at the original serving size first).
- A calculator — the one on your phone works perfectly.
- Measuring cups and spoons; a kitchen scale is a nice extra if you have one.
- A pen or note app to jot down your multiplier and any tweaks you make.
- About **30 minutes** for the math, scaling, and the first round of cooking.

Step 1 — Find your scaling factor

Every later step depends on this one number, so get it right first. The math is simple: divide the number of people you want to feed by the number the recipe was originally written for. That decimal is your **scaling factor** — your "multiplier" — and you'll use it again and again. Think of it as the single dial that turns the whole recipe up or down.

For example, a recipe that feeds 6 is being stretched to 14 guests: 14 ÷ 6 = 2.33. Write that number somewhere visible — top of the recipe page, on a sticky note, or in your phone. **If your recipe doesn't list a serving size,** count what it's "meant to feed" — chicken for 4 means 4 portions of protein — and round your factor **up** rather than down. Leftovers beat running out every time.

Example: original serves 6, you need 14 → factor = **2.33**; original serves 8, you need 4 → factor = **0.5**.

You'll know it worked when you have one decimal number that you can call out (or whisper) every time you measure an ingredient.

Step 2 — Multiply every dry, fatty, and bulk ingredient

Now that you have your factor, scale almost everything in the recipe: flour, sugar, rice, pasta, chopped vegetables, oil, and butter. Multiplying by 2.33 means **1 cup becomes 2.33 cups** — that is, just over 2¼ cups. Don't try to round too early; the decimals will wash out once everything goes in the pan. This is where most of the "volume" of your scaled dish comes from.

Example: 1 cup flour × 2.33 = about 2¼ cups; 2 tbsp butter × 1.5 = 3 tbsp.

If your calculator dies or you lose track, that's fine — for most spices and bulk ingredients, being within 10% of the math is close enough. The flavor of these ingredients scales reliably because they're part of the structure, not the seasoning.

You'll know it worked when you have a complete pile of measured ingredients, each one adjusted to your new guest count, ready to mix or cook.

Step 3 — Treat salt, leavening, and bold spices more gently

Here's the first place the math breaks down. **Salt, baking soda, baking powder, yeast, dried chili, and raw garlic** are flavor amplifiers, not bulk — they get loud fast when concentrated. Multiplying them by the full factor usually makes a scaled-up dish taste harsh, bitter, or chemical. The fix: multiply them by about **75–90%** of your factor instead. If your factor is 2.33, scale salt by 2, not 2.33. This single step is what separates a "fine" scaled recipe from one that genuinely tastes right.

Example: 1 tsp salt in the original × factor 2.33 (full math) would be ~2.33 tsp; gentler multiplier 2 → about 2 tsp is the right target.

If you're scaling DOWN, do the opposite: don't cut these ingredients by the full math ratio, or the dish goes flat. Use about 70% of the calculated amount, then taste at the end and add more if needed.

You'll know it worked when the seasoning in your scaled dish tastes similar to the original — not saltier, not duller, just right.

Step 4 — Pick a pan that fits the new volume

This is the step most home cooks forget, and it's where most scaled-up failures come from. Your recipe was designed for a specific pan shape and depth, and switching serving size without changing the pan leads to raw middles and burnt edges. The rule of thumb: aim for roughly the **same depth of food** in the pan as the original recipe. A cake for 6 in an 8-inch round is about 2 inches deep; when you scale up for 12, a 9×13-inch rectangular pan gets you there instead of trying to fit it all in a too-small round.

Example: original 9×13 pan for 8 servings → 16 servings → use **two** 9×13 pans, or one large roasting pan.

If you don't have a bigger pan, use two of the original size side by side — that's the safest backup. Cooking two smaller dishes is almost always better than cramming everything into one overflowing pan.

You'll know it worked when the food fills the pan to about the same depth as the original recipe, not piled higher or sitting in a too-thick layer.

Step 5 — Adjust cooking time and oven temperature carefully

Time doesn't scale with servings — heat has to travel through the food, and bigger batches just take longer to heat **all the way through**. Keep the oven temperature the same as the recipe says, but start checking for doneness 10–15 minutes **before** the original cooking time. Then add cooking time in small chunks. A doubled casserole might take about 1.3× longer, not 2× longer. A doubled cake may need only 5–10 extra minutes, not double.

Example: original bake time = 45 min → start checking at **30 min**, then every 5 min after, using a toothpick or thermometer.

If the dish is very thick or dense (lasagna, bread, dense meat loaf), expect closer to a 1.5× time increase. **For thin or mostly-liquid dishes** (soup, custard, sauce), the time barely changes — just rely on the visual cues the original recipe gave you.

You'll know it worked when the food is fully cooked inside — not just colored on the outside — at a time that's "a bit longer than the original," not doubled.

Step 6 — Taste, adjust, and save your notes for next time

A scaled family recipe is a working document, not a finished product. The first scaled batch is always a learning experience, and that's the point. When the dish is done, taste it and write down what you'd change next time — more salt, less bake time, a bigger pan. Most home cooks find that three rounds in is when a scaled recipe finally feels "dialed in" for any crowd.

Example: "Scaled 2× for Thanksgiving. Next time: bigger pan, 10 min less in oven, +½ tsp salt."

If the first attempt misses the mark, that's not a failure — it's data. Salt too low? Add a pinch at the table. Too dry? Spoon a little sauce on. The next batch will be better because of this one.

You'll know it worked when you have a working scaling factor and a few small tweaks written down for next time — and a dish your guests are still talking about.

Common mistakes

- **Doubling salt and baking powder along with everything else.** The dish turns harsh, bitter, or flat. Fix: scale these by 75–90% of your factor, then taste.
- **Using the same pan and hoping for the best.** A too-full pan cooks unevenly — burnt edges, raw middle. Fix: aim for the same depth of food as the original, or split into two pans.
- **Multiplying the cook time** by the same factor. Time doesn't scale that way. Fix: keep the oven temperature, start checking 10–15 minutes early, and rely on visual cues or a thermometer instead of the clock.

What this means for you

- **In everyday life:** you can host impromptu dinners, potlucks, or holiday gatherings without panic-buying takeout as backup. Scaling becomes a small mental habit, not a major project.
- **For family recipes and gatherings:** the recipe stays *yours* forever — you just adjust it to the crowd. Grandma's cookie recipe works for 4, 14, or 40 once you know the trick.
- **If you're just getting started:** don't try to tackle a wedding cake on day one. Pick one reliable family recipe and try multiplying it by 1.5× this week. The math gets easier with practice, and so does the intuition.

Try it now

Open the calculator app on your phone, choose a recipe you trust (something simple is best — cookies, soup, or a one-pan dinner), and work out the factor you need for next week's meals. Multiply the dry ingredients and salt by hand, just to feel how short the math actually is. The ingredient list is the easy half — and now you know exactly where the tricky half lives.

Plan a Relaxing Weekend Getaway From Scratch in One Sitting

6 steps · Read online

> **In short:** You can plan a calm, realistic two-day weekend trip in one focused session by asking an AI assistant to draft the itinerary, the budget, and a backup plan for bad weather — then you tweak the parts that matter to you. An AI cannot book flights or reserve hotels for you, so those final steps stay with you.

A relaxing weekend does not need weeks of research. With a clear set of questions and a blank chat window, you can go from "I have no idea where to go" to "I know what Saturday morning looks like" in under an hour. This guide walks you through that session step by step, so the only thing left at the end is packing.

Before you start

- **An AI assistant you already use** — any chat-based tool will work (ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek). No special plan needed; the free tier is usually enough for this.
- **A rough idea of where you are leaving from** — your city, region, or even just "I am open to driving up to 3 hours."
- **Your weekend dates** — which two days you are planning for.
- **Rough budget in your head** — a single number like "around $300 total" is plenty. The assistant will help you split it.
- **About 45 minutes** — most of which is thinking, not typing.

Step 1 — Set up the planning prompt

Open your AI assistant's chat window in a browser or app and start a new conversation. A blank chat box usually sits at the bottom of the screen, with a text field and a small arrow or paper-plane icon to send. Type your request as one clear paragraph that covers three things: where you want to start from, what kind of trip you want, and any hard limits. If you do not see a chat box, look for a button labeled something like **New chat**, **Start a conversation**, or +. Send the message and wait for the reply — usually within seconds, sometimes up to a minute on busy days.

💬 Example: "I live in Austin, Texas, and I want a relaxing weekend trip for the second weekend of March. I have a car, around $400 total to spend, and I do not want to fly. I am traveling with one other adult and we both like quiet nature and good food, not nightlife."

You'll know it worked when the assistant replies with a structured plan — usually some kind of intro paragraph followed by suggested destinations or a day-by-day breakdown.

Step 2 — Ask for a realistic two-day itinerary

Reply to the assistant's first answer and ask for a tighter, hour-by-hour plan for the two days. Tell it how you actually like to spend mornings and evenings — for example, slow mornings, an early dinner, no alarms — and ask it to bake those habits in. The assistant will usually rewrite the plan with time blocks like "9:00am — coffee and a walk" or "4:00pm — check in." If the reply is still too vague, tell it: "Make it more specific, hour by hour, for two days." It will usually respond with a sharper version.

💬 Example: "Please redo it as an hour-by-hour itinerary for Saturday and Sunday. Keep mornings slow, plan one main activity per day, and put dinner before 8pm. We do not drink, so please skip bars and breweries."

You'll know it worked when you can read down the list and picture your weekend — you know what you are doing at 10am, at 3pm, and at dinner.

Step 3 — Add a rough budget breakdown

Ask the assistant to put a price tag on each part of the plan: lodging (a place to sleep), food, fuel or local transport, activities, and a small "extras" bucket for things like coffee or a souvenir. It will not know exact prices in your area, so it will give ranges — that is fine. Your job is to read the numbers and adjust anything that feels too high or too low for your own life. If the totals do not fit your budget, tell the assistant to swap in cheaper options; it usually will.

💬 Example: "Now give me a rough budget breakdown in USD for the whole weekend. Split it into lodging, food, transport, activities, and a small extras bucket. Keep the total around $400 for two people."

You'll know it worked when the items add up to roughly your budget and you can see where the biggest costs are. If the total is way over, that is fine — it just means the next step is where you adjust.

Step 4 — Build a backup plan for rain or tiredness

Now ask the assistant: "If it rains all weekend, or if we are tired by Saturday afternoon, what is the slower indoor version of this plan?" This is the part most people skip, and it is the part that decides whether your weekend actually feels relaxing. The assistant will usually rewrite the same trip

with museums, cafes, bookstores, or a long lunch instead of hiking or sightseeing. Save this version somewhere — copy it into a notes app, an email to yourself, or just a second message in the same chat.

💬 Example: "Now give me a Plan B for the same weekend if it rains nonstop or if we are too tired to do the outdoor stuff. Keep it calm, indoors, and cheap."

You'll know it worked when you have two short plans sitting side by side — the sunny one and the rainy one — and you can switch between them without thinking.

Step 5 — Have the assistant flag what to book in advance

Send one more message asking the assistant to list the things you need to book or reserve ahead of time — usually lodging, any restaurant with a long wait, and any timed-entry tickets. It will often add a small tip about cancellation policies. Treat this list as your short to-do before the weekend: book the sleep spot first, then anything time-sensitive, then everything else can wait until you arrive. If your assistant offers to do the booking itself, it cannot — that part is on you.

💬 Example: "Based on this plan, what should I book in advance? Please list the items in the order I should handle them this week."

You'll know it worked when the reply is a short, ordered list — usually three to five items — that you can work through in one sitting.

Step 6 — Save the plan and close the laptop

Copy the final plan — both the main version and the backup — into a single note. Your phone's notes app, a Google Doc, or even a pinned email all work. Give the note a simple title like "Weekend — March 14–15" so you can find it later. Then close the chat window. The point of doing this in one sitting is to stop second-guessing; the plan you saved is the plan you follow. If something changes later, you can reopen the chat and adjust, but you do not have to.

💬 Example: "Thanks, please summarize everything into one short message I can paste into my notes app — Plan A, Plan B, the budget, and the booking list."

You'll know it worked when you can put your phone away, look up from your desk, and feel like the weekend is already half-planned.

Common mistakes

- **Asking for too many destinations at once.** "Plan a trip" works better than "compare five destinations" — the assistant gets lost. Pick a place, or ask it to suggest two and choose one before going deeper.
- **Skipping the budget conversation.** Without a number, the assistant will quietly suggest hotels and restaurants that are too nice for real life. A budget line in the first prompt saves a redo later.
- **Forgetting the backup plan.** The whole point of "relaxing" is low stress. If your only plan requires sunshine and energy, you have not really planned — you have just hoped. Always leave the chat with a Plan B.
- **Treating the assistant's prices as real.** It will give numbers, but they are guesses based on what it learned during training. Always check the actual price on the booking site before you pay anything.

What this means for you

- **In everyday life:** A short, repeatable planning habit. Any time a small trip feels too big to start, one focused chat turns it into a weekend you can actually look forward to.
- **For your business or work:** The same shape works for planning a conference trip, a client visit, or even a personal day off. The ingredients are the same: a starting point, a structure, a budget, and a backup.
- **If you're just getting started:** Use your own city as the "destination" the first time. Plan a quiet Saturday in your own neighborhood. It is the cheapest way to learn how the assistant thinks before you trust it with a longer trip.

Try it now

Open your AI assistant right now and send one message: "Help me plan a calm Sunday in my own city this weekend. I want one outdoor walk, one good meal, and one slow afternoon at home. Keep the total under $50." Read the reply, save it to your notes, and you have just practiced the whole workflow in two minutes.

Brainstorm birthday gift ideas with Le Chat for any budget and personality

6 steps · Read online

> **In short:** Describe one person in a short prompt, ask Le Chat for gift ideas grouped by price, and you'll get a brainstormed list in seconds. Le Chat can't see live stock or place orders — you'll still buy the gift yourself.

Scrolling an online store for an hour trying to find the "right" birthday gift is a familiar kind of stress. A short conversation with Le Chat (the personal assistant from Mistral) can turn a few details about a person into a tidy, budget-sorted list of ideas you can actually act on. Here's how to set it up so the suggestions feel personal — not generic.

Before you start

- A free Le Chat account. Sign up at chat.mistral.ai with an email and password — the free plan is enough for casual brainstorming. A paid plan exists if you want higher usage limits, but it's not required for this.
- A web browser, or the Le Chat app on iOS or Android.
- A rough idea of the person's hobbies, age, and how much you'd like to spend.
- About 10 minutes total.

Step 1 — Open Le Chat and start a new chat

Open your browser and go to **chat.mistral.ai**, or tap the Le Chat app on your phone. If you're not signed in, look for a **Sign in** or **Log in** button near the top-right of the page — entering your email and a password is usually enough. Once you're in, the screen shows a clean chat box, often with a friendly line like **"How can I help you today?"** at the top. Find the **"+ New chat"** or **"New conversation"** button — usually a pencil icon in the sidebar, sometimes a button in the top-left — and click it so your gift brainstorm gets its own fresh thread. If you're on mobile and don't see a "+" button, look for a small chat-bubble icon with a plus sign instead; both do the same thing. **You'll know it worked when** the chat area is empty and ready for your first message, with no previous conversation showing.

Step 2 — Describe the person in detail

Before asking for gifts, give Le Chat something to work with. Type a short paragraph about the birthday person — their age range, what they enjoy, what they already own, and anything they care about. The more specific you are, the better the suggestions. Type into the message box at the bottom of the screen and press Enter (or tap the send-arrow icon). If your screen looks different and you don't see a text box, look for a **"Send a message"** or **"Ask anything"** field near the bottom. After you send, Le Chat will reply with a short acknowledgment or a quick follow-up question — that's normal, not an error. **You'll know it worked when** Le Chat responds to your description instead of saying it doesn't have enough information. 💬 Example: "My sister is turning 32. She loves cooking, has a tiny apartment kitchen, and already owns a stand mixer. She hates clutter and reads a lot of mystery novels."

Step 3 — Ask for ideas grouped by budget

Now ask for gift suggestions sorted by how much you want to spend. Type a clear request that names the price ranges you care about — three tiers is usually enough: a small "thoughtful" gift, a mid-range "main" gift, and a bigger "splurge" option. Send the message the same way you did in Step 2. Le Chat will produce a list grouped under each price tier, often with a one-line reason why each idea fits the person. If the reply doesn't appear instantly, wait a few seconds — Le Chat is typing. **You'll know it worked when** you see a structured list with price ranges as headers (like "budget-friendly", "mid-range", "splurge"), not just one flat list of random items. 💬 Example: "Based on that, suggest birthday gift ideas for her grouped into three price tiers: a small thoughtful gift, a mid-range main gift, and a bigger splurge option. For each idea, tell me in one sentence why it suits her."

Step 4 — Push back, refine, and personalize

The first list is a starting point, not the final answer. Read through it and pick one or two ideas that almost fit but need tweaking — then ask Le Chat to adjust. Maybe an idea is too generic, the price is slightly off, or you want something more personal. Type a follow-up message in the same chat and Le Chat will rewrite or expand on those items using everything you already told it. This is where the brainstorm really pays off: Le Chat remembers the details from your earlier messages, so you don't have to repeat them. If you ever feel stuck, ask Le Chat to "give me 5 wilder ideas" or "5 quieter, calmer ideas" to widen or narrow the style. **You'll know it worked when** Le Chat's reply directly addresses your tweak — for example, names a smaller version of the original item, or swaps a product for something more handmade. 💬 Example: "I love the cookbook idea but she already has a shelf full of cookbooks. Can you replace it with something equally thoughtful in the same price range? Also, swap one of the splurge ideas for something experience-based instead of a physical object."

Step 5 — Save or export the ideas you like

Once you have a shortlist of 3–5 ideas you actually like, get them out of the chat so you don't lose them. Look for a **copy** icon near the assistant's message — often in the top-right of the reply, sometimes as two overlapping squares — and tap it. That puts the whole reply on your clipboard so

you can paste it into a note app, email, or phone note. If you don't see a copy icon, look for a **"Share"** or **three-dot ⋯ menu** near the message; that usually has options to copy, share, or export. On mobile, you can also long-press the message text to copy just that part. **You'll know it worked when** the reply text appears in your notes app, your email draft, or wherever you pasted it, exactly as Le Chat wrote it. 💬 Example: Ask Le Chat: "Give me a clean final list of the top 4 gift ideas, each with a one-sentence reason and the price range, so I can paste it into my notes."

Step 6 — Double-check before you buy

Le Chat is a brainstormer, not a shopper. Before you actually spend money, do a 60-second reality check on your top pick. Open a search engine or your favorite store and confirm the item exists, the price is roughly what Le Chat said, and it's something you can actually get where you live. Some suggestions will be fictional brand names, vague categories, or items that are out of stock — that's normal, just swap in a real alternative. If a suggestion doesn't pan out, drop back into Le Chat and ask for two more ideas in the same price tier. **You'll know it worked when** you've got one real, available product in mind and a rough budget confirmed — that's your gift sorted, in a fraction of the time you'd spend scrolling.

Common mistakes

- **Too vague a description.** Writing only "my mom, 60" will get generic gift ideas. Include at least three specific details — a hobby, a constraint, a taste. The richer the prompt, the more personal the list.
- **Asking for everything at once.** "Give me 50 ideas across all budgets" usually returns a long, unfocused list. Asking for **about 10 ideas across three tiers** is much easier to skim and refine.
- **Trusting the list without checking.** Le Chat can suggest items that don't exist, are out of season, or are wildly overpriced. Always do a quick real-world price check before buying, especially for the splurge tier.

What this means for you

- **In everyday life:** When a birthday sneaks up on you, a 10-minute brainstorm session beats an hour of scrolling. You'll walk into the store — or open a checkout page — already knowing what you want.
- **For your work or business:** The same prompt pattern — *describe the person, sort by budget, refine with constraints* — works for client gifting, welcome kits, or picking thoughtful team rewards. It's a small system you can reuse every quarter.
- **If you're just getting started:** Don't worry about writing the "perfect" prompt. A messy, honest description of the person is better than a polished one. Le Chat works with what you give it.

Try it now

Open Le Chat and paste this in one message: *"My friend is turning 28. He loves cycling, lives in a small apartment, drinks black coffee, and already owns a good laptop. Suggest birthday gift ideas in three price tiers (a small thoughtful gift, a mid-range option, and a bigger splurge), one sentence per idea."* See what comes back — you'll have a real shortlist in under a minute.

How to create a custom bedtime story starring your child with any AI

6 steps · Read online

> **In short:** Any modern AI assistant — ChatGPT, Gemini, Claude, Copilot, Le Chat, or others — can write a custom bedtime story starring your child in about five minutes. You type a few clear sentences about who to include and what should happen, and the assistant drafts a story you can read aloud or save.

In the next few minutes, you'll have a story where your kid is the main character, their goldfish is the sidekick, and the adventure happens in their favorite park — or anywhere else you imagine. This guide works with any AI assistant you already use, and you don't need any special skills, just the ability to type a few sentences.

Before you start

- **An AI assistant you can type into.** Use ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or any other chatbot you have an account with. A free plan is enough.
- **A device with internet access** — phone, tablet, or computer all work.
- **About five minutes** for the first try.
- **Three small pieces of information** to feed the assistant: your child's first name, a sidekick (a pet, stuffed animal, or imaginary friend), and one favorite place.

That's it. No special software, no payment needed for the basic version, no technical knowledge.

Step 1 — Open a fresh chat with your AI assistant

Go to the website or app of the AI assistant you use, and **start a new, empty conversation**. Most assistants put a big button or a "**New chat**" option at the top of the screen. If you don't see that, look for a **pencil icon**, a **plus sign**, or a button that says something like **"New conversation"** or **"Start new chat"**. A blank chat is important because the assistant should focus only on your story — old messages from a previous chat can confuse it. 💬 Example: nothing to type yet — just open the app and tap the new-chat button. You'll know it worked when you see an empty message box with no earlier messages above it.

Step 2 — Tell the assistant what role you want it to play

In that empty message box, type one short sentence that asks the AI to act as a **bedtime story writer**. A **prompt** is just the instruction you give an AI — think of it like asking a friend to help. Asking the AI to play a specific role helps it write in the right style: gentle, warm, and suitable for kids. If the button on the screen is labeled something like **"Send"**, **"Ask"**, or has a **paper airplane icon** ✈️, that's the one to press after you type. 💬 Example: `You are a kind bedtime story writer for young children. Please write in simple, gentle language with a happy ending.` You'll know it worked when the assistant replies confirming the role, often with a sentence like "Sure, I'd love to help!" or a short acknowledgment.

Step 3 — Add the three key ingredients: the child, the sidekick, the place

Now send a second message that names who is in the story and where it happens. Keep it short and specific — the more detail you give, the more personal the story feels. A **sidekick** is just the character's helper or best friend in the adventure. Include your child's first name, the sidekick (their pet, a teddy bear, or even an imaginary dragon), and one favorite spot (the backyard, the playground, grandma's house). If your assistant uses different button names, look for a **send arrow**, a **return/enter key**, or a circle with an upward arrow. 💬 Example: `Please write a story starring my 6-year-old daughter Lily, her hamster named Biscuit, and their backyard with the big oak tree. Lily is shy but very curious.` You'll know it worked when the assistant starts writing a story that mentions Lily by name within the first few sentences.

Step 4 — Ask for a bedtime-friendly length and a gentle lesson

Every AI can write short or very long stories. For bedtime, a short one works best so your child doesn't get restless. Send another short message asking for a story that takes about five minutes to read aloud, with **one soft lesson** woven in — like being brave, sharing, or trying something new. A **lesson** here just means a small, kind idea the story quietly teaches, not a lecture. If the assistant gives you a story that's too long, just say "shorter, please" and try again — there's no penalty for asking. 💬 Example: `Please make it about 500 words long — the length of a 5-minute bedtime read — and gently show how being brave helps, without being preachy.` You'll know it worked when the story comes back at roughly the right length and ends with a calm, peaceful feeling rather than a cliffhanger.

Step 5 — Read it, then ask for changes

Read the story out loud to yourself. If a part feels off — a character does something scary, a word is too hard, or the ending feels rushed — tell the assistant what to fix. This back-and-forth is normal and is how you **collaborate** with the AI (meaning you work on it together, taking turns). You might ask for a different ending, swap the sidekick for a unicorn, or add a sibling. The assistant rewrites the whole story in seconds. If the assistant offers choices like **"Regenerate"** or **"Try again"**, those are buttons that produce a fresh version using the same instructions. 💬

Example: `This is lovely, but please change the ending so Lily and Biscuit fall asleep peacefully under the oak tree instead of running home.` You'll know it worked when the new version includes your requested change and still has Lily's name in it.

Step 6 — Save or share the story

Once you're happy, copy the story into your phone's notes app, email it to yourself, or print it. Some assistants let you **share** the whole conversation through a **link** (a web address you can send to someone) — look for a **share icon**, which often looks like an arrow pointing up out of a box, or three connected dots. You can also screenshot the story page. The same prompt will work again tomorrow with different names, so you've just learned a skill you can repeat every night. 💬 Example: tap the **share icon**, choose **"Copy link"** or **"Send to email"**, and save the story in a folder called "Bedtime stories." You'll know it worked when the story is saved somewhere you can find it again, even after you close the chat.

Common mistakes

- **Forgetting to start a fresh chat.** Old messages can leak into the new story and confuse the assistant. If your story suddenly mentions your grocery list, that's why — open a new chat and try again.
- **Asking for too much at once.** "Make it exciting, funny, scary-but-not-too-scary, with songs and a dragon" overwhelms the assistant. One or two clear requests per message works much better.
- **Skipping the read-aloud test.** Stories that look fine on screen can sound clunky when spoken. Always read it once out loud before reading it to your child — your ear catches problems your eyes miss.

What this means for you

- **In everyday life:** You have a free, endless supply of personalized bedtime stories. The same trick works for birthday parties, long car rides, or "I'm bored" afternoons — change the names and place.
- **For your business or work:** This same prompt technique — asking an AI to play a role, then giving clear ingredients — is the foundation of using AI for emails, reports, and lesson plans. You've just practiced the most important AI skill there is.
- **If you're just getting started:** You've now used an AI assistant for something creative, fun, and useful. That feeling of "I made this work" is the same feeling you'll have the next time you try something new with AI.

Try it now

Open your AI assistant right now, start a new chat, and paste this single message to get your first story in under a minute:

```
You are a kind bedtime story writer for young children. Please write a 500-word bedtime story starring my child named [YOUR CHILD'S NAME], their pet [PET NAME], set in [FAVORITE PLACE]. Make it gentle, end with them falling asleep peacefully, and quietly show that being brave is good.
```

Replace the three bracketed parts with real details, hit send, and read the result out loud. That's your first custom bedtime story — and the start of a habit you can repeat as many nights as you like.

How to organize a chaotic shared family calendar in one afternoon

6 steps · Read online

> **In short:** Pick one calendar app everyone can see, create a shared family calendar, and spend about two hours adding every school event, work shift, and after-school practice. After that, set phone reminders so nobody asks "where am I supposed to be?" at 7am.

It is 6:47 on a Tuesday morning. Someone is hunting for a permission slip. Someone else is asking whether soccer practice got moved. The dog needs to be walked before anyone leaves. This is the moment a shared family calendar was invented for, and you can build one this afternoon while the kids are at practice and the laundry is running.

Before you start

- A free calendar app that everyone in the household can install: Google Calendar (works on Android, iPhone, and any browser), Apple iCloud Calendar (built into every iPhone and Mac), or Outlook Calendar (if your household lives on Windows or uses Outlook at work). All three have a free tier and let several people see and edit the same calendar.
- Each person's phone, signed in to the same account family members will share. (Children do not need their own email — they can appear as a placeholder once you add them.)
- A rough 90–120 minutes for the first pass. Updating after that takes about 10 minutes a week.
- A rough list of recurring commitments. Skim the school newsletter, the team schedule email, and your work rota — you will copy-paste these in Step 4.

Step 1 — Pick one calendar app the whole family will actually open

Open the calendar app that is already on your phone, or install one from your app store. If your family uses Android phones and a Windows computer, Google Calendar is usually the easiest because the free Google account it lives on is also what Android expects. If everyone has an iPhone, Apple iCloud Calendar works without any extra sign-up. If your workplace runs on Microsoft, the Outlook calendar keeps work shifts and family events in the same place.

Pick the option that matches what your family already uses — the best calendar is the one nobody has to remember to open. Sign in on your own phone and confirm you can see your existing events before moving on.

Example: *"We have three iPhones and a Mac — let's use Apple iCloud Calendar so we don't have to install anything new."*

You'll know it worked when you open the app and see your existing appointments, birthdays, or that empty workweek you have been ignoring.

Step 2 — Create a new calendar called "Family"

Inside the app, look for a way to add a new calendar. On Google Calendar this lives behind the plus icon next to "Other calendars" in the left sidebar on the web; in Apple Calendar on Mac it is under the `File` menu; in Outlook it is `Add calendar → Create new`. The exact wording varies by app and device, but every major app has an "add calendar" button somewhere.

Name it clearly — "Family" works well. Pick a bright color so it stands out from your work calendar the moment you open the week view.

Example: *"Calendar name: Family. Color: orange. Description: shared with mom, dad, Mia, and Leo."*

You'll know it worked when the new calendar appears in your sidebar with its own color and is empty.

Step 3 — Invite each family member to the new calendar

Now share the new calendar. In Google Calendar on the web, hover over the calendar name and look for a settings or share option, then add each family member's email. In Apple Calendar on Mac, look for a "Shared Calendar" toggle in the calendar info panel and paste iCloud email addresses. In Outlook, right-click the calendar and pick "Sharing and permissions". If you cannot find the exact button, search the app's help for "share calendar with family" — most apps walk you through it in two taps.

Each person accepts the invite on their phone. For younger children, add a placeholder name like "Mia (kid)" rather than creating a real account — they only need to see the events, not own them.

You'll know it worked when each person's phone shows the family events you add in the next step, usually within a minute or two.

Step 4 — Add the big repeating events first

Now the satisfying part. Think of the calendar as four layers and tackle them in order. Start with the hardest layer — fixed weekly commitments — because they happen every week and never change. Add soccer practice every Tuesday and Thursday, school assembly every Monday morning, your partner's shift pattern, your own standing meetings. Each app lets you set an event to repeat: look for a phrase like "Does not repeat" and change it to "Weekly" or "Custom" with the right days.

Set a reminder on each one. Most apps default to 10 minutes before, which is too short when you are getting three kids out the door. Look for the reminder setting and push it to 60 or 90 minutes, plus a second alert 15 minutes before.

Example: *"Tuesday 16:00–17:30, Soccer practice, Mia. Repeats every Tuesday. Reminder: 90 min before (popup) and 15 min before (notification)."*

You'll know it worked when the first event shows up on the shared calendar for everyone and a reminder pops up on your phone before the activity starts.

Step 5 — Add one-time events from this week's emails

Open your email and your partner's email (with permission, or just ask them to forward the relevant ones). School newsletters, dentist appointments, the birthday party on Saturday — copy each one into the shared calendar as a single event. If it has a location, add it. If the email includes a sign-up link or a note, paste the link into the event's notes or description box.

This step is messy the first time, and that is fine. The point is not perfection — it is getting the next 14 days out of everyone's head and onto one screen.

Example: *"Saturday 14:00–16:00, Leo's birthday party at the community center. Notes: RSVP link — bring a swimsuit."*

You'll know it worked when you can ask "what is on Thursday?" and tap the family calendar to find the answer in three seconds.

Step 6 — Set a 10-minute Sunday review habit

A shared calendar only stays calm if someone owns it. Pick one person — usually the one who already notices missed permission slips — to spend ten minutes every Sunday adding the next two weeks. Treat it like brushing your teeth: small, regular, easy to skip if you are tired.

Set a recurring reminder on your own phone: "Sunday 19:00 — family calendar check". Open the calendar, scan for missing items, add them, done.

Example: *"Sunday 19:00–19:10, Family calendar review, repeats weekly. Reminder: at start time."*

You'll know it worked when Sunday-night panic about forgotten events quietly stops happening.

Common mistakes

- **Skipping the color step.** Without a clear color for the family calendar, events blend into your work calendar and you still do not see them. *Fix:* pick the brightest color the app offers and use it only for family events.
- **Setting reminders 10 minutes before.** That is enough time to read a notification, not enough time to find a shoe and leave the house. *Fix:* set a long lead reminder (60–90 min) plus a short one (10–15 min) on every important event.
- **Forgetting the partner's work shifts.** If shift work runs your house, an unshared shift calendar quietly undoes everything. *Fix:* ask your partner to share their work calendar with yours, or add the recurring shifts to the family calendar in Step 4.

What this means for you

- **In everyday life:** the Tuesday-morning "where am I supposed to be?" texts disappear, and quiet Sundays get a little less tense.
- **For your business or work:** the same habit — a single shared calendar with proper reminders — protects meetings from double-bookings the same way it protects carpools.
- **If you're just getting started:** do Steps 1–4 today and stop there. The rest can wait until Sunday.

Try it now

Open the calendar app on your phone, create one event called "Soccer practice — Mia" for next Tuesday at 16:00, set a reminder 90 minutes before, and share it on the new family calendar. That single event is your proof the system works — the rest is just more of the same.

Drafting a Basic Family-Tree Outline with AI: A Beginner's Guide

6 steps · Read online

By the end of this guide you'll have a plain-text family-tree sketch on your screen — parents, grandparents, siblings, cousins, the people you actually know — drafted for you by an AI assistant in about ten minutes. This is for anyone curious about family history who wants a simple starting point without filling in a full genealogy database on day one.

> **Honest note before you begin:** an AI assistant works in text, not in polished diagrams. What you'll get is a clean text outline (think of it like a typed list with indents and lines) that you can later paste into a free tree-drawing website if you want something visual. That's the real value of starting here — a tidy draft you can grow.

Before you start

- An account with any modern AI assistant — ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek, or similar. The free tier is fine.
- Five quiet minutes to memory-jog — names of parents, grandparents, aunts, uncles, and any cousins you remember. Dates and places aren't required.
- A notes app, a fresh document, or a notebook where you'll paste the final text.
- Roughly 10–15 minutes in total.

Step 1 — Jot down the relatives you actually know

Before you open any AI, take five minutes to list the people you're confident about. Grab paper or open a notes app. Write down: you; your parents and their first names; your grandparents on each side (maternal and paternal); your siblings and half-siblings if any; your aunts, uncles, and any cousins you remember. Don't worry about dates or places yet — names and relationships are enough to start a draft.

Example: "Me — Sara. Parents: Mom = Linda, Dad = Tom. Mom's parents: Gran = Joan, Pop = Bill. Dad's parents: Nan = Margaret, Grandpa = Alan. Siblings: brother Jake. Aunts/uncles: Aunt Mary (Mom's sister), Uncle Rob (Dad's brother). Cousins: Lily (Aunt Mary's daughter)."

If your memory is hazy on a branch, just leave it blank and add it later. You'll know it worked when you have a rough list — even a messy one is a perfectly good place to begin.

Step 2 — Open any AI assistant and start a fresh chat

Open the website or app of your chosen assistant — ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek, or whichever you already use. Sign in if asked, then start a brand-new conversation. A fresh chat matters because older messages can distract the AI. Look for a button like **"New chat"**, **"New conversation"**, or a small **pencil-and-paper icon**, usually near the top-left or top-right of the screen. Click or tap it once.

After the new chat opens, you'll see an empty text box waiting for your message. If your assistant offers different modes (like "Search" or "Thinking"), the default mode is fine for this task.

You'll know it worked when the screen shows a blank conversation with just the message box at the bottom and no previous replies hanging around.

Step 3 — Send a clear, structured prompt

Click or tap inside the message box and type a short request that includes your list from Step 1. Tell the AI you want a family-tree outline drawn with text — using indents, dashes, or a simple tree shape made of letters and lines. Mention that you're after a starting draft, not a finished genealogy chart, and that you'll fill in gaps later. Then send the message by pressing **Enter** or tapping the **paper-plane / up-arrow icon**.

Example: "Here's my family as I know it: me (Sara), Mom Linda, Dad Tom, brother Jake, Mom's parents Joan and Bill, Dad's parents Margaret and Alan, Aunt Mary (Mom's sister) and her daughter Lily, Uncle Rob (Dad's brother). Please draw a simple text-based family tree outline with three generations — me, parents, grandparents — using indents or a basic tree shape. Note any branches I haven't filled in. This is just a starting draft."

You'll know it worked when the AI replies with a formatted block of text showing people organized by generation, instead of just a flat sentence listing names.

Step 4 — Review the draft and ask for fixes

Read the AI's reply carefully. The first draft is almost never perfect — a name might be on the wrong line, a relative might be missing, or the formatting might be hard to read. Don't worry. Simply reply in the same chat and tell the AI what to change. Be specific: name the person, say where they belong, and how you want them shown. The AI works best with concrete instructions, not vague feedback like "make it better".

💬 Example: "Thanks. Please move Aunt Mary to sit under Mom's parents as their daughter, and add a cousin called Lily as her child. Also use dashes and greater-than symbols (├── └──) instead of dots so it copies cleanly into my document."

You'll know it worked when the AI sends a revised version that matches what you asked for. Keep tweaking until the structure feels right — two or three small rounds is normal.

Step 5 — Ask for an export-friendly version

Once the structure looks right, ask the AI for one final version that's easy to paste into a notes app, an email, or a text file. Mention you want plain text only — no special fonts, no images, no fancy code blocks that might break when pasted. Most assistants can produce a clean version in seconds. If you want a proper visual chart later, this text version can be copied into a free online tree-drawing site, but for now the text itself is your family-tree sketch.

💬 Example: "One more version please — same tree, but plain text only, no markdown or code blocks, so I can paste it straight into my notes app without fixing the formatting."

You'll know it worked when the AI's reply is plain text you can highlight, copy, and paste anywhere without formatting glitches or stray symbols.

Step 6 — Save your sketch and plan to grow it

Paste the final text into your notes app or a document and give the file a name you'll remember, like "Family tree — draft 1". This is your starting point. From here you can add more relatives over time — great-grandparents, in-laws, step-relatives, or children and nieces and nephews. Each time you add someone, return to the same AI chat (which still holds your earlier context) and ask for an updated version. The AI can also help you draft questions to ask older relatives, or summarize old letters and documents you find.

💬 Example: "Add these new people to the tree: my partner Sam, our daughter Mia, and Sam's parents Pat and Chris. Please show the in-law branch as a separate section so it doesn't get mixed up with my birth family."

You'll know it worked when you have a saved text file that grows each time you add to it, and you no longer have to keep everything in your head.

Common mistakes

- **Forgetting to start a fresh chat.** Older conversations can make the AI mix up unrelated topics. Always open a new chat for a new project.
- **Asking for too much in one go.** Asking for "every relative I might ever have, with dates and places" overwhelms the AI and gives a vague answer. Stick to what you actually know first.
- **Treating the AI's first draft as final.** It almost never is. Plan for two or three rounds of small fixes — that's normal and gives a much better result.

What this means for you

- **In everyday life:** A simple family-tree sketch makes it easier to talk to older relatives about family history — you can show them what you already know and ask them to fill the gaps.
- **For your business or work:** If you work in genealogy, heritage research, or local history, a quick AI-drafted outline is a fast starting point before investing in a proper database tool.
- **If you're just getting started:** Don't worry about being thorough on day one. A messy draft with the names you remember is a perfectly good place to begin — you can add detail whenever you like.

Try it now

Spend five minutes jotting down the names you know — parents, grandparents, siblings, one or two aunts or uncles. Then open any AI assistant you already use, start a new chat, and paste the example prompt from Step 3 with your own list filled in. You'll have your first family-tree sketch before your coffee goes cold.

How to Troubleshoot Your Home Wi-Fi Step by Step with an AI Assistant

6 steps · Read online

> **In short:** An AI assistant can walk you through a technician-style checklist (restart, isolate the device, test the line, check the router) and help you figure out what's wrong before you spend an hour on hold with your internet provider.

When your Wi-Fi drops in the middle of a video call or your phone refuses to load anything, the temptation is to call your internet provider and wait. Most home Wi-Fi problems, though, follow a handful of common patterns, and a chatbot like ChatGPT, Gemini, or Claude can walk you through the same diagnostic checklist a technician would use, in plain language, while you stand next to your router.

Before you start

- **An AI assistant you already use** — ChatGPT, Gemini, Claude, Copilot, or any other chatbot. No special account or paid plan is required for this task.
- **Your phone or laptop**, so you can type to the assistant and run the steps it suggests.
- **Access to your Wi-Fi router** (the small box your internet provider installed). You'll need to be near it for some steps.
- **Roughly 15–30 minutes.** Most fixes happen in the first few minutes; the rest is just confirming what works.

Step 1 — Tell the AI exactly what's wrong

Open your AI assistant and describe the problem in one short paragraph. Mention when it started, which devices are affected, and what you've already tried. The clearer you are, the better the questions and suggestions it can give you back. If it asks follow-up questions, answer them — this is the AI narrowing down the cause, the same way a technician would. Example: "My home Wi-Fi has been cutting out for the past two days. My laptop loses connection every 20–30 minutes, but my phone stays connected fine. I haven't changed any settings. Where should I start?" You'll know it worked when the assistant responds with a short list of likely causes or a first diagnostic step, instead of a generic answer.

Step 2 — Run the "is it just this device?" check

Ask the AI: "How do I tell whether the problem is with one device or with the whole Wi-Fi network?" It will usually walk you through picking up your phone and seeing if it has internet. If your phone works fine on Wi-Fi but your laptop doesn't, the issue is probably the laptop, not your network. If every device is slow or offline, the problem is more likely the router or your internet line. This one test saves a lot of wasted effort later. If your assistant doesn't ask, you can also try: "What should I check first to narrow down whether it's my device, my router, or my internet provider?" You'll know it worked when you can say out loud which of the three it probably is — device, router, or provider.

Step 3 — Restart the router the right way

Ask the AI to guide you through restarting your router properly. The correct method is: unplug the power cable, wait a full 60 seconds (so the internal memory clears), then plug it back in. Wait another two or three minutes for the lights to settle. Look for a button or a small recessed hole labeled "reset" — do **not** press that one, it wipes your settings. A regular power cycle fixes a surprising number of home Wi-Fi problems. If your router has a power button instead of a cable, hold it down until the lights go off, wait 60 seconds, then press it again. You'll know it worked when the lights on the router return to their normal pattern (usually one steady "internet" light) and your devices reconnect automatically.

Step 4 — Test the actual internet speed and signal

Ask the AI: "How do I check my current internet speed and Wi-Fi signal strength from home?" It will point you to a free speed test site (you can type something like "speed test" into any search engine to find one) and explain that you should run the test once near the router and once in the room where you usually use the internet. A big difference between the two numbers means the Wi-Fi signal is weak in that spot — often solved by moving the router higher, away from walls, or off a metal shelf. If both numbers are low even near the router, the problem is likely your provider's line. You'll know it worked when you have two speed numbers (one near the router, one where you usually sit) and a rough idea whether the slowdown is the signal or the line.

Step 5 — Ask the AI to interpret the results

Copy the numbers and what you observed back into the chat and ask, "Based on this, what should I try next?" This is where the assistant earns its keep — it can usually tell you whether to move the router, change the Wi-Fi channel, switch from the 2.4 GHz to the 5 GHz band, or accept that the line itself is the problem. If you see a phrase like "channel congestion" or "5 GHz band" for the first time, ask the AI to explain it in plain words before you change anything. A good chatbot will gladly translate its own jargon. You'll know it worked when the assistant gives you one or two concrete next actions, instead of a long list of possibilities.

Step 6 — Decide whether to call your provider

Ask one last question: "Based on everything we've checked, does this sound like something I can fix myself, or should I contact my internet provider?" Be honest about what you've found — the assistant can only help if you tell it the truth. If the speed test was bad right next to the

router, the Wi-Fi works on your phone but not your laptop, or every device has been down since a storm, those are useful details to share. When the AI says it's time to call, write down the two or three key things you tried, because the person on the phone will ask. You'll know it worked when you have either a working Wi-Fi connection, or a clear short summary you can read to the support agent so you don't have to start from zero.

Common mistakes

- **Pressing the tiny "reset" hole on the router by accident.** That wipes your settings and you'll need your provider's login details to set it up again. Power-cycle by unplugging instead, and only use the reset hole if the AI specifically tells you to.
- **Restarting only the router and not the device.** If your laptop still won't connect after the router is back, restart the laptop too. Stale network settings on the device are a common silent cause.
- **Trusting a speed test taken on Wi-Fi from far away.** A slow number in the far room tells you the signal is weak, not that your provider is slow. Move closer to the router before you blame anyone.
- **Letting the AI guess instead of asking.** If an answer feels vague, reply with more detail about what you actually saw — lights, error messages, devices affected. The clearer the input, the better the diagnosis.

What this means for you

- **In everyday life:** the next time your Wi-Fi drops during a movie, a game, or a work call, you can spend ten minutes running through this checklist before deciding whether to call support — and you'll know exactly what to say when you do.
- **For your business or work:** if you work from home, a slow or unreliable connection can quietly cost you hours a week. Having a fast self-diagnosis routine means fewer dropped calls and less time spent waiting on hold while deadlines slip.
- **If you're just getting started:** the single easiest first move is to bookmark one AI assistant on your phone, so the next time something breaks you can open it in seconds instead of searching blindly.

Try it now

Open your favorite AI assistant right now and type: "My Wi-Fi keeps dropping on one device. Can you walk me through how to figure out whether it's the device, the router, or my internet provider?" Read the first three suggestions it gives you — that's your personal starting checklist the next time something goes wrong.

Photos, video & music

Generate unique profile pictures for social media with Midjourney

7 steps · Read online

Want to make your online presence pop with a profile picture that's uniquely *you*? This guide will show you how to use Midjourney, an artificial intelligence (AI) tool, to design eye-catching images for all your social media profiles. By the end, you'll be able to create custom profile pictures or even social media banners that reflect your personality or brand, making a memorable first impression.

Before you start

- **A Discord account**: this guide uses Midjourney inside Discord, a free communication app where communities gather. Discord still works, but most people now use the web app at midjourney.com instead — if you prefer that, follow our web guide instead of this one. Either way you need the paid plan below.
- **A paid Midjourney subscription**: there is no free option. Midjourney ended its free trial in March 2023, so you cannot generate even one image without a plan. The cheapest tier is called Basic — check Midjourney's own pricing page for the current cost.
- **Device**: A computer, tablet, or smartphone with the Discord app installed or access to the Discord website.
- **Time commitment**: Allow about 15-20 minutes for your first few attempts to explore options and get comfortable.

Step 1 — Get set up in Discord

First, you need to join the Midjourney community on Discord. Discord is like a digital clubhouse where groups can chat, share, and collaborate. Midjourney operates as a **bot** (a computer program that lives inside Discord and responds to your commands) within its own dedicated **server** (the specific 'clubhouse' for Midjourney users). To begin, open your Discord app or visit the Discord website in your browser. Look for the "Explore Public Servers" option, often represented by a compass icon on the left sidebar. From there, search for "Midjourney" in the search bar and click to join its official server. Once you're in, you'll see many chat channels. Look for channels named "newbies," "newcomer-rooms," or "generate-channels" in the sidebar, as these are where new users can create images. If you don't immediately see "Explore Public Servers" or a compass icon, look for a '+' sign to "Add a Server" on the left sidebar, and then select "Join a Server" or "Explore Public Servers" from there. You'll know it worked when you see the Midjourney server listed on your left sidebar in Discord, and you can click into a "newbies" or "generate" channel.

Step 2 — Start your first image with a prompt

Now that you're in a generation channel, it's time to tell Midjourney what you want to create using a **prompt** – this is the text instruction you give to the AI. Think of it like giving a detailed request to an artist. In your chosen "newbies" or "generate" channel, type `/imagine` into the message box at the bottom. As you type, you'll see a pop-up with `/imagine prompt:` appear. Click on this, or press Enter. The word "prompt" will now be highlighted, waiting for your description. This tells the Midjourney bot you're ready to **generate** (create) an image. After you've typed your description, simply press Enter to send it. If it looks different, make sure you're typing `/imagine` in a chat channel, not in a private message to the Midjourney Bot yet. If the pop-up doesn't appear, try typing a space after `/imagine` before entering your text. You'll know it worked when your message `/imagine prompt: [your description]` appears in the chat, and the Midjourney bot starts processing your request (often showing a percentage as it works). 💬 Example: `/imagine prompt: a joyful wombat wearing a tiny chef hat, cartoon style, vibrant colours`

Step 3 — Crafting your profile picture prompt

To get a great profile picture, your prompt needs to be specific. Think about the style, subject, and any particular details you want. For profile pictures, a square **aspect ratio** (the width-to-height proportion of an image) is best. You can specify this with `--ar 1:1` at the end of your prompt. Consider who or what your profile picture is for: Is it for a professional network, a gaming profile, or a personal social media page? Describe the subject, their mood, the art style (e.g., "photorealistic," "vector art," "watercolour"), colors, and background. Adding details like "portrait," "headshot," or "close-up" can help Midjourney focus. You'll know it worked when you've carefully thought about the details and added `--ar 1:1` to ensure your image will fit a typical square profile picture slot. 💬 Example: `/imagine prompt: a happy person with short curly hair, wearing glasses, smiling gently, digital art portrait, soft pastel colours, simple background, --ar 1:1`

Step 4 — Reviewing your generated grid

After you send your prompt, the Midjourney bot will take about a minute to process it. What you'll see next is a 2x2 **grid** – a set of four distinct image options based on your prompt, all presented together as a single image. Underneath this grid, there will be several buttons. You'll notice buttons labeled U1, U2, U3, U4 and V1, V2, V3, V4. The 'U' buttons are for 'Upscale', meaning you want to make one of the four small images larger and more detailed. The 'V' buttons are for 'Variations', meaning you want to create four *new* images that are similar in style or composition to one you liked from the grid. Buttons U1 and V1 correspond to the top-left image, U2 and V2 to the top-right, U3 and V3 to the bottom-left, and U4 and V4 to the bottom-right. You'll know it worked when you see your 2x2 grid of four images appear in the chat, along with the U and V buttons below it.

Step 5 — Selecting and upscaling your favorite

Once you have your grid, it’s time to pick your favorite! Look closely at all four options and decide which one best suits your vision for a profile picture. If you like one of the images, click the corresponding 'U' button beneath the grid (e.g., U1 for the top-left image). This tells Midjourney to **upscale** that specific image, creating a larger, more refined version that's suitable for download. Midjourney will then process this request and present you with the single, higher-resolution image. If you prefer a slightly different take on one of the images, you can click a 'V' button instead to generate more variations. If it looks different and you don't see buttons, make sure you're looking at the most recent image grid generated by the Midjourney bot, as Discord chat moves quickly. Sometimes you may need to scroll up slightly. You'll know it worked when a single, larger version of your chosen image appears in the chat, often with new buttons like "Web," "Vary (Subtle)," or "Zoom Out" below it.

Step 6 — Saving your new profile picture

Once your chosen image has been upscaled, you'll see it displayed as a single, high-resolution picture in the chat. Now you can save it to your device! To save the image, simply click on it. This will open the image in a larger preview. From there, right-click (on a computer) or long-press (on a phone/tablet) and select "Save Image As..." or "Download Image." You can then choose where to save the file on your device. Once saved, it’s ready to be uploaded to your social media profiles! You can also click the "Web" button to view and download it directly from your Midjourney gallery in your web browser, which sometimes offers more options. If the options look different when saving, search for keywords like "Share," "Export," or "Save to Photos/Gallery" in the menu. You'll know it worked when the image file is successfully downloaded to your device, and you can open it in your photos or gallery app.

Step 7 — Refine your image or try again

Sometimes your first attempt might not be exactly what you envisioned. That's perfectly normal with AI art! If you want to refine an image you've upscaled, you can use the "Vary (Subtle)" or "Vary (Strong)" buttons below the upscaled image to create slightly different versions. If you want to make bigger changes, it's often better to start with a new `/imagine` prompt, adjusting your description with lessons learned from your previous attempts. Experiment with different keywords, art styles, or even add **negative prompts** (e.g., `--no blur` to avoid blurry images) to guide the AI more precisely. The Midjourney bot cannot directly edit specific parts of an image, so small tweaks often require generating new variations or starting fresh. The key is **iteration** – trying different prompts and variations until you get something you love. You'll know it worked when you successfully generate new variations or a completely new image that moves closer to your ideal profile picture.

Common mistakes

- **Not being specific enough in your prompt**: A vague prompt like "person" will give you very generic results.
  - **The fix**: Add descriptive details about appearance, mood, setting, and art style. Think about the expression, clothing, background, and overall aesthetic (e.g., "a determined young professional, wearing a navy blazer, standing in front of a blurred city skyline, warm lighting, photorealistic").
- **Forgetting the aspect ratio**: If you don't add `--ar 1:1` to your prompt, Midjourney might give you a rectangular image that doesn't fit well as a profile picture.
  - **The fix**: Always include `--ar 1:1` at the end of your prompt for square images.
- **Using the wrong Discord channel**: Trying to generate images in general chat channels or private messages to the bot (before enabling it for DMs) won't work.
  - **The fix**: Always use a designated "newbies" or "generate" channel on the Midjourney server, or if you've enabled it, in a direct message with the Midjourney Bot itself.

What this means for you

- **In everyday life:** You can create unique profile pictures for Instagram, Facebook, LinkedIn, or even dating apps that truly capture your personality and help you stand out from the crowd. No more struggling to find the perfect photo!
- **For your business or work:** You can design professional-looking headshots, team profile images, or creative avatars for internal communications or company social media pages, ensuring a consistent and polished brand image.
- **If you're just getting started:** Try generating a simple profile picture of yourself with a specific art style, like "cartoon portrait of a person with long brown hair, smiling, warm colors, --ar 1:1".

Try it now

Open Discord, go to a Midjourney generation channel, and type `/imagine prompt: a futuristic astronaut with a friendly face, vibrant colours, cartoon style, --ar 1:1` to create your very first unique AI-generated profile picture.

How to generate AI images in Picsart: a beginner's step-by-step guide

7 steps · Read online

> **In short:** Picsart has a built-in AI image generator that turns your text prompts into pictures, then lets you keep editing in the same conversation. The AI tools usually sit behind a paid plan, and the engine works best with still images — not video.

If you've ever opened Picsart to fix a photo, design a poster, or make a quick thumbnail, the app already has an AI image generator tucked inside. Type a sentence, and Picsart turns it into a picture. Don't like something in the result? Stay in the same chat and ask for a change — the AI keeps the rest of the image intact. Here's how to try it, step by step, even if you've never touched an AI image tool before.

Before you start

- A Picsart account — sign-up is free, but AI image generation usually requires a paid plan. Check Picsart's current pricing page before you start.
- The Picsart app on your phone or tablet, or the web version in your browser at picsart.com.
- A rough idea of what you want to make: a poster, a product mock-up, a photo edit, or a wallpaper.
- About 10 minutes for your first image.

Step 1 — Open Picsart and find the AI image generator

Open the Picsart app and sign in. Once you're on the home screen, look for a tile, button, or section whose name suggests AI image generation — common labels include **AI Image Generator**, **AI Playground**, or **Text to Image**. If you don't see any of those on the main grid, tap the + (plus) button at the bottom-center of the screen, or look for a **Create** tab along the bottom row of icons — the AI tools usually live in one of those two places. On the web version at picsart.com, look in the top menu for something like **AI Tools** or **Create**.

Example: this step has nothing to type — just finding the right entry point.

You'll know it worked when a screen appears with a text field where you can describe the image you want (usually called the prompt box) and at least one option to either start from text or upload a photo.

Step 2 — Pick your starting point: text, photo, or template

Inside the AI image tool, you'll usually see a few big choices near the top of the screen: **Text to Image** (start from a written description), **Image to Image** (start from a photo you upload), or a ready-made template you can fill in. Tap whichever fits what you want — for your first try, choose **Text to Image** if it's available. If the screen also shows a model selector — a small dropdown, a chip near the prompt box, or a tile in a side panel — pick whichever option is marked as the latest or recommended. Picsart rotates AI models over time, and the list changes often, so we won't name a specific one here.

Example: this step is a tap, not a typing step — the key choice is **Text to Image** for a first run.

You'll know it worked when the screen shows an empty prompt box and a clearly labeled button like **Generate**, **Create**, or an arrow icon.

Step 3 — Choose a canvas size (if the app asks)

Before you write your prompt, some versions of the tool show a size or shape picker. If you see it — usually a small icon shaped like a square, a rectangle, or a phone outline, or a dropdown near the prompt box — pick the shape that matches where you'll use the picture. A square shape works for Instagram-style posts, a tall rectangle works for phone wallpapers and Stories, and a wide rectangle works for banners and YouTube thumbnails. If the app also shows a quality choice (sometimes labeled with letters or numbers like 1K, 2K, or 4K, where higher means sharper but slower), pick the lowest one for a quick test so you don't burn through your credits.

Example: for a square post, choose the square option; for a phone wallpaper, choose the tall rectangle; if you see a quality choice, pick the lowest setting for a quick test.

You'll know it worked when a small preview, a number, or a labeled chip on the screen shows the shape and size you picked.

Step 4 — Write your prompt and generate the image

In the prompt box, describe the picture you want in plain, specific language. Name the subject, the mood, the colors, the lighting, and any text you want visible on the image — Picsart says its AI image generator handles on-image text reasonably well, which is useful for posters and diagrams. The more concrete you are, the closer the result will be to what you imagined. Then tap the **Generate** button (sometimes labeled **Create** or shown as an arrow icon). After a few seconds, one or more images appear on screen — pick the one you like best, or tap **Regenerate** if none feel right.

Example: "A friendly poster for a Saturday morning yoga class, soft pastel colors, with the words 'Saturday Morning Yoga' clearly readable in the center, hand-drawn style."

You'll know it worked when a finished image appears on screen and matches the canvas shape you picked in Step 3.

Step 5 — Refine the image in the same chat

Here's where the in-chat editing saves time: you don't have to start over when something's off. Stay in the same conversation and send a follow-up message describing one specific change. The AI keeps the rest of the picture intact and adjusts only what you asked for. You can chain several refinements one after another — change colors, swap a background, add a logo, rewrite the on-image text — without losing the previous version. Be specific: name the exact color, the exact word, the exact object. If you can't find the chat box, look for a field at the bottom of the screen that says something like **Message**, **Ask anything**, or **Describe a change**.

Example: after the yoga poster above, send: "Make the background soft pink, change the words to 'Bring a friend!', and add a small sun in the top-right corner."

You'll know it worked when the new image shows your specific change and the rest of the picture still looks like the previous version.

Step 6 — Edit an existing photo (optional)

If you'd rather start from a picture you already have, look for an **Upload** or **Image** button near the prompt box — usually shown as a small picture icon or a + sign. Tap it and pick a photo from your device. Once the photo appears, describe the edit you want in the prompt box. For finer control, look for a brush or markup tool near the image — common labels are **Mask**, **Edit**, or a small paint-bucket or brush icon. With that tool, drag your finger across the part of the photo you want changed; the painted area is the only part the AI will redraw. Then write a short note in the prompt box describing what should appear there. Be specific about what to remove, replace, or add.

Example: upload a group photo, then type: "Remove the person in the red shirt on the right." Or paint over the background with the brush tool and write: "Replace this background with a sunset beach."

You'll know it worked when the edited photo appears with your specific change applied and the rest of the image still looks natural.

Step 7 — Save and export your image

When you're happy with the result, look for a **Save**, **Download**, or **Export** button — usually at the top-right corner of the image or inside a ⋯ (three dots) menu. The app may ask you to pick a file format: choose **JPG** for everyday sharing, or **PNG** if you need a transparent background (for layering the image on top of other designs). The image then saves to your device's photo gallery or your computer's Downloads folder. If you want to keep editing the file inside Picsart later, look for a **Save to project** or **Save as draft** option instead.

Example: not a typing step — the key choice is **JPG** for everyday sharing, **PNG** only if you need a transparent background.

You'll know it worked when the image file appears in your phone's gallery (or your computer's Downloads folder) and opens like any normal picture.

Common mistakes

- **Picking the wrong starting mode.** If you upload a photo but meant to start from text (or the other way around), the result will look off. In Step 2, double-check you've chosen the right mode — **Text to Image** or **Image to Image** — before tapping Generate.
- **Asking for vague edits.** "Make it better" or "fix it" gives the AI nothing to act on. Always name the specific change: a color, an object, a word, a mood. The more concrete the instruction, the closer the result.
- **Forgetting the cost.** Each generation usually spends credits or counts against your plan's allowance. A higher-quality image costs more than a lower-quality one, and each follow-up edit in the chat counts as another generation. Check your remaining credits before kicking off a long chain of edits — otherwise the tool may stop mid-conversation.

What this means for you

- **In everyday life:** You can make a custom birthday card, a flyer for a school bake sale, or a fun wallpaper for your phone — in minutes, without any design training.
- **For your business or work:** You can mock up a quick social-media post, an event poster, or a product visual to share with a colleague before commissioning a designer for the final version.
- **If you're just getting started:** Begin with one simple text-to-image prompt (Step 4) at the lowest quality setting. Get a feel for what the tool does well before moving on to photo edits or higher-quality output.

Try it now

Open Picsart, find the AI image generator, and create one square image of your favorite season — one sentence prompt, no edits, just see what comes back. Takes about two minutes.

Auto-write captions for your photos using Hugging Face's free image captioning tool

7 steps · Read online

> **In short:** Hugging Face hosts free AI demos in its "Spaces" section, including image captioning tools that describe a photo in one short sentence. You upload an image, wait a few seconds, and copy the text — no coding needed.

By the end of this guide, you'll know how to turn a pile of unlabeled vacation or family photos into one-line descriptions, ready to drop into a social media post, paste as alt text, or use to search your own library.

Before you start

- A free Hugging Face account (you can sign up with an email, or sign in with a Google or GitHub account if you have one already).
- A modern browser on a laptop or desktop — Chrome, Edge, Firefox, or Safari all work. Phones work too, but a bigger screen makes uploading easier.
- 5–15 photos saved on your device, ideally as JPEG or PNG and under ~5 MB each so they upload quickly.
- Roughly 10 minutes for your first batch.

Step 1 — Create a free Hugging Face account

Open your browser and go to **huggingface.co**. Look for a **Sign up** button or link, usually near the top-right of the page. Type your email, pick a password, and confirm your inbox when the verification email arrives.

A free account lets you try any of the public demos. The page should land on your profile or the home feed once you're done.

Example: the verification email is usually titled something like "Confirm your Hugging Face account".

If it looks different: if you already have an account, click **Log in** instead of Sign up — and if your browser saved an old login, you may skip straight to the homepage.

You'll know it worked when you see your profile picture or initials in the top-right corner of the page.

Step 2 — Find an image captioning Space

From the top menu, click **Spaces** — think of a Space as a small web app built by someone in the Hugging Face community. In the search bar on the Spaces page, type **image captioning** and press Enter. You'll see a list of community-built demos; look for ones with clear names related to image captioning (if it's available, you might find options with names like **BLIP Image Captioning** or **ViT-GPT2 Image Captioning**).

Don't worry about picking the "best" one. Most of these demos do the same basic job: read your photo and write one sentence about it.

Example: searching for "BLIP image captioning" may return a Space built by Salesforce, which has been one of the popular captioning demos on the site.

If it looks different: if the search bar isn't obvious, try the **Filters** button, choose the **Computer Vision** category, and scroll until you spot a captioning demo.

You'll know it worked when you land on a page whose title mentions "Image Captioning" and shows an upload area or a button that says something like "Upload" or "Drop an image here".

Step 3 — Open the demo and wait for it to load

Click on the chosen Space to open it. The first time you visit, the demo may take 10–30 seconds to wake up — this is normal, since free demos often go to sleep when no one is using them. A progress bar, a spinner, or words like "Running" or "Starting" usually appear while it loads.

Once it's ready, you'll see an upload area and, on most Spaces, a short blurb explaining what the model does.

Example: a typical blurb is "A short image captioning model based on BLIP, fine-tuned on COCO" — COCO is just the name of a popular photo dataset the model learned from.

If it looks different: if you only see a gray box for more than a minute, refresh the page once. If it still won't load, try a different captioning Space from the search results in Step 2.

You'll know it worked when the upload area becomes active — usually it says "Drop image here" or shows a small image icon.

Step 4 — Upload your first photo

Drag a single photo from your desktop into the upload area, or click the area to open a file picker and choose your picture. The photo appears as a small preview on the page. Don't upload dozens at once — most free demos are built to caption one image at a time.

Wait a few seconds. The caption will appear in a text box or label below the image.

💬 Example: a beach sunset photo might return something like "a sunset over the ocean with orange clouds", and a family group shot might return "a group of people smiling and standing together".

If it looks different: if nothing happens after about 30 seconds, look for a **Submit** or **Run** button near the upload area — some Spaces require you to confirm the upload before they process it.

You'll know it worked when a short English sentence describing your photo appears in or near the upload area.

Step 5 — Try a few more photos to see the range

Pick 3–5 different photos — a landscape, a portrait, a pet, a city street, a close-up of food — and run them through the demo one at a time. This helps you see what the tool handles well and where it gets confused. You'll likely notice it's strong at obvious scenes (beaches, dogs, kitchens) and weaker at very specific or unusual images.

💬 Example: a clear photo of a coffee cup on a wooden table might come back as "a cup of coffee on a table" — accurate, but plain.

If it looks different: don't be surprised if the caption is generic ("a person standing outside") for busy scenes. That's the model playing it safe rather than guessing at details it can't see clearly.

You'll know it worked when you have a small list of captions that roughly match each photo's main subject.

Step 6 — Copy and store your captions

Highlight the caption text in the demo's output box and copy it with **Ctrl+C** on Windows or **Cmd+C** on Mac. Paste it into a notes app, a Google Doc, or a spreadsheet right next to the photo's filename. Repeat for each photo. There's no built-in batch export — the work is mostly copying and pasting one at a time.

💬 Example: a simple spreadsheet column might read `beach_sunset.jpg → "a sunset over the ocean with orange clouds"`.

If it looks different: if the caption text isn't selectable, right-click the output area and choose **Copy** from the menu instead.

You'll know it worked when you have a tidy list of one-line descriptions you can reuse later.

Step 7 — Use your captions in real life

Now do something useful with them. For social media posts, paste a caption under your photo and tweak the wording if you want more personality. For alt text (the short description screen readers use for images on websites), paste it directly into the image's alt-text field. For organizing your own photo library, drop the captions into a folder named "_captions.txt" so you can search them later.

💬 Example: a LinkedIn post becomes "Wrapped up a long week with a walk on the beach 🌅 — a sunset over the ocean with orange clouds. Felt good to breathe."

If it looks different: alt-text fields look slightly different on every site — on a WordPress blog it's under "Image attributes", on X or Twitter it's the small alt field you fill in while uploading.

You'll know it worked when you stop staring at unlabeled photos thinking "what was this one again?"

Common mistakes

- **Uploading giant RAW files straight from a camera.** Free demos often struggle with files over ~10 MB, and some reject RAW formats entirely. Convert to JPEG first, or resize using your phone's built-in photo editor.
- **Expecting one upload to caption a whole album at once.** Most demos only handle one image at a time. Captioning 200 vacation photos means 200 separate uploads — plan for it to take a while, or pick a few favorites first.
- **Trusting the caption blindly.** Image captioning AI can mislabel things — it might call a labrador a "golden retriever" or skip a person in the background. Always glance at the caption before publishing or pasting it as alt text.

What this means for you

- **In everyday life:** captioning a personal photo for the first time feels like giving your camera a tiny voice. You can finally find that one sunset photo from three summers ago by searching the captions.
- **For your business or work:** if you run a blog, a small shop, or any website with lots of photos, alt text is something search engines and accessibility tools really do look at. Auto-captioning gives you a starting point in seconds — you then edit for accuracy and tone.

- **If you're just getting started:** skip the jargon and just play. Upload one photo of your pet, laugh at whatever the AI writes, and you've already learned the basics.

Try it now

Pick one favorite photo from your camera roll, open huggingface.co in your browser, search for "image captioning" in Spaces, upload the photo, and copy the caption. Two minutes, one photo, one sentence — that's the whole loop.

How to remove a photo background cleanly with Meta AI on your phone

7 steps · Read online

> **In short:** Open Meta AI in WhatsApp, Messenger or Instagram, upload a photo, and ask it to "remove the background." Meta AI will usually generate a transparent version you can save as a PNG. You'll get a clean cutout in about a minute, though fine details like frizzy hair or sheer fabric can still need a touch-up.

You already chat with Meta AI in WhatsApp, Messenger or Instagram. This guide adds one more thing you can do there: turn any photo into a clean cutout (the photo with the background taken away) and drop it onto a new background, all from your phone, no Photoshop, no extra app. The whole workflow takes about five minutes once you've done it once.

Before you start

- **The app:** Meta AI built into **WhatsApp, Messenger, or Instagram** on your phone (iPhone or Android). If you can already chat with Meta AI in any of these, you're set.
- **A photo on your phone:** JPEG or PNG, ideally with a clear subject (a person, a product, a pet) and decent lighting. Busy, low-contrast backgrounds are the hardest to cut.
- **Roughly 5–10 minutes** for the first try, less after that.
- **Free.** No paid plan is required for the workflow below. Meta may rate-limit (cap how many times you can ask in a short window) on busy days, so have a backup photo ready.
- **What Meta AI can and can't do here:** it can usually remove a simple background in one go, but it is not pixel-perfect. Fine hair, glasses, and see-through materials can leave a faint halo (a soft glow around the edges of the cutout). It's great for social posts and product shots, less great for printing at poster size.

Step 1 — Open a Meta AI chat

Tap the **circular blue/purple Meta AI icon** — in WhatsApp it's usually at the top of your Chats tab or behind a **blue "AI" button** in the bottom-right; in Messenger it's the top of the search bar; in Instagram it's in the search/DMs area. If you don't see it, look for a **sparkle or four-pointed star icon** or check the **app's top menu**. The chat screen opens, showing a short welcome message and the message box at the bottom. You'll know it worked when the box at the bottom says something like "Message Meta AI" or shows a send arrow.

Step 2 — Upload the photo you want to cut out

Tap the **paperclip or "+" icon** next to the message box to attach a file. From the menu that appears, pick **"Gallery"** (Android) or **"Photo & Video Library"** (iPhone) and choose the photo. The image appears inside the chat as a small preview above the message box. If your app doesn't show a paperclip, look for a **camera icon** instead — the option is the same. Don't send the message yet; the prompt you add in the next step shapes what Meta AI does with the image.

Example: choose a clear, well-lit photo of a friend standing in front of a plain wall, or a product on a tabletop.

You'll know it worked when the photo preview is sitting above the message box, ready to go with your next sentence.

Step 3 — Write the prompt that asks for the cutout

Tap the message box and type a short, direct instruction. Be specific: say "remove the background" rather than just "edit this." If you already know where you want to use the image, mention that too — it helps the model (the AI "brain" that reads your words and produces the answer) pick a sensible default background color.

Example: `Remove the background from this photo and give me the subject on a transparent background.` Or, if you'd like a new background in the same step: `Remove the original background and place the person on a soft white studio background.`

Tap **send** (the paper-plane or arrow icon). Meta AI will usually reply within 10–30 seconds with a new version of your image. You'll know it worked when a new picture appears in the chat that looks like your subject on a flat color or transparent-looking background.

Step 4 — Check the edges before you save

Pinch to zoom in on the edges around the subject — especially around hair, fingers, glasses, and anything thin. Look for a faint color halo (a thin glowing outline) where the old background is bleeding through, or a slightly hard "stamped" edge. If the cut is clean, move to Step 5. If you see issues, stay in the same chat and ask for a fix in plain words.

Example: `Zoom in around the hair — I can still see a faint gray halo. Please clean up the edges.` Or: `The fingers on the right hand look chopped. Please redo with finer edge detail.`

You'll know it worked when the second image shows noticeably cleaner edges than the first, with no obvious halo or missing fingers.

Step 5 — Save the new image to your phone

Long-press the new image Meta AI produced. A small menu appears with options like **"Save," "Save to Photos,"** or **"Download."** Tap the save option. The image lands in your phone's Photos or Gallery app, usually in a folder called **Recents** or **Downloads**. If long-press doesn't show "Save," look for a **three-dot ⋯ menu** on the image instead, or tap the image to open it full-screen and then tap the **download/save icon** in the corner.

You'll know it worked when the new image appears in your Photos app as its own file, separate from the original you uploaded.

Step 6 — Drop the cutout onto a new background

Open any simple editor on your phone — the built-in **Photos** app on iPhone, or **Google Photos** on Android both work. Open the saved cutout, then tap **Edit**, and look for an option like **"Add Photo"** or **"Overlay / Sticker"** that lets you put a second image behind the first. Pick your new background — a sunset you took, a logo, a solid color. Resize and reposition, then save.

Example: a beach sunset from your camera roll, so your friend looks like they're standing on vacation.

You'll know it worked when the final saved image shows your original subject, with clean edges, sitting naturally on the new background you chose.

Step 7 — Keep your reusable prompt in one easy-to-find place

Don't retype the prompt every time. Save the working version from Step 3 in your phone's **Notes app** or in the **Notes** feature of WhatsApp/Messenger itself, then copy-paste it next time. For even faster reuse, ask Meta AI in the same chat: `Save this as a reusable prompt I can paste next time.` Meta AI will usually reply with a tidy, copy-ready version you can keep.

Example: `Remove the background from this photo and give me the subject on a transparent background. Keep the subject exactly the same — only the background should change.`

You'll know it worked when, next time, you can paste the saved prompt, attach a photo, and get a clean cutout in under a minute.

Common mistakes

- **Asking vaguely.** "Edit my photo" is too open. Always name the change: "remove the background," "place the subject on white," "tighten the edges." A specific instruction gets a specific answer.
- **Uploading a tiny, blurry, or low-contrast photo.** If the subject and the background are similar colors (a black dog on a dark couch), Meta AI will struggle. Pick a photo with a clear difference between subject and background.
- **Saving the wrong image.** The chat shows both your original upload and the new generated image. Long-press the *new* one, not the one you sent. If you saved the original by mistake, just repeat Step 3 and save the reply.

What this means for you

- **In everyday life:** Clean cutouts are perfect for WhatsApp or Instagram stickers of your pet, a quick profile picture with a softer background, or a fun "I went to Paris" photo where you actually went to the kitchen. Five minutes of fiddling beats a desktop app you don't have time to learn.
- **For your business or work:** If you sell a few products on a side hustle, you can knock out consistent product photos for a catalog or marketplace listing without paying a designer. The edges won't be magazine-perfect, but they're good enough for a small storefront and free.
- **If you want to go further:** Try a two-step prompt in one chat — first remove the background, then ask `Now place the subject on a marble countertop with soft daylight from the left.` The model remembers the previous image in the same conversation, so each new instruction builds on the last one instead of starting over.

Try it now

Open Meta AI in WhatsApp, Messenger or Instagram right now. Pick a photo of a person or object that is already on your camera roll, paste this prompt, and send:

`Remove the background from this photo and give me the subject on a transparent background.`

Save the result to your phone. That's a complete, reusable workflow you can do again in under a minute next time.

Create a multi-room music playlist with Alexa+

7 steps · Read online

> **In short:** With Alexa+ you can set up a voice routine so a single phrase — like "Alexa, play my dinner playlist" — starts the same music in every room at once. It uses the Alexa app's Routines feature.

Imagine saying "Alexa, play my dinner playlist" and the same music starts in the kitchen, lounge and bedroom at once. You'll learn how to set up one playlist that plays across multiple rooms with a single voice command.

Before you start

- An **Amazon account** (free).
- The **Alexa app** installed on your phone or tablet (Android or iOS).
- At least one **Amazon Echo device** (like an Echo Dot, Echo Show or Echo Studio) in each room.
- A **music streaming service** already linked to Alexa (Amazon Music, Spotify, Apple Music, etc.).
- Roughly **15 minutes** of your time.

Step 1 — Open the Alexa app and go to Devices

Open the Alexa app on your phone or tablet. At the bottom of the screen you'll see a row of icons — look for the one that looks like a **house outline** (this is the Devices section). Tap it. You should see a list of all your Echo devices, often grouped by room.

You'll know it worked when the screen shows your rooms and devices.

Example: Look for the house-shaped icon at the bottom of the Alexa app and tap it.

Step 2 — Pick a room to start

Choose the room where you want the music to begin. Tap the room name (for example, "Kitchen"). Inside that room view, look for a menu option — this is typically a **three-dot menu (⋯)** in the top-right corner, though on some app versions it may appear as a **gear icon** or a **menu with three lines** instead.

You'll know it worked when a menu slides up or opens from the bottom of the screen.

Example: Open the Kitchen room and look for the ⋯ menu in the top-right corner.

Step 3 — Find the routine option in the menu

From the menu that just appeared, look for an option to add this room to a routine — the exact wording may be **"Add to Routine"** or something similar depending on your app version. Tap it.

If you're not sure which option it is, look for any menu item that mentions "routine" or "routine" in English — this is the feature that links your voice phrase to your music action.

You'll know it worked when a routine setup screen opens.

Example: Tap **Add to Routine** (or the equivalent option in your app version).

Step 4 — Choose a voice trigger

In the routine setup screen, look for the option to set a **voice trigger** — on many versions this is labeled **"Voice"** and shown with a microphone icon. Tap it. A text box appears where you can type the phrase you want to say to Alexa to start the playlist.

If the voice option isn't visible, look through the list of available trigger types for something related to speaking or voice commands.

You'll know it worked when you see a text box labeled something like "Enter your phrase" and a way to confirm and move to the next step.

Example: Tap the voice trigger option and then tap inside the text box.

Step 5 — Type your wake-up phrase

Type a short phrase you'll say to Alexa, for example: **"Alexa, play dinner playlist"**

Confirm the phrase and move to the next step. Alexa may show a summary of the phrase you've entered.

You'll know it worked when the summary confirms your exact phrase.

Example: Type "Alexa, play dinner playlist" and confirm it.

Step 6 — Add the music action

In the routine setup, look for an action type related to music — this may be labeled **"Music"** (often with a musical note icon) or something similar depending on your app version. Tap it. A new screen opens where you can pick the playlist and the devices that should play it.

If you don't see a music option right away, look through the list of available actions for anything about playing audio or music.

You'll know it worked when you see options to choose a playlist and to choose which devices or rooms should play it.

Example: Tap the music action and then tap the option to choose a playlist.

Step 7 — Pick your playlist and devices

Select the playlist you want to play everywhere. Then choose all the rooms where you want the music to play (for example, Kitchen, Lounge, Bedroom) — look for a device or room picker. Save the routine.

You'll know it worked when the routine summary shows your chosen playlist and all the rooms you selected.

Example: Select "Dinner Vibes" playlist and tick Kitchen, Lounge and Bedroom.

Common mistakes

- **Wrong room selected** — If the music only plays in one room, double-check Step 2 and make sure you picked the correct room before saving the routine.
- **Phrase already used** — If Alexa says the phrase is taken, add a word like "my" or change the wording slightly (for example, "Alexa, start dinner playlist").
- **Device not selected** — If a room is missing, go back to Step 7 and make sure every room is ticked before saving.

What this means for you

- **In everyday life:** Set up a "Good morning" routine that plays your favorite news podcast in the bedroom and kitchen at the same time.
- **For your business or work:** Create a "Focus time" routine that plays lo-fi beats in every meeting room whenever you say "Alexa, start focus time".
- **If you're just getting started:** Start with one room and one playlist; once it works, add more rooms in the same routine.

Try it now

Open the Alexa app, go to Devices, pick a room, tap the menu icon → look for a routine option → choose voice trigger → type your phrase → add a music action → pick your playlist and devices → Save. Say the phrase out loud to test it.

Sharpen Your Photos with AI Using Upscayl

6 steps · Read online

Do you have a cherished photo that's just a little bit blurry, or an old snap you wish looked sharper and ready for a larger print? This guide will show you how to use the free AI image super-resolution application, Upscayl, to quickly enhance your images. By the end, you'll be able to give your photos a professional, crisp finish, making them poster-worthy in minutes.

Before you start

- **Install the free Upscayl application:** You'll need to download and install Upscayl on your computer. It's an open-source tool, which means it's developed by a community and freely available for use.
- **A computer with a dedicated graphics card (GPU) is recommended:** While Upscayl can run on most systems, a graphics card (like an NVIDIA, AMD, or Intel Arc GPU) will significantly speed up the processing of your images, making the enhancement much faster.
- **A photo to enhance:** Choose a digital image that's slightly blurry or low-resolution that you'd like to improve.
- **Time:** Allow around 5-7 minutes to follow the steps and enhance your first image. Upscayl is free to use with no hidden costs.

Step 1 — Launch Upscayl

First, open the Upscayl application on your computer. Look for its distinctive icon on your desktop, in your Start menu (Windows), Applications folder (macOS), or program launcher (Linux), and double-click it. When it launches, you'll typically see its main window, often featuring a clean interface with a central area to load images and various options organized on the side. This is your starting point for transforming your photos. **If it looks different:** If it's your very first time launching Upscayl, you might see a quick setup or welcome screen, possibly asking to confirm default file paths or to download initial AI models. Follow any clear on-screen prompts to complete this setup before proceeding. **You'll know it worked when:** The Upscayl application window fully appears on your screen, indicating it's ready for you to select an image and begin the enhancement process.

Step 2 — Load Your Image

Now it's time to bring in the photo you want to improve. In the Upscayl window, you'll usually find a prominent button labeled "Open Image" or a clear "Drag and Drop" area in the center of the screen. Click the "Open Image" button, and your computer's file browser window (like File Explorer on Windows or Finder on macOS) will pop up. Navigate to where your blurry photo is stored on your computer, select it with a click, and then click "Open." Your chosen image will then appear in the main Upscayl window, usually displayed as a clear preview. **If it looks different:** If the image doesn't load immediately, an error appears, or the preview area remains empty, ensure your photo is in a common format like JPG, PNG, or WEBP. You might also try restarting the application and attempting to load the image again. **You'll know it worked when:** Your selected photo is clearly visible within the Upscayl application's main preview area, confirming it has been loaded successfully and is ready for enhancement. 💬 Example: Select a file named `old_family_photo.jpg` from your `Pictures` folder.

Step 3 — Choose an AI Model

With your photo loaded, you need to tell Upscayl which "recipe" or AI model to use for the enhancement. Think of an AI model as a specialized algorithm (a set of computational instructions) that the AI follows to intelligently reconstruct and add detail to your image. Each model is trained differently and can produce slightly varied results depending on the original image. On the left side of the Upscayl window, you'll see a "Models" section with a list of AI options like "Real-ESRGAN," "Remacri," or "Ultramix_Balanced." Click on a model to select it; for general photos, "Real-ESRGAN" or "Ultramix_Balanced" are excellent starting points known for good all-round performance. **If it looks different:** If you don't see many models listed, or options are greyed out, you might need to download additional models within Upscayl's settings. Look for a "Download Models" button or an option to manage models in the application's configuration or preferences menu. **You'll know it worked when:** The name of the model you clicked is clearly highlighted or appears as the active selection, indicating that Upscayl will use this specific AI algorithm for processing.

Step 4 — Select Upscale Factor

Next, you'll decide how much larger you want your image to become. The "upscale factor" determines how many times the image's width and height will be multiplied, significantly increasing its total number of pixels and its overall resolution. For example, a 2x factor makes the image twice as wide and twice as tall (quadrupling the total pixels), while a 4x factor makes it four times wider and four times taller (sixteen times the total pixels). Below the models section, you'll typically find options for "Upscale Factor," often presented as buttons like 2x, 3x, 4x, or sometimes a custom size. Select "4x" for a significant improvement in detail and resolution for most images, providing a good balance of quality and processing time. **If it looks different:** Some AI models are specifically designed for particular upscale factors, so if an option like "4x" is greyed out, try selecting "2x" or "3x" instead. Alternatively, you might need to switch to a different AI model (as described in Step 3) that supports a higher factor. **You'll know it worked when:** The "4x" option (or your chosen upscale factor) is clearly selected or highlighted in the Upscayl interface, confirming your desired output size.

Step 5 — Start the Enhancement

Once your photo is loaded and you've selected your preferred AI model and upscale factor, it's time to let the AI do its work. Look for a prominent button, usually positioned at the bottom right of the Upscayl window, labeled "Upscayl" or "Enhance." Click this button to begin the super-resolution process. Depending on your computer's processing power (especially if you have a dedicated graphics card like an NVIDIA RTX GPU) and the original image size, this can take anywhere from a few seconds to a minute or two. You might see a progress bar, a spinning icon, or a message indicating that the AI is actively processing your image. **If it looks different:** If the "Upscayl" or "Enhance" button appears greyed out or inactive, double-check that you've successfully loaded an image and made definite selections for both the AI model and the upscale factor. All required options must be chosen to activate the enhancement button. **You'll know it worked when:** The progress indicator disappears, and Upscayl either displays the enhanced image in a side-by-side view or automatically prompts you to save it.

Step 6 — Review and Save Your Enhanced Image

After the AI has finished processing, you'll see your newly enhanced image. Take a moment to compare it with the original; Upscayl often provides a side-by-side comparison or an easy toggle to switch between the two. Notice the improved sharpness, clarity, and detail in areas that were once blurry or pixelated. If you're pleased with the result, Upscayl will usually prompt you with a "Save As" dialogue box, which is your computer's standard file explorer. Choose a location on your computer to save your crisp new poster-ready photo, and give it a descriptive name so you can easily find it later. **If it looks different:** If a "Save As" prompt doesn't automatically appear, look for a "Save" or "Export" button within the Upscayl application itself. Some versions might also automatically save the enhanced image to a default folder (like "Pictures" or "Downloads") if no prompt is explicitly given. **You'll know it worked when:** You have a new, sharper, and higher-resolution version of your photo saved as a separate file on your computer, ready for use. While Upscayl can significantly enhance your photos, it cannot perfectly invent details from an extremely blurry or tiny original image that were never present in the first place, but it can significantly improve clarity.

Common mistakes

- **Not downloading AI models:** Upscayl relies on various specialized AI models to perform its magic, and not all are included in the initial application installation.
  - **The fix:** If you see few or no options in the "Models" section (Step 3), go into Upscayl's settings or preferences and look for a "Download Models" or "Manage Models" option to acquire more AI algorithms.
- **Choosing the wrong AI model:** Different models are optimized for different image types (e.g., general photos, illustrations, specific art styles). Using a model designed for, say, anime on a family portrait might give unnatural or distorted results.
  - **The fix:** Experiment! If an image looks strange after processing, try re-loading the original and selecting a different model. "Real-ESRGAN" or "Ultramix_Balanced" are generally good choices for everyday photos and realistic images.
- **Expecting miracles from extremely low-quality sources:** While AI super-resolution is powerful, it cannot invent information that simply isn't there in the original image. An image that is completely out of focus, heavily pixelated, or tiny from the outset will still have significant limitations even after enhancement.
  - **The fix:** Aim for input photos that are only moderately blurry or low-resolution. The better the starting material (even if just slightly), the more impressive and natural the AI's final result will be.

What this means for you

- **In everyday life:** You can rescue cherished old photos that are a bit blurry, prepare family pictures for larger prints or digital display, or make your social media images stand out with professional-level sharpness and detail.
- **For your business or work:** Enhance product photos for e-commerce listings, sharpen images for presentations or marketing materials, or improve scans of documents and graphics to make them clearer and more legible for your audience and colleagues.
- **If you're just getting started:** Pick one of your favorite slightly blurry photos – perhaps a holiday snap or an old family picture – and follow this guide. Start by selecting the "Real-ESRGAN" model and a 4x upscale factor. Seeing the transformation firsthand is the best way to understand the power of AI image enhancement.

Try it now

Grab a photo that needs a little boost, launch the free Upscayl application, and follow these steps to give it a fresh, crisp look in the next few minutes.

How to create and edit images in ChatGPT

6 steps · Read online

By the end of this guide, you'll be able to type a simple description and get a unique image from ChatGPT, then adjust it to suit your needs – all without needing any complicated design software. This is perfect for anyone who wants a quick visual for a blog post, a social media banner, or a personal creative project. Just remember, while ChatGPT can generate fantastic images, it can't perform complex photo manipulation like a dedicated graphic design program.

Before you start

- **An active OpenAI account:** You'll need to sign up at `openai.com` if you don't already have one.
- **Access to ChatGPT:** You can use the ChatGPT web interface (in your browser) or the official ChatGPT mobile or desktop app. Image generation is available to both free and paid users, though free users may have usage limits.
- **A basic internet connection:** ChatGPT runs online, so no extra software downloads are required, but a stable internet connection is essential.
- **Rough total time:** This guide should take about 5-10 minutes to follow.

Step 1 — Start a new conversation for image generation

Open ChatGPT and begin a fresh conversation. Image generation is built directly into ChatGPT's main interface, so you don't need to access any special mode or tool — simply type your image request in the chat.

What you see: The standard ChatGPT chat interface, with a text input box at the bottom of the screen. **Action:** Either start a new chat (click the "New chat" button) or use an existing conversation. There's no need to find a special image mode or custom GPT — just type your description directly. **What happens next:** ChatGPT will recognize when you want to generate an image based on your prompt and will process your request. **If it looks different:** If ChatGPT responds with text instead of an image, make sure your prompt clearly describes what you want to see (e.g., "Generate an image of..." or "Create a picture showing..."). The more specific you are, the better. **Example:** "You click 'New chat' and immediately type: 'Generate an image of a peaceful cottage garden.'" **You'll know it worked when:** You see a message like "Generating image..." or a progress indicator, followed by the image appearing in the chat.

Step 2 — Write your image description (prompt)

Now it's time to tell ChatGPT what picture you want. The instruction you give the AI is called a **prompt**, which is like giving a briefing to a human artist. Be specific about the main subject, desired style, colors, and any particular details.

What you see: A clear text input field. **Action:** Type a detailed description of the image you want ChatGPT to create. Mention key elements, the overall mood, color palette, and artistic style. You can start with phrases like "Generate an image of..." or "Create a picture showing..." **What happens next:** Once you press Enter or send your message, ChatGPT will process your request and begin generating the image. This usually takes a few seconds. **If it looks different:** If ChatGPT returns text instead of attempting to create an image, double-check that your prompt is clear and explicitly asks for an image. Try starting with "Generate an image of..." **Example:** `Generate an image of a peaceful, sun-drenched cottage garden in the American countryside, filled with native wildflowers, a small birdbath, and a maple tree in the background, painted in an impressionistic art style with soft, warm colours.` **You'll know it worked when:** You see a message like "Generating image..." or a progress bar indicating that the AI is working on your visual.

Step 3 — Review the generated image(s)

After a short wait, ChatGPT will present the image(s) it created based on your prompt. Take a moment to look at them closely.

What you see: One or more small images (thumbnails) displayed directly within your chat conversation. **Action:** Click on an image to view it in full size and examine the details. Think about what you like and what could be improved. **What happens next:** The image will expand, usually showing more detail. You might also see options to download or provide feedback. **If it looks different:** Sometimes ChatGPT might generate a single image instead of multiple options, or occasionally, if your prompt was unclear, it might offer text suggestions for refinement instead of images. If the image seems completely off, it might be worth going back to Step 2 and rewriting your prompt entirely. **Example:** "You notice the maple tree looks more like an oak. You also see a small download icon, usually a downward arrow or three dots, when you hover over the image." **You'll know it worked when:** Your chat window displays one or more visual results of your prompt, and you can click them to view them larger.

Step 4 — Refine with an edit request

It's rare to get a perfect image on the first try. Luckily, you can easily ask ChatGPT to make adjustments. You don't need to rewrite the whole prompt; just tell it what to change.

What you see: The previously generated image(s) still visible in the chat, and the active text input box below. **Action:** Type a new, concise prompt requesting specific changes to the *previous* image. Start with a phrase like "Edit the image to..." or "Make the tree taller..." **What happens next:** ChatGPT understands you want to modify the existing image and will generate a new version incorporating your feedback, usually keeping the overall composition while applying your tweaks. **If it looks different:** If ChatGPT generates an entirely *new* image that doesn't seem related to the previous one, your refining prompt might have been too broad, or it might have lost the context. Try being more direct, e.g., "Adjust *the last image* by..." 💬 **Example:** `Edit the image to make the maple tree in the background taller and add a small, rusty tin shed near the cottage.` **You'll know it worked when:** A new image appears in the chat, showing the requested adjustments applied to the original concept.

Step 5 — Save your final version

Once you're happy with the results, it's time to save your unique creation to your device.

What you see: The final, refined image displayed in the chat. **Action:** Hover your mouse over the image (on desktop) or tap the image (on mobile) to reveal a download icon. This is often a downward-pointing arrow or a three-dot menu icon that opens more options, including 'Download'. Click or tap this icon. **What happens next:** Your browser or app will prompt you to save the image file to your computer or phone's downloads folder. You might be able to rename the file before saving. **If it looks different:** Some versions might require a right-click (desktop) or a long-press (mobile) on the image itself to bring up a context menu with a "Save image" or "Download" option. 💬 **Example:** "You click the downward-pointing arrow in the top-right corner of the image, and a file named 'image.png' appears in your browser's download bar." **You'll know it worked when:** The image file is successfully saved to your device, and you can open it with your usual photo viewer.

Step 6 — Understand usage guidelines

As you create images, it's important to be aware of OpenAI's usage policies. These rules help keep the platform safe and responsible.

What you see: No specific visual change, but your actions adhere to the platform's terms. **Action:** Be mindful of the content you ask ChatGPT to create or edit. Avoid prompts for anything that could be considered explicit, harmful, misleading, or illegal. **What happens next:** By following these guidelines, you ensure your images are generated successfully and your account remains in good standing. **If it looks different:** If your prompt violates the policy, ChatGPT will usually refuse to generate the image and might give you a message explaining why. Sometimes it might just generate a blank or highly abstract image. 💬 **Example:** "You avoid asking for images of specific copyrighted characters or content that promotes hate speech." **You'll know it worked when:** Your image is generated without any warning messages or content refusal from ChatGPT, allowing you to use it freely (within its licensing terms).

Common mistakes

- **Vague or short prompts:** Providing too little detail often leads to generic or unexpected results that don't match your vision.
  - **Fix:** Treat ChatGPT like an artist's brief. Include the subject, style, colors, mood, and any specific elements. For example, instead of "a cat," try "a fluffy ginger cat sleeping on a sunny windowsill, drawn in a cartoon style with vibrant colors."
- **Ignoring policy warnings:** Attempting to bypass content filters can lead to image generation being blocked or, in persistent cases, account warnings.
  - **Fix:** Always respect the content policies. If an image is flagged, try rephrasing your request to be more appropriate, or choose a different subject.
- **Not being specific enough about edits:** When refining an image, vague requests like "make it better" may not produce the desired result.
  - **Fix:** Be specific about what you want to change — for example, "make the background darker" or "add more flowers in the foreground."

What this means for you

- **In everyday life:** Easily create custom invitations for a family gathering, design a unique background for your phone, or illustrate a personal story without needing to be an artist.
- **For your business or work:** Quickly generate social media graphics, blog post headers, unique presentation slides, or marketing visuals, saving time and potentially design costs.
- **If you're just getting started:** The easiest first move is to open ChatGPT and type a one-sentence description of something you love – like "Generate an image of a cute koala hugging a eucalyptus leaf, in a colorful watercolour style."

Try it now

Spend the next two minutes opening ChatGPT, and typing a simple prompt such as **"Generate an image of a cup of coffee on a wooden table, sketch style"**. Download the result and see how fast a quality picture can appear.

How to Sort and Clean Up a Messy Phone Photo Library with Gemini

6 steps · Read online

> **In short:** Gemini can look through photos you upload, flag near-duplicates, group similar shots, and pull readable text out of old screenshots — all without touching your library. You stay in control: Gemini suggests, you decide what to delete.

By the end of this guide, you'll know how to use Gemini as a careful second pair of eyes for your photo mess — the kind where you have seven nearly-identical sunset shots and a folder full of screenshots you can never search. This is for anyone whose phone gallery has gotten out of hand and wants a low-risk way to start sorting through it.

Before you start

- A free Google account (any Gmail address works). The free Gemini tier is enough for small batches; a paid Google AI plan lets you upload more photos at once and keeps longer conversations in memory — check Google's official site for current plan options.
- The Gemini app on your phone, **or** a browser pointed at the Gemini website.
- About 20–30 minutes for your first batch.
- A folder of photos you actually want to clean up. Pick one small group to start with — don't try the whole library on day one.

Step 1 — Open Gemini and start a fresh chat

Tap the Gemini app icon, or visit the Gemini website in your browser. If you're not signed in, sign in with your Google account. Look for a **"New chat"** button — usually a plus sign or a small pencil icon at the top of the chat list — and tap it. A blank conversation opens. Starting fresh keeps the AI focused on the photos you're about to share, not on a leftover question from yesterday's weather check.

If you don't see "New chat," try the **three-line menu** (sometimes called a hamburger icon) in the top corner — the option to start over is usually inside it. On some devices the same button looks like a small **square with a pencil** inside instead.

Example: tap **New chat**, then wait for the empty text box to appear.

You'll know it worked when you see a clean text box with no messages above it and a paperclip or upload icon somewhere nearby.

Step 2 — Gather a small batch of photos to upload

Don't try your whole library on day one. Pick **10 to 20 photos** that feel like they might overlap — the seven sunset shots from your last trip, the burst-mode photos from your kid's birthday, the screenshots you saved last month. In your phone's gallery, tap **Select** (or hold one photo until checkboxes appear), choose the photos, then tap **Share**, and pick **Gemini** from the list of apps. On a computer, drag the photos straight into the chat box.

If the share menu doesn't show Gemini, copy the photos to your clipboard and paste them into the chat. On older Android phones the upload icon looks like a **paperclip** 📎; on iPhones it's often a small **plus sign +** at the bottom of the text box.

Example: open your gallery, select the 8 beach photos from last weekend, tap **Share → Gemini**.

You'll know it worked when the photos appear as small thumbnails in your message and Gemini replies with something like "Got it — I can see 8 photos."

Step 3 — Ask Gemini to flag duplicates and similar shots

Now comes the useful part. In the same chat, type a short, clear question about the photos you just shared. The clearer your instruction, the better the answer. You can ask Gemini to find photos that look almost the same, group them by scene, or just describe what makes each one different. It can read the image content directly — the colors, the framing, who's in the shot — and reason about which ones are worth keeping.

If Gemini's response is too generic, follow up with a more specific question. After a vague "they all look similar," try "Which one is sharpest?" or "Which one best shows both people smiling?" If the free tier says it's reached a limit, work in even smaller batches, or look into a paid plan for larger uploads.

Example: "Here are 8 photos from the same sunset. Can you group any that look like near-duplicates, and tell me which one of each group is the sharpest or best composed?"

You'll know it worked when Gemini gives you a numbered list or short table, like "Group A: photos 1, 3, and 5 are nearly identical — keep #3."

Step 4 — Pull searchable text out of an old screenshot

Screenshots are useless until you can search them. Take a single screenshot — a recipe, an address, a confirmation number, a Wi-Fi password you saved months ago — and share it with Gemini the same way you did in Step 2. Then ask Gemini to read the text out of it and rewrite it in a clean,

searchable form. This works for receipts, WhatsApp messages, store hours, anything that ended up trapped inside an image.

The text won't be added back to your photo library automatically — you'll copy it out yourself in Step 6. And Gemini can misread small or blurry text, so for anything important like medical info or legal documents, always double-check the original screenshot.

Example: "Please transcribe the text in this screenshot. Format it as: store name, date, total amount, and any reference number."

You'll know it worked when Gemini replies with readable text you can copy and paste, and that text matches what you actually see in the screenshot.

Step 5 — Ask for a simple cleanup plan

Once Gemini has analyzed a batch, ask it to summarize what it found into a short, action-friendly list. You're looking for a "do this, then this, then this" answer — not a long essay. The point is to turn Gemini's observations into a checklist you can actually use in your phone's gallery. This is also where Gemini's "context window" (the amount of text and images it can keep track of in one conversation) matters: if your batch is too big, it may start forgetting earlier photos.

If the plan is too long, ask for the top three actions only. If it's missing something obvious, ask a follow-up like "Did you see any photos that look like receipts or notes I should transcribe?"

Example: "Based on what you saw, give me a short checklist of which photos I can safely delete and which I should keep. Keep it under 5 bullet points."

You'll know it worked when you have a short, scannable list that feels doable in under five minutes.

Step 6 — Take the plan into your Photos app

Now switch to your phone's gallery or Google Photos. Open the same batch of photos. Use Gemini's checklist as your guide: delete the ones it flagged as duplicates, keep the sharpest of each group, and save the transcribed text somewhere searchable — a notes app, a contact entry, or a dedicated "Screenshots I've read" album. This is the step where the cleanup actually happens, and it's the step only you can do.

For an extra safety net, turn on your gallery's recently-deleted folder so any mistake can be undone within 30 days. On Google Photos it's usually at **Library → Trash**; on iPhone it's at **Albums → Recently Deleted**.

Example: open the 8 beach photos, delete 5 based on Gemini's notes, star the best one, and paste the transcribed receipt text into your notes app.

You'll know it worked when your batch is smaller, the photos you kept are the ones you actually want, and the old screenshots have become searchable text.

Common mistakes

- **Uploading your whole library at once.** Gemini gets less accurate with very large batches, the conversation becomes hard to follow, and you may hit free-tier limits fast. Work in groups of 10–20.
- **Letting Gemini "decide" what to delete.** Gemini can suggest, but it can't actually delete anything from your phone — and it shouldn't. Always make the final call yourself, in your own gallery, with your finger on the button.
- **Trusting transcribed text blindly.** Especially for numbers, dates, and addresses — Gemini reads most of it right, but a single wrong digit on a confirmation code can cost you a refund. Always glance back at the original screenshot.

What this means for you

- **In everyday life:** That folder of 200 mystery screenshots from the last three years can finally become a few searchable notes — addresses you can find, receipts you can expense, articles you can actually read.
- **For your business or work:** If you snap whiteboards, business cards, or reference material on the go, this same flow turns them into text you can paste into documents, spreadsheets, or your CRM.
- **If you're just getting started:** Pick **one** folder — say, screenshots from last month — and run the full six steps on just that folder. Twenty minutes later, you'll have a feel for the process before touching anything bigger.

Try it now

Open Gemini right now, share it the last five photos you took, and ask: "Which of these is the best one, and why?" You'll get a tiny but real taste of how Gemini sees your library — and you'll have deleted nothing by accident.

Safely Rename a Batch of Photos with a GitHub Copilot Script

6 steps · Read online

> **In short:** With GitHub Copilot inside VS Code, you can describe a renaming task in plain English and get a working Python script back — no coding required. The trick is to test on copies first so your originals stay safe.

By the end of this short guide, you'll have asked Copilot to write a small Python script, tested it on copies of your photos, and renamed an entire folder without typing a line of code yourself. If you have a folder of vacation shots named `IMG_4829.jpg`, `IMG_4830.jpg`, and so on, this is for you.

Before you start

- **A computer** running Windows, macOS, or Linux.
- **VS Code installed** — the free code editor from Microsoft (download from `code.visualstudio.com`).
- **The GitHub Copilot extension** installed in VS Code, with a Copilot account signed in. Copilot offers a free tier and paid plans; check the official GitHub page for current availability and which features each tier includes.
- **Python installed** on your computer (download from `python.org`). Python is the language your script will be written in — think of it as plain English instructions for your computer.
- **A folder of photos** you'd like to rename, plus enough free disk space to make a copy. The whole task takes about **15 minutes**.

Step 1 — Make a working copy of your photos

Right-click your photo folder in your file manager, choose **Copy**, then **Paste** somewhere easy to find, like your Desktop. Rename the copy to something like `photos-test`. This is the folder you'll experiment on.

On macOS, you can also select the folder and press Option-drag to duplicate it. On Windows, hold Ctrl while dragging. You'll see a second folder appear next to your originals, with all the same files inside.

If it looks different: If your photos live in cloud storage like Google Drive or iCloud, the same Copy/Paste idea works — just make sure the copy finishes syncing before you start renaming anything. Some file managers hide the "Paste" option inside an Edit menu; look there if right-click doesn't show it.

You'll know it worked when you see two folders: your untouched originals and a clearly named test copy.

Step 2 — Open the test folder in VS Code

Open VS Code, then go to **File → Open Folder…** and pick your new `photos-test` folder. VS Code's left sidebar (called the **Explorer**) now lists every photo in the folder.

A workspace in VS Code is simply the folder you opened — think of it as putting that folder "on your desk" so the editor can see everything inside it. Nothing is moved or copied; VS Code is just looking.

If it looks different: On the VS Code welcome screen, there's a big **Open Folder** button you can click instead. If the Explorer sidebar is hidden, use **View → Explorer** to bring it back.

You'll know it worked when the names of your test photos appear in the left sidebar.

Step 3 — Open the Copilot chat panel

In VS Code, open the **Copilot Chat** panel. The keyboard shortcut is **Ctrl+Alt+I** on Windows and Linux, or **Cmd+Shift+I** on macOS. A chat window appears — usually on the side of the window.

Copilot is an AI coding assistant: you describe what you want in plain language, and it suggests code. You don't have to understand the code right now — that's the whole point of having an AI helper.

If it looks different: If the shortcut doesn't open anything, look for a small chat-bubble icon in the VS Code sidebar, or use the menu **View → Copilot Chat**. Some older versions put Copilot under **File → Preferences** instead.

You'll know it worked when a chat box is visible with a blinking text cursor, ready for your message.

Step 4 — Describe your renaming task in plain English

Click into the chat box and type what you want. Be specific about the folder, the file types, and the new naming pattern. Copilot reads your message and proposes a Python script in the chat.

The more clearly you describe the rule, the better the script will be. If your rule is "rename to `vacation-001.jpg`, `vacation-002.jpg`, …", say that exactly. If you want the date taken included, ask for that too.

Example: `Write a Python script that renames every .jpg and .png file in the current folder to "vacation-2024-001.jpg", "vacation-2024-002.jpg", and so on. Skip files that already match this pattern.`

If it looks different: Copilot might ask a follow-up question like "What folder should I use?" — just answer in the same chat. If it offers several code blocks, pick the one that looks closest to what you asked for and ignore the rest.

You'll know it worked when a code block appears in the chat with a Python script inside it.

Step 5 — Save the script and run it on the test folder

When Copilot shows a script, look for a small **"Insert into new file"** button above the code (or a similar option like **"Apply"** or **"Copy"**). Click it, and VS Code creates a new file. Save this file inside your test folder as `rename.py`.

Now open VS Code's built-in terminal from the top menu: **Terminal → New Terminal**. A text panel slides up from the bottom. Type `python rename.py` and press Enter.

The terminal is a text-based way to talk to your computer directly — instead of clicking icons, you type commands and read text replies.

If it looks different: If `python` isn't recognized, try `python3 rename.py` instead — that's more common on macOS and Linux. On Windows, you may need to install Python from `python.org` and tick **"Add Python to PATH"** during setup for the command to work.

You'll know it worked when the terminal shows your photos being renamed one by one, and the file list in the Explorer updates.

Step 6 — Verify, then run the same script on your real folder

Open the test folder in your regular file viewer (Finder, Explorer, or Files) and check the new names. Open a few photos to make sure none are broken — the picture files themselves shouldn't change, only their names.

If the names look right, you have two safe options: copy your `rename.py` script into your real photo folder and run it there, or ask Copilot to adjust the script so it points to a different folder. Either way, your originals stay untouched.

If it looks different: If a photo won't open, the file may have been partially affected — restore it from your test copy or your originals. If the names look wrong, go back to Step 4 and ask Copilot to fix the rule, then re-test on the test folder before touching real files.

You'll know it worked when your real photo folder shows the clean new names, and every photo still opens normally.

Common mistakes

- **Skipping the test folder.** It's tempting to point the script straight at your real photos, but one wrong rule can scramble hundreds of filenames at once. Always run the script on copies first, then on the real folder only after you're sure the names are right.
- **Running the script twice without changing the names.** The second run will keep building on your filename pattern (`vacation-2024-001.jpg` could become `vacation-2024-001-001.jpg`). Either move the renamed files out, or include the "skip if already named" rule in your prompt.
- **Trusting Copilot's first answer without reading.** Copilot is helpful but not perfect. If the script does something you didn't ask for — like deleting files or changing the file extension — rewrite your prompt to be more specific and try again on the test folder.

What this means for you

- **In everyday life:** A folder of 200 vacation photos can go from chaos to `paris-day1-001.jpg`, `paris-day1-002.jpg`, and so on in under a minute. No more guessing which `IMG_4829.jpg` was the Eiffel Tower shot.
- **For your work or business:** Product photos, invoice scans, and event pictures often arrive with cryptic camera names. A two-minute script saves the headache of clicking through 80 files one by one, and the same approach scales to thousands of files.
- **If you're just getting started:** Don't worry about learning Python itself — you're learning how to ask an AI for a script. The skill that matters here is describing your rule clearly, not writing the code by hand.

Try it now

Open your test folder in VS Code, open the Copilot chat, and paste this exact message.

`Write a Python script that lists every .jpg file in the current folder and renames them to "sample-001.jpg", "sample-002.jpg", in the order they appear. Print the new names before doing it.`

Run it on a tiny folder of 3 to 5 test files. Seeing those photos rename themselves is the moment this stops feeling like magic and starts feeling like a tool you control.

How to Turn a Long Podcast Transcript Into Shareable Notes With Claude

6 steps · Read online

> **In short:** You can paste a podcast transcript into Claude and ask for a 5-bullet summary, the best quotes, and an action list — usually in under a minute. You'll need the transcript as text (Claude can't listen to audio), and very long ones may need a paid plan.

A podcast you actually enjoyed drops, and now you want to share the best parts without making your friend sit through 90 minutes. Or you covered it for work and need notes by tomorrow. Here's how to use Claude to turn that wall of text into something a busy person will actually read.

Before you start

- A Claude account (the free tier works for shorter transcripts; very long ones may need a paid plan)
- Access to **claude.ai** on a web browser or the Claude mobile app
- A transcript of the podcast as plain text — either from the show's website, a transcription service, or your own export
- About 10 minutes total

Step 1 — Get the transcript as plain text

Before Claude can help, you need the transcript as plain text — words you can copy and paste. Many podcasts publish transcripts on their website or in the show notes. If yours doesn't, run the audio through a transcription service (there are free and paid options) and export the result. A simple .txt or .docx file works best; if you only have a .vtt file with timestamps, that's fine too — Claude can skip past the timecodes. If you only have the audio file and no transcript yet, pause here and get one first; Claude reads text, not sound. Example: a clean opening line — "Host: Today we're talking about how small teams ship faster than big ones." You'll know it worked when you have the spoken words sitting in a document or on your clipboard, ready to paste in the next step.

Step 2 — Open a fresh chat in Claude

Go to **claude.ai** on your computer, or open the Claude app on your phone. Look for a button that starts a brand-new conversation — it usually looks like a pencil icon, a plus sign, or carries a "New chat" label, often near the top-left of the screen. Click or tap it once, and the page clears to an empty conversation with a text box at the bottom. A fresh chat matters because mixing your podcast transcript with old questions can muddle the summary later. If you don't see a clear "new chat" option, try logging out and back in, or closing the browser tab and reopening it. Example: no typing yet — you're just clearing the deck. You'll know it worked when the chat shows an empty message thread and a blinking cursor in the input box.

Step 3 — Paste the transcript and tell Claude what it is

Click into the empty text box, paste the full transcript with your keyboard (Ctrl+V on Windows, Cmd+V on Mac), and then add one short line underneath that explains what you're sharing. That line matters — without it, Claude doesn't know it's looking at a podcast and may try to edit the text as if it were your draft. If your transcript is very long and the paste doesn't fit, look for a paperclip or attachment icon near the text box and upload the file instead; files count toward the same content Claude can read. Example: after pasting, add this on a new line — "The text above is a podcast transcript. Please read it carefully before I ask my next questions." You'll know it worked when Claude replies with something like "Got it — I've read the transcript. What would you like to know?" or simply waits for your next instruction.

Step 4 — Ask for a 5-bullet summary

Stay in the same chat and type a follow-up request for a summary. Ask for a specific number of bullets and a specific audience — vague prompts get vague summaries. Claude usually answers in under a minute for a one-hour transcript, though very long ones may be too much for Claude to read all at once on the free tier. If the answer stops mid-sentence, just reply "keep going" and Claude will resume from where it left off. Example: "Now please summarize this podcast in exactly 5 bullets for a busy teammate who hasn't listened. Lead with the single most useful idea." You'll know it worked when you see five short, self-contained bullets that someone could read in 60 seconds and actually understand the episode.

Step 5 — Pull the best quotes and an action list

Ask for two more things in the same chat: the most quotable lines from the transcript, and a short list of action items — concrete things the listener might actually do after hearing the episode. Ask for them as separate, clearly labeled sections so you can copy each one cleanly. If a quote feels bland, ask for a more surprising or specific one. If the action list sounds generic, ask Claude to tie each item back to a moment in the transcript. Example: "Now give me two more sections: (1) the 3 most quotable lines, with who said them; (2) a 4-item action list for someone who wants to apply this tomorrow. Label each section clearly." You'll know it worked when you have three distinct blocks — summary, quotes, action list — that you can copy out separately.

Step 6 — Save and share the notes

Copy the three blocks out of the chat and paste them into your notes app, an email, or a Slack message. One honest note: every quote Claude pulls is something the speaker actually said in the transcript you gave it — but if you paraphrase a quote before sending, double-check it against the source, because Claude can occasionally mis-attribute a line. If you're on a paid plan, look for a "Share" option that gives your recipient a link to the whole conversation instead. Either approach works — pick the one your teammate will actually open. 💬 Example: "Summary: [paste 5 bullets]. Best quotes: [paste 3 lines]. Action list: [paste 4 items]. Source: [episode name and link]." You'll know it worked when you hit send and your recipient replies with something other than "what is this?"

Common mistakes

- **Pasting without telling Claude what the text is.** If you just drop in a transcript and ask "summarize this," Claude may treat it as your own writing and try to edit it. Always add one short line identifying it as a podcast transcript first.
- **Asking for too much in one prompt.** "Summarize this, pull quotes, give me five tweet ideas, and write a LinkedIn post" usually produces shallow output for everything. One clear request per turn gets much better results.
- **Trusting the quotes without checking.** Claude is reading your transcript, so the quotes are real — but if you rewrite them in your own words before sending, you can accidentally change the meaning. For important quotes, copy them verbatim and verify.

What this means for you

- **In everyday life:** instead of telling a friend "you have to listen to this episode," you can send the three best quotes and one sentence on why it matters — and they will actually read it.
- **For your business or work:** a 90-minute interview becomes a 60-second read for your team in roughly 10 minutes of work, transcribing included.
- **If you're just getting started:** this is one of the easiest first projects with Claude. The transcript gives the assistant a clear job to do, and you'll see useful output within minutes.

Try it now

Find the transcript of one podcast episode you've already listened to (anywhere from 20 minutes to an hour long), paste it into a fresh Claude chat, and ask for the 5-bullet summary using the prompt from Step 4. That single action is the whole first run — and it usually takes less time than making a cup of coffee.

How to Summarize a Long YouTube Video in 60 Seconds with Any AI

6 steps · Read online

> **In short:** Paste a YouTube link into any AI assistant and ask for a summary — in about a minute you get the main points, key quotes, and a one-sentence takeaway without watching. The catch: not every assistant reads YouTube links directly, so we cover that fallback too.

Found a one-hour interview you don't have time to watch? A 40-minute lecture you need for work tomorrow? You don't have to skip it. Any modern AI assistant can read the page a YouTube link points to and pull out the important parts. Here's the whole workflow, start to finish.

Before you start

- **A YouTube link** — the URL of the video you want summarized. It works on desktop, mobile, or in the YouTube app.
- **An AI assistant** you already have access to — ChatGPT, Gemini, Claude, Copilot, Le Chat, or any other mainstream option. Free or paid plans both work, though some assistants may need a paid plan to read external links reliably.
- **Roughly 2–3 minutes** for the whole thing, including reading the summary.

Step 1 — Copy the YouTube video link

Open the YouTube video you want summarized and look at the address bar at the top of your browser (or the link shown under the video title in the YouTube app). A YouTube link starts with **youtube.com** or **youtu.be** and ends with a jumble of letters and numbers — that's the video's unique ID. Tap and hold (on phone) or click and drag (on computer) to highlight the whole address, then copy it. On the YouTube mobile app, tap **Share** under the video and choose **Copy link** instead — it's faster. If you're on a TV or game console and don't see a copy option, open the video in a phone or computer browser to grab the link from there. 💬 Example: Your copied link will look something like youtube.com/watch?v=AbCdEfGh123 — the letters and numbers at the end are the video's ID. You'll know it worked when the address is highlighted or your phone briefly says "Copied."

Step 2 — Open your AI assistant

Open the website or app for the AI you normally use. Most readers already have at least one of these: ChatGPT, Gemini, Claude, Copilot, or Le Chat. Open a new chat — look for a button like **New chat**, **New conversation**, or a + or pencil icon in the corner of the screen. Starting fresh matters because a long previous chat can confuse the assistant about what you want right now. You'll know it worked when you see an empty message box, usually at the bottom of the screen, with a prompt that says something like "Message…" or "Ask anything" — your exact wording may be different.

Step 3 — Paste the link and ask for a summary

Click in the empty message box and paste the YouTube link you copied. Then type your request. A short, clear instruction works best — tell the assistant what you want and how long the answer should be. Press **Enter** (or tap the send arrow). The assistant will fetch the video's public page, read its title and description, and (on most assistants) the transcript too. A short summary appears in under a minute. 💬 Example: "Here's a YouTube link: [paste your link]. Summarize this video in 5 bullet points and keep it under 150 words." You'll know it worked when the assistant replies with bullet points or a short paragraph that names the video's main topic and the people or ideas involved.

Step 4 — Pull out a few key quotes

The summary gives you the gist, but sometimes you want the actual words someone said — a striking line, a statistic, or a memorable phrase. Ask for those next. Quotes are useful when you want to share the video with a friend or quote it in your own work. Type a follow-up question that specifically asks for verbatim lines. The assistant will return them with quotation marks and usually say who said them. 💬 Example: "Now pull out 3 key quotes from that video, with the speaker's name for each one. Keep each quote to one sentence." You'll know it worked when the reply contains lines wrapped in quotation marks and the speaker is named beside each one.

Step 5 — Get a one-sentence takeaway

A bullet-point summary is great for studying, but a single sentence is easier to remember and share. Ask for the takeaway last. The assistant will compress everything it just read into one line you can use as a reference or send to a colleague. If you have a specific audience in mind — your team, your teacher, your family — say so, because the assistant will word the takeaway differently for each. 💬 Example: "In one sentence, what's the single most important takeaway from that video, written so I can send it to my manager?" You'll know it worked when the reply is a single complete sentence (not a paragraph) that captures the main idea.

Step 6 — Save or share your summary

Now that you have the points, quotes, and takeaway, decide where they should live. Most assistants have a **Copy** button at the top or bottom of each reply — tap it to grab the text and paste it into your notes app, email, or document. Some assistants also have a **Share** button that creates a

link you can send to someone. If you don't see either button, you can manually select the text and copy it the same way you copied the YouTube link in Step 1. You'll know it worked when the summary text is sitting in another app — your notes, an email draft, or a message — and you can read it again without opening the AI chat.

Common mistakes

- **Pasting the link without any instructions.** The assistant may reply with "What would you like me to do with this?" — always tell it you want a summary, in what length, and for whom.
- **Assuming every assistant reads YouTube links directly.** A few assistants only read the video's title and description, not its full transcript. If your summary feels thin or vague, ask the assistant to look up the video's transcript, or copy the transcript yourself from YouTube (look for an option like **Show transcript** under the video description) and paste that text instead of the link.
- **Trusting the summary on a topic you know nothing about.** Summaries are usually accurate, but AI can still mix up names, dates, or numbers. For important decisions — anything medical, legal, financial, or that affects other people — verify the quote or fact by checking the original video yourself.

What this means for you

- **In everyday life:** Next time someone sends you a 30-minute tutorial or a long podcast episode on YouTube, decide in 60 seconds whether it's actually worth your time. Skip the noise, keep the useful stuff.
- **For your business or work:** Before a meeting, paste the link to a recorded presentation or webinar and ask for the five main points plus the speaker's recommendations. You walk in prepared without sitting through the whole recording.
- **If you're just getting started:** Pick one YouTube video you've been putting off, follow the six steps above, and see how much you actually retain from the summary versus the full video. Most people are surprised how often the bullet points are enough.

Try it now

Open your AI assistant right now, paste the link to any YouTube video you've been meaning to watch, and type: "Summarize this video in 5 bullet points and pull out 2 short quotes." You'll have a usable summary in about a minute — that's the whole skill.

Sharpen Old Videos with NVIDIA AI Super-Resolution Online

5 steps · Read online

Do you have cherished family memories or important old clips that look blurry and dated on modern screens? This guide will show you how to use a free online AI tool to take those low-quality videos and make them look crisp and clear again, breathing new life into your digital archives. You'll be able to transform your footage without needing any special software or a powerful computer.

Before you start

- **A short video clip:** Find a low-resolution video file from your computer or phone. For your first test, a clip under 10 seconds and in a common format like MP4 works best.
- **A modern web browser:** Use an up-to-date browser like Chrome, Edge, Firefox, or Safari on your computer or tablet.
- **A stable internet connection:** You'll be uploading and downloading video files, so a reliable connection helps.
- **A free online account (if prompted):** While many NVIDIA AI demos are free to try instantly, some might ask you to sign up for a basic guest or developer account. This is usually free and quick.
- **Rough total time:** Expect to spend about 5-10 minutes on your first video, including setup and processing.

Step 1 — Find the online AI demonstration

To begin, you need to locate the special online space where NVIDIA showcases its innovative AI tools for anyone to try. Open your web browser and search for "NVIDIA AI Video Super Resolution online demo" to find the dedicated demonstration page. A **model** here is like a very smart computer program trained to do a specific task, such as identifying objects or, in this case, making videos look better. The **Super-Resolution** model uses AI to cleverly guess and draw in missing details, making a blurry image sharp, rather than just stretching it, which would only make it blurrier.

You'll usually see a welcoming page describing the AI's capabilities, sometimes with example videos. If you're prompted to create a free account, follow the instructions to sign up quickly using your email address. If it looks different, you might land on a general NVIDIA AI page; if so, look for a "Demos" or "Labs" section, then specifically for "Video Super Resolution" to find the right tool. You'll know it worked when you see the main demonstration area with clear options to upload a video, or you've completed any required sign-up.

Example: Open your web browser and search for "NVIDIA AI Video Super Resolution online demo". Look for a link from an official NVIDIA website or a reputable AI demonstration platform like Hugging Face or Replicate, often with "NVIDIA" or "Super Resolution" in the title.

Step 2 — Upload your low-resolution video

Once you're on the demonstration page, you'll find a workspace designed for you to interact with the AI. This area often includes some sample videos that show off what the AI can do, but you want to test your own. Locate the designated area on the page that allows you to add your video file. You'll typically see a large button or a dotted box labeled "Upload your video," "Choose File," or "Drag and drop video here."

Clicking this will usually open your computer's file browser window, allowing you to locate and select your video. Once selected, the file will begin uploading to the online platform. If it looks different, some sites might use a "drag and drop" zone instead of a button; in this case, simply drag your video file directly from its folder on your computer into that specified area. You'll know it worked when your video file name, a small preview image of your video, or a progress bar indicating the upload status appears on the screen.

Example: Look for a button or box labeled "Upload your own video," "Choose File," or "Drag and drop video here." Click it, select your short video clip (e.g., an MP4 file under 10 seconds), and click "Open" or "Upload."

Step 3 — Choose your upscaling settings

Now it's time to tell the AI how much sharper and larger you want your video to become. Select the desired upscaling factor to tell the AI how much larger and clearer to make your video. The interface will typically offer different "upscaling factors," which determine how many times bigger the AI will make your video. For instance, "2x" would double its size and detail, while "4x" would make it four times larger and significantly more detailed. High-definition video, like **4K**, refers to a very clear picture quality with many more pixels (the tiny dots that make up an image) than older, standard-definition videos, making it look much sharper on modern screens.

Look for a section, often a drop-down menu or a set of radio buttons, labeled "Upscaling Factor," "Resolution," or "Quality." Choosing an option, like "4x," will usually update the display, showing your selection as active, and might even provide an estimated output resolution. If you don't see exact "2x" or "4x" options, look for quality presets like "High," "Ultra," or "Best quality." You'll know it worked when the selected upscaling factor or quality level is highlighted, checked, or clearly displayed as your current choice on the demonstration page.

Example: On the demonstration page, find the settings panel or enhancement options. Select "4x Upscaling" or "Ultra Quality" to get the maximum possible detail out of your old video file.

Step 4 — Run the AI process

With your video uploaded and the desired quality settings chosen, the AI is ready to perform its task. Initiate the AI processing to transform your video. Locate a prominent button, usually at the bottom or side of the workspace, labeled "Process," "Generate," "Start Upscaling," or simply "Run." This step triggers the **inference** process – a technical term for when the AI program applies its learned knowledge (its "training") to your video to enhance it. The AI will analyze each frame, cleaning up blurry edges, smoothing out visual "noise" (graininess), and enhancing colors.

Clicking this button will send your video and settings to NVIDIA's AI servers for processing. The screen will typically display a progress bar, a percentage counter, or a "Processing..." message, indicating that the AI is actively working on your footage. If it looks different, some online demos might show a spinning icon, a small animation, or simply a text update like "Enhancing video..." without a detailed progress bar. This process can take anywhere from a few seconds to a few minutes, depending on your video length and the selected upscaling quality. You'll know it worked when the processing indicator disappears, and the page transitions to show the results, usually presenting your original and enhanced videos side-by-side.

Example: Click the "Generate," "Process," "Start Upscaling," or "Run" button to begin the transformation of your video.

Step 5 — Review and download your new video

After the AI has finished its work, the page will present the results. Examine the AI's results and download your newly enhanced video. The screen will now display your original video next to the AI-enhanced version, often playing them simultaneously or allowing you to toggle between them. You'll likely notice the improved clarity, sharpness, and reduced blur in the enhanced footage. This visual proof often highlights the significant improvement made by the AI.

Below or beside this comparison, you'll find a clear button to save your work. Clicking this button will start the download of the high-resolution, upscaled video file to your computer's default downloads folder. If it looks different, some platforms might only show the enhanced video playing in its own window with a single "Download" or "Save" button. This tool can significantly improve video quality, but it cannot invent entirely new scenes or information that wasn't subtly present in the original footage. You'll know it worked when your web browser indicates a file download is in progress, and you can locate the new, larger video file in your computer's "Downloads" folder.

Example: Review the side-by-side comparison of your original and enhanced videos. Once you are satisfied, click the "Download," "Save," or "Export" button to save your new, high-quality video to your computer.

Common mistakes

- **Uploading long videos:** These free online tools are primarily designed for quick tests or short clips, not full-length movies. Uploading long footage will likely cause the process to fail or time out.
  - **The fix:** Always crop your video down to a few seconds (e.g., 5-10 seconds) first using a basic video editor on your phone or computer before uploading.
- **Expecting miracles from extremely dark or damaged footage:** AI needs some visual information to work with. If your original video is very dark, pixelated, or heavily damaged, the AI might make incorrect "guesses," leading to strange or unnatural visual patterns.
  - **The fix:** Start with videos that, while low-resolution, still have visible shapes and colors. AI can enhance, but it can't create information from nothing.
- **Using unsupported video formats:** While most tools accept common formats like MP4, older or less common video types might not be supported. This can cause uploads to fail.
  - **The fix:** If your video is in an unusual format, use a free online video converter to change it to an MP4 file before attempting to upload it.

What this means for you

- **In everyday life:** You can restore precious old family videos captured on early digital cameras or mobile phones, making them look sharp enough to share on modern high-definition screens or for family slideshows without the blur.
- **For your business or work:** Quickly polish old archival footage, low-quality product demonstration videos, or customer testimonial clips to make them appear more professional and suitable for current presentations or marketing materials.
- **If you're just getting started:** This guide offers a fantastic, risk-free way to experience how artificial intelligence can dramatically improve media. You don't need expensive software or a powerful computer, just a web browser and an old video.

Try it now

Grab a short, low-quality video clip from your phone or computer. Go back to the online NVIDIA AI Super-Resolution demo (or find another one with a quick search) and try upscaling it to 4x quality. See the magic of AI for yourself!

Create Royalty-Free AI Music Loops for Podcasts and Videos

5 steps · Read online

Imagine having custom background music for your podcast, YouTube video, or even a presentation, without needing a huge budget or any musical training. By the end of this guide, you'll know how to use Artificial Intelligence (AI) to generate unique, royalty-free music loops that perfectly match the mood of your content. This guide is for anyone who wants to add a professional touch to their projects without the hassle. It's important to know that while AI can create fantastic music, it cannot read your mind or fully understand complex musical theory, so clear instructions are key!

Before you start

- **A device with internet access:** You'll need a computer, laptop, or tablet with a stable internet connection. Most AI music tools, like the one we'll use, operate through a web browser.
- **A Suno.ai account:** This guide uses **Suno.ai** as a practical example. Visit the Suno.ai website and sign up using a Google, Discord, or email account. This process is typically quick.
- **About 10-15 minutes of time:** This is a rough estimate for following the steps and generating your first few music loops.
- **A creative idea:** Have a general sense of the mood, style, or genre of music you're looking for. Do you need something calm, energetic, mysterious, or upbeat?

Step 1 — Sign up or log in to Suno.ai

First, you need to access an **AI music generator**, which is a special online tool that uses **Artificial Intelligence (AI)** — computer programs that can learn and solve problems — to compose original music based on your descriptions. For this guide, we'll use **Suno.ai**, a popular and user-friendly option. Open your web browser (like Chrome, Firefox, or Safari) and navigate to the Suno.ai website (usually by typing "suno.ai" into your browser's address bar). You'll typically see prominent buttons to "Sign up" or "Log in." Choose the option that suits you, often by linking an existing Google or Discord account, or by creating a new account with your email.

If it looks different: If you're on a mobile device, the layout might be slightly condensed, but you'll still find clear sign-up or log-in options. If you're using a different AI music generator, the exact buttons might vary, but look for similar "Get Started" or "Join" prompts. **You'll know it worked when** you are successfully redirected to a main dashboard or a prominent creation screen, indicating you are logged in and ready to begin generating music.

Step 2 — Find the creation area and understand the prompt

Once you're signed in to Suno.ai, you'll typically land on a page that features a large text box, often labeled "Describe your music here" or "What do you want to hear?". This is your **prompt box** – a designated area where you'll type your instructions for the AI. A **prompt** is simply the text description you give to the AI, telling it exactly what kind of music you want it to create. Think of it like giving a detailed brief to a composer. Look around this area for other settings like a checkbox for "Instrumental" or a "Generate" button.

If it looks different: Some AI music generators might require you to click a "Create" or "Generate" button first to reveal the prompt box. Regardless of the tool, your goal is to locate the main text input field where you can describe your desired music, usually a long, empty box. **You'll know it worked when** you see a clear text input field on your screen, indicating that this is where you'll type your musical vision for the AI.

Step 3 — Craft your music idea (your 'prompt')

Now it's time to tell the AI what to compose! In the prompt box, type a clear and descriptive explanation of the music you're looking for. Consider these elements: the **genre** (like folk, electronic, rock), **mood** (happy, reflective, energetic), **tempo** (how fast or slow the music should be), and **instrumentation** (which instruments you want to hear, such as piano, guitar, or drums). The more specific you are, the better the AI can match your vision. Remember, we're aiming for a **music loop**, which is a short piece of music designed to play seamlessly over and over, so clarity is crucial.

Example: `A calm, ambient electronic loop for a meditation podcast, featuring soft synth pads, a subtle bassline, and a gentle, pulsing rhythm. Focus on a peaceful and reflective mood.` **If it looks different:** Some AI music tools might offer specific dropdown menus for genre, mood, or instrumentation instead of requiring you to type them into the main prompt. If these dedicated options exist, it's generally best to use them as they can guide the AI more precisely. **You'll know it worked when** you have typed your detailed musical description into the prompt box, ready for the AI to process.

Step 4 — Generate your music and listen

After you've carefully written your prompt, look for the "Generate," "Create," or "Make Music" button, usually located below or next to the prompt box. Click this button to send your instructions to the AI. The tool will then process your request, using its understanding of musical patterns and your description to produce sound clips. This process typically takes a few seconds to a minute, depending on the tool and the complexity of your request. The tool may generate one or more variations of your music.

If it looks different: While the AI is generating, you might see a progress bar, a spinning icon, or a message indicating that the music is "processing." If you are using a free plan, you might only get one generation per prompt or have a daily limit on the total number of generations you can make. **You'll know it worked when** one or more generated music tracks appear on your screen, each typically accompanied by a play button. Press the play button on each track to listen and determine if it matches your original vision.

Step 5 — Refine and download your chosen loop

It's common not to get the perfect result on the very first try, so don't hesitate to experiment! If the music isn't quite right, go back to your prompt (from Step 3) and make small, precise adjustments. For instance, you could add details like "add a subtle orchestral swell" or change a mood word from "calm" to "serene." Once you find a music loop you love, look for a download button, often represented by an arrow pointing downwards, a cloud icon, or the word "Download." Click this to save your music file to your device. Most tools allow you to save the file as an **MP3** (a common, compressed audio format that balances quality and file size) or **WAV** (a higher-quality, uncompressed audio format ideal for professional use).

If it looks different: The download icon might be found within a three-dot menu (...) next to the track, or there might be a specific "Export" option. On a free tier, some tools may restrict downloads or offer lower-quality files. Always review the tool's terms regarding "royalty-free" usage to ensure you're allowed to use the music for your specific purpose, especially if it's for commercial content like monetised videos; some may require attribution (crediting the AI). **You'll know it worked when** the music file appears in your device's downloads folder or a designated location, ready for you to use in your podcast, video, or other creative projects.

Common mistakes

- **Being too vague in your prompt:** Just saying "happy music" doesn't give the AI enough detail to understand your specific needs.
  - **The fix:** Be much more specific. Add details about instruments ("upbeat acoustic guitar, light drums, simple piano melody"), tempo ("fast tempo"), and exact mood ("joyful and bright, like a morning sunrise").
- **Forgetting about "loop" functionality:** If the tool has a specific "loop" setting, or if you don't explicitly ask for it, you might get a song with a distinct ending that won't repeat seamlessly.
  - **The fix:** Always include "loop" or "seamless loop" in your prompt, e.g., "a seamless electronic loop." If the tool has a dedicated "loop" checkbox or setting, ensure it's ticked before generating.
- **Not checking usage rights:** Assuming all AI-generated music is free for any use can lead to unexpected copyright issues, especially if you plan to monetise your content.
  - **The fix:** Always confirm the AI music generator's terms and conditions, particularly for "royalty-free" usage and commercial licensing. "Royalty-free" usually means you don't owe ongoing payments, but some tools might still require you to credit the AI (attribution).

What this means for you

- **In everyday life:** You can easily create unique background music for your personal videos, social media stories, or even just for studying or relaxing. This allows you to personalize your digital content in a snap, making your home movies feel more professional or crafting the perfect atmosphere for your focus sessions.
- **For your business or work:** You can quickly produce bespoke, original music for marketing videos, online courses, corporate presentations, or business podcasts. This saves significant time and helps you avoid the complex licensing fees and legalities often associated with traditional stock music libraries.
- **If you're just getting started:** Begin with very simple prompts, like "upbeat acoustic guitar loop" or "calm piano instrumental." Listen to what the AI produces and then gradually add more detail or experiment with different genres.

Try it now

Open Suno.ai on your device, log in, and type a simple prompt like "energetic electronic loop for a workout video" into the prompt box. Click 'Generate' and listen to the results!

Remove Distractions from Photos with Google Photos Magic Eraser

5 steps · Read online

Have you ever taken a fantastic photo, only to have it ruined by a random person in the background, an overflowing bin, or an unsightly power cord? Google Photos' Magic Eraser feature uses a bit of clever Artificial Intelligence (AI) to help you instantly remove those distractions, leaving your pictures looking clean and professional. This guide will walk you through each simple step to make your photos shine.

Before you start

- You'll need the **Google Photos app** installed on your smartphone or tablet.
- You'll need a **Google Account** signed in to the app.
- Pick a photo from your library that has a clear object or person you'd like to remove.
- This feature is generally **free and available** on most modern devices via the Google Photos app, though very heavy or advanced use might have limitations on some free accounts.
- This whole process should take about **5 minutes**.

Step 1 — Open Google Photos and choose a photo

To begin, find and tap the **Google Photos app icon** on your device. It usually looks like a colorful pinwheel. Once the app is open, you'll see your photo library, often displayed as a grid of small picture previews, also known as thumbnails. Tap on the specific photo you want to edit.

What you see: Your device's home screen or app drawer, then your photo library in Google Photos, followed by your chosen photo appearing in full-screen view.

What happens next: The selected photo will open up, filling your screen, ready for you to view or edit.

If it looks different: If your photos aren't showing in a grid, look for a 'Library' or 'Photos' tab at the bottom of the screen. Tap it to see all your pictures.

You'll know it worked when your chosen photo fills the screen, showing all its details.

Step 2 — Find the Magic Eraser tool

With your photo open, you need to access the editing tools. Look for an **'Edit' button** at the bottom of the screen – it often looks like a pencil icon or a set of sliders. Tap this button to enter the editing mode. Once there, you'll see a row of different editing options like 'Crop', 'Adjust', or 'Filters'. Swipe through these options from right to left until you find and tap the **'Tools' section**. Inside 'Tools', you'll then see the 'Magic Eraser'.

What you see: Your photo with various editing icons or tabs appearing along the bottom. After tapping 'Edit', you'll see a new set of options, and then a menu specifically for 'Tools' will slide into view, revealing 'Magic Eraser'.

What happens next: The app transitions from viewing your photo to its comprehensive editing interface. Tapping 'Tools' brings up a selection of smart editing features, including the Magic Eraser.

If it looks different: The 'Edit' button might be a pencil or a 'Settings' icon. If you don't immediately see 'Tools' in the main editing bar, try looking under an 'Adjust' tab or swiping further along the bottom options.

You'll know it worked when you see 'Magic Eraser' clearly listed among the other editing tools, ready for you to select.

Step 3 — Let the AI suggest removals or draw your own

Now that you've selected 'Magic Eraser', the fun begins! The Artificial Intelligence (AI), which is like a clever computer program designed to think and learn, will immediately start analyzing your photo. It's looking for common distractions like people, power lines, or objects that stand out from the background. The AI will then present "suggestions" by highlighting these areas with a subtle, shimmering outline.

What you see: Your photo with the Magic Eraser interface active. You might notice flashing outlines around objects that the AI thinks you might want to remove. Below the image, you'll also see options to 'Erase' (using a brush) or 'Undo'.

What happens next: The AI has completed its scan. If it found anything, the suggested distractions will be highlighted on your screen. If no suggestions appear, that's perfectly fine – you can simply proceed to manually draw over any unwanted areas yourself.

If it looks different: If the highlights don't appear, or the AI doesn't suggest anything, don't worry. This just means you'll guide it yourself. The brush tool will still be active, ready for you to use.

You'll know it worked when the screen shows your photo with potential distractions outlined, or the brush tool is visible, indicating Magic Eraser is ready for input.

Step 4 — Erase unwanted elements

With Magic Eraser active, you have two ways to remove things. If the AI highlighted an object, simply **tap on the highlighted area**. The AI will then try to intelligently remove it and fill the gap with what it thinks the background should look like. For anything the AI didn't highlight, or for more precise control, use your finger to **draw a line or doodle directly over the unwanted object** on the screen. Try to cover the entire object you wish to remove.

What you see: As you tap or draw, the selected part of your photo will disappear almost instantly. The app will attempt to blend the surrounding pixels (the tiny dots that make up your image) to create a natural-looking background where the object once was.

What happens next: The unwanted element vanishes from your photo. You'll see the cleaned-up version right away, giving you a clearer, distraction-free image.

If it looks different: If the result isn't perfect, or if it leaves a blurry patch, don't panic. This tool works best on clear, distinct objects. Simply tap the 'Undo' arrow (often pointing left) at the bottom or top of the screen to reverse the last change and try again.

You'll know it worked when the unwanted object is gone, and the area where it used to be is filled in, making your photo look much tidier.

Step 5 — Refine and save your work

After erasing, take a moment to look closely at your photo. Sometimes, a first pass might leave a slight blur or an incomplete edge. If you need to make further adjustments, you can zoom in by pinching out on the screen with two fingers, and then use the brush tool again for finer detail. You can also use the 'Undo' arrow if you change your mind about a particular erasure. Once you're happy with your edits, look for a **'Done' button** or a tick icon, usually at the bottom-right, and tap it. Then, to save your masterpiece, tap **'Save copy'** (or sometimes just 'Save') at the bottom of the screen.

What you see: Your edited photo still on screen, possibly with 'Done' and 'Save copy' buttons. After tapping 'Save copy', you might see a brief 'Saving' message.

What happens next: Tapping 'Done' exits the Magic Eraser tool, taking you back to the main editing screen. Tapping 'Save copy' processes your changes and creates a brand-new version of your photo in your Google Photos library, leaving your original photo untouched.

If it looks different: The 'Done' button might be a simple checkmark. 'Save copy' might simply be labeled 'Save', or it could appear as a menu option after you tap 'Done'.

You'll know it worked when you see a confirmation message like 'Saved' or 'Copy saved', and you can then find the new, cleaned-up version of your photo alongside the original in your Google Photos library.

Common mistakes

- **Over-erasing**: Accidentally drawing over too much of your photo, or parts you actually wanted to keep, can make the result look unnatural.
    - **The fix:** Always use the 'Undo' arrow immediately if you make a mistake. It will revert your last action, allowing you to try again with more precision.
- **Unnatural background fill**: Sometimes, the AI can't perfectly guess what should be behind an object, leaving a blurry or repetitive pattern.
    - **The fix:** This often happens with complex backgrounds. Try erasing in smaller, more precise sections. If it still doesn't look right, remember that Magic Eraser works best on less complicated backgrounds.
- **Forgetting to save the copy**: After spending time perfecting your photo, it's easy to tap back out of the app without saving, losing all your hard work.
    - **The fix:** Always remember to tap 'Done' and then 'Save copy' (or 'Save') once you're satisfied. Google Photos will usually prompt you if you try to exit without saving, but it's good practice to make it a habit.

What this means for you

- **In everyday life:** Instantly clean up your holiday snaps by removing stray tourists, make family photos picture-perfect by getting rid of background clutter, or simply tidy up a shot of your pet for social media.
- **For your business or work:** Enhance product photos for online listings by removing dust or distracting elements, clear up real estate images by erasing power lines or bins, or prepare professional-looking profile pictures for your business's social media.
- **If you're just getting started:** Look for a photo with a single, clear, distinct object against a relatively uniform background – like a small sign on a wall, or a distant person against the sky. This will give you the best chance of a perfect 'erase' on your first try!

Try it now

Find a photo in your Google Photos library that has a simple, easy-to-remove object, like a small rubbish bin in the distance or a loose cable on the floor. Follow the steps above to magically erase it and see the difference!

How to use Kling AI video templates in Picsart Flow

5 steps · Read online

> **In short:** Picsart Flow includes Kling-powered video templates that turn a still photo, a character reference image, or a short script-plus-footage combo into a short AI-generated clip. You can try them without editing experience, though longer videos and higher resolution usually need a paid plan.

Want to make a short video but don't have filming skills, a camera, or editing software? Kling is a video-generation model that creates short clips from a still photo, a character image, or a written script combined with your existing footage. Picsart Flow hosts it as ready-made templates, so you mostly just upload, type a sentence, and wait. This guide walks through all three template styles one at a time. None of this requires a video background — only a clear idea of what kind of clip you'd like to end up with.

Before you start

- A free Picsart account (you can sign up with email, Apple, or Google — check the latest sign-up options on the Picsart login page).
- The Picsart app on your phone, or the Picsart web editor on a laptop. Picsart Flow tools may live in different spots in each, so this guide uses language like "look for" rather than exact menu names.
- A few source images saved to your device. For best results, use clear, well-lit photos where the subject is centered.
- Around 15–30 minutes for your first full run-through (most of that is waiting for the AI to render).
- A realistic expectation: AI video is impressive but imperfect. People may look slightly off, hands can blur, and motion can wobble — that's normal for current tools.

Step 1 — Open Picsart Flow and find the Kling templates

Tap the Picsart app to open it, then look along the bottom or side menu for a section called something like **Flow**, **AI Tools**, or **Video**. The exact label depends on your version and platform, so if you don't see "Flow", search inside the app for "Kling" or "AI video" — that usually leads you to the same area. Once you land on the Kling page, scroll until you see a row of template tiles. You'll typically find three styles: one that animates a single photo, one built around a character reference image, and one that combines your footage with a script.

Example: search terms to try inside the app — `Kling`, `AI video`, `photo to video`.

You'll know it worked when you see a screen with several preview tiles, each labeled with the input it needs (photo, image, or footage).

Step 2 — Use the photo-to-video template

Pick the **photo-to-video** template and tap it. The app will ask you to upload one image from your camera roll. Choose a clear photo where the main subject (a person, a pet, a product) is easy to see. After upload, a small text box appears asking what should happen in the video — that's called a **prompt** (the instruction you give the AI). Type one short sentence describing the motion you want, then tap the generate button. The app will show a progress indicator while Kling works. This step usually takes 30 seconds to a few minutes, depending on clip length and server load.

Example: `The woman slowly turns her head to the right and smiles, with soft sunlight on her face.`

You'll know it worked when a short video preview appears in place of the still photo. If it looks wrong, tap regenerate or tweak the prompt — small wording changes make a big difference.

Step 3 — Try the character reference sheet template

Now pick the **character reference** template, sometimes labeled "character sheet" or "consistent character." This one is designed for keeping the same face or outfit across multiple clips — useful for short stories, social media series, or product mascots. The app will ask you to upload one or more images of the same character from different angles. Add a short prompt describing the scene, then generate. Kling tries to keep facial features and clothing consistent, which is harder than it sounds for AI, so your first result may need two or three tries.

Example: `A young woman with short black hair, wearing a yellow raincoat, walks through a misty forest at dawn.`

You'll know it worked when the generated clip shows a person who actually resembles your reference images, in roughly the same clothes and setting. If the face drifts, regenerate or add more reference images — more input usually helps consistency.

Step 4 — Combine footage with a script

The third template, often called **video-to-video** or **script + footage**, lets you feed Kling a short video clip you already have plus a written description of what should change. Think of it as a soft redraw: the AI keeps the general shape and motion of your clip, but re-styles it. Upload a short clip (a few seconds is enough), then type a script describing the new mood, setting, or style. This is the template most people use for stylized social posts — turning a daytime clip into a moody evening scene, for example. Generation takes longer here because the model has more to process.

Example: `Change the setting to a neon-lit Tokyo street at night, keep the same walking motion.`

You'll know it worked when your uploaded clip now looks visually different (new lighting, color, or setting) while the original motion still feels familiar.

Step 5 — Preview, trim, and export

Once any template finishes, tap the result to open a preview screen. Most flows let you trim the start and end by dragging handles on a small timeline bar. If the clip is too long, shorten it; if the opening is awkward, trim the first second. When it looks right, look for an **Export**, **Save**, or **Download** button — usually at the top or bottom of the preview. On the free plan, you may get a watermark or a lower resolution; higher resolutions and watermark-free exports are typically behind a paid plan. The exact limits change often, so check Picsart's current pricing page if this matters for your project.

Example: trim suggestion — `keep only seconds 2 through 6` (drag the handles until those seconds are highlighted).

You'll know it worked when a finished video file appears in your phone's photo library or downloads folder.

Common mistakes

- **Uploading a blurry or busy photo.** Kling struggles with tiny faces and cluttered backgrounds. Use the clearest, simplest image you have.
- **Writing prompts that are too long.** Two sentences is plenty. Long paragraphs confuse the model and often produce muddled motion.
- **Expecting perfect hands or text.** AI video still fumbles fingers, writing, and small details. If those matter, plan to crop them out of frame or pick a different shot.

What this means for you

- **In everyday life:** you can turn a favorite vacation photo into a short clip to share with family, without learning video software.
- **For your business or work:** small businesses can produce quick product clips, social teasers, or stylized B-roll (extra footage used to support a story) without hiring a video editor — useful for one-person shops and busy marketing teams.
- **If you're just getting started:** begin with the photo-to-video template on a single clear selfie or pet photo. It's the fastest way to see what Kling can do in under five minutes.

Try it now

Open Picsart, search for the photo-to-video template, and upload one clear photo of someone smiling. Type: `The person gently blinks and turns to look out a window, soft natural light.` Tap generate, wait a minute, and watch your first Kling clip come to life.

How to create images with AI: a beginner's first try

6 steps · Read online

Imagine typing a simple sentence and watching a brand-new, unique image appear on your screen in seconds. It's like having a digital artist ready to bring any idea to life, no drawing skills required. Whether you want to design a custom greeting card, create a visual for a work presentation, or just visualise a creative thought, **AI image generators** make it possible. By the end of this guide, you will know how to choose a free tool, write a description that the AI understands, and be aware of how you can use your new creations.

Before you start

- **A computer or smartphone** with an internet connection and a web browser.
- **A free account** with an AI image generation platform. Many trusted companies offer free ways to create images — for this guide, we suggest exploring **Microsoft Copilot** (at copilot.microsoft.com) or **Canva** (at canva.com), as both have options to create images without needing a paid subscription. You'll need to sign up for a free account if you don't already have one.
- **Roughly 10-15 minutes** of time to follow along with the steps and create your first images.

Step 1 — Choose and open a free tool

Your first action is to pick a free online tool and open it in your web browser. You don't need to spend money to try this out, as many trusted companies offer free **AI image generators** — these are smart computer programs that act like a digital artist, translating your written words into pictures. We'll use Microsoft Copilot as our primary example.

Open your web browser and go to `copilot.microsoft.com`. If you're not logged in, you'll see a screen asking you to sign in with a Microsoft account or create one. Once you're signed in, you'll see the Copilot interface. **If it looks different:** If you're using Canva, open `canva.com` and log in. Then look for AI-related features, often found in the design creation flow or through a search function. **You'll know it worked when** you see the main interface of your chosen tool, ready for you to input commands or search for tools.

Step 2 — Find the image generator

Now that you're in the tool, you need to tell it you want to create an image. In Microsoft Copilot, look for the image generation feature — this may be a dedicated button, a mode toggle, or you can simply describe what you want in the chat. Type a simple instruction into the input area. When you do this, you're giving a **prompt** to the AI – this is just a fancy word for the instruction you write to tell the AI what you want to see. **If it looks different:** In Canva, explore the AI image creation options, which may be labeled differently or located in various places within the design interface. **You'll know it worked when** the tool responds by acknowledging your image request, or when a dedicated image generation interface appears with a specific text box for your image description.

Example:

> "Create an image of"

Step 3 — Write your first detailed "prompt"

To get a great result, your **prompt** (the instruction) needs to be more than just one or two words. Think of it like giving a detailed recipe to a chef – the more specific you are, the better the dish will be! A good prompt usually has three main ingredients:

1. **The Subject:** Who or what is the main focus of the picture (e.g., "a fluffy ginger cat").
2. **The Details:** What they are doing, the colors, and the background (e.g., "wearing a blue knitted jumper, sitting on a wooden park bench, bright sunny day").
3. **The Style:** Do you want a realistic photo, a watercolour painting, a 3D cartoon, or something else (e.g., "soft watercolour painting style").

Now, type your complete, detailed description into the text box and send it. The AI will then start processing your request. **If it looks different:** If the AI seems confused, try rephrasing your request slightly, such as "Generate a picture that shows..." instead of "Create an image of...". **You'll know it worked when** you see a message like "Generating your image..." or a loading indicator, confirming the AI is working on your request.

Example:

> "A fluffy ginger cat wearing a blue knitted jumper, sitting on a wooden park bench, bright sunny day, soft watercolour painting style."

Step 4 — Generate and review your images

After you send your detailed prompt, the AI will get to work. It usually creates a few different image options based on your description. This process can take anywhere from a few seconds to about half a minute, depending on the tool and how busy it is.

Once the AI has finished, you'll see a collection of images appear on your screen, usually displayed side-by-side or in a grid. Take a moment to look at them. Don't worry if they're not absolutely perfect yet – this is just the first draft! This step shows you the AI's interpretation of your words. **If it looks different:** Some tools might generate only one image at a time, or show a progress bar as they work. If the images look completely wrong, double-check your prompt for any typos or unclear instructions. **You'll know it worked when** you see your newly generated images displayed clearly on your screen, all inspired by the description you provided.

Step 5 — Refine or download your favorite

You don't have to settle for the first set of images. Most AI tools allow you to refine your results, just like talking to a helpful assistant. You can ask for changes to the images you see or generate completely new ones with a tweaked prompt.

If you want to make changes, simply type a follow-up instruction. If you're happy with an image, look for a "Download" or "Save" button, which is often an arrow pointing downwards or a disk icon, usually found near the individual image or at the bottom of the collection. Clicking this will save the image to your device. **If it looks different:** The download button might be hidden behind a "..." (three dots) menu or a right-click option. If you're refining, some tools let you select a specific image to refine, while others apply changes to the whole set. **You'll know it worked when** you either see a new set of images based on your refinements, or you've successfully saved an image file to your computer or phone.

Example:

> "Change the jumper color to yellow, and add a small friendly magpie sitting on the bench next to the cat."

Step 6 — Review your image for watermarks and usage

Before you share your AI masterpiece, it's helpful to understand two quick concepts: **watermarks** and **usage rights**. Take a close look at the images you've generated.

A **watermark** is a faint, semi-transparent logo or symbol placed in the corner of an image to show which tool created it. Many free tools include these automatically to keep the images easily recognizable as AI-made. Next, think about **usage rights**, which are the rules for how you can use the images. As a general rule, images you make on free accounts are perfect for personal fun, sharing with friends, or school projects. However, if you plan to use them to make money (like on a commercial website or product packaging), you must check the tool's terms of service (their official list of rules) as some free tiers do not allow commercial use. **If it looks different:** Some tools might offer a paid plan to remove watermarks, or their terms of service might be difficult to find. Look for a "Terms" or "Legal" link, usually at the bottom of the website. **You'll know it worked when** you can identify if your image has a watermark, and you have a general understanding of whether your chosen free tool allows personal or commercial use of your creation.

Common mistakes

- **Being too vague:** Typing "a dog" gives the AI too much room to guess. You might end up with a cartoon pug when you wanted a realistic golden retriever.
  - **The fix:** Be specific! Add details about the breed, color, action, and setting (e.g., "A golden retriever puppy frolicking in a field of sunflowers, hyperrealistic photo style").
- **Expecting perfect text spelling:** AI image generators are brilliant with shapes and colors, but they often struggle to spell words inside images. If you ask for a sign that says "Welcome", the AI might print scrambled letters or strange symbols.
  - **The fix:** It's usually best to avoid asking the AI to generate text within your images. Create the image first, then add text later using a separate image editing tool.
- **Forgetting that AI does not know real people:** For safety and privacy reasons, free tools will not let you generate realistic images of real public figures, celebrities, or your friends.
  - **The fix:** Stick to fictional characters, general subjects, or broad descriptions (e.g., "a young woman reading a book" rather than "Sarah reading a book").

What this means for you

- **In everyday life:** You can instantly create custom birthday invitations, organize your home with beautifully styled labels, or make a unique phone wallpaper that no one else has – all with just a few words.
- **For your business or work:** You can quickly mockup visual ideas for presentations, create colorful backdrops for social media posts, or brainstorm logo concepts without spending hours searching through generic stock photo libraries.
- **If you're just getting started:** Remember that experimenting is half the fun. There are no wrong answers, and the more you practice describing your ideas, the better and more accurate your results will be.

Try it now

Let's get you started right now. Open Microsoft Copilot or Canva in a new tab, find the image generator function, and type this simple prompt:

*"A friendly koala wearing a green scarf, holding a warm mug of tea, soft digital illustration style"

How to Edit Your Phone Photos with Simple Text Instructions in Meta AI

6 steps · Read online

> **In short:** Send Meta AI a photo, type what you want changed ("remove the person on the left", "make the sky look like sunset"), and it returns a new version. Simple edits usually work on the first try; complex changes may need a couple of attempts.

Ever take a great photo and then notice a stranger walking through the background? Or snap a sunset shot where the sky came out dull and gray? You don't have to learn a photo editor to fix these — Meta AI can do small edits just by typing what you want. This guide walks through how, on your phone, in a few minutes.

Before you start

- A phone with **Meta AI** available — built into WhatsApp, Instagram, Facebook, or Messenger. If you don't see it in any of those, the same chat works at **meta.ai** in your phone's browser.
- A Meta account (the same one you use for Facebook or Instagram works).
- One photo saved on your phone that you'd like to edit — landscape, portrait, or a quick snapshot, any of them.
- A working internet connection.
- About **5 minutes** for your first edit, and a little patience if the first attempt isn't quite right.

> **One honest note:** Meta AI is great at small, focused edits — sky color, brightness, removing a clear object. It is not a replacement for a full photo editor like Photoshop. If your request is very complex, the result may surprise you, and that's normal.

Step 1 — Open a chat with Meta AI

Open WhatsApp, Instagram, Facebook, or Messenger, then find the contact or chat labeled **Meta AI**. It's usually marked with a little blue or purple ring around its profile picture. In WhatsApp, tap the green new-chat button and look for "Meta AI" near the top of your contacts. In Instagram, start a new direct message and search for "Meta AI". A one-on-one chat opens, with a text box at the bottom and a paperclip or + button on the left for attaching things. If you can't find Meta AI in any app, open your phone's browser and go to **meta.ai** — sign in with your Facebook or Instagram account if asked, and you'll land in the same chat. 💬 Example: Open WhatsApp → tap the new-chat icon → tap "Meta AI" in the contacts list. You'll know it worked when a chat thread opens with "Meta AI" at the top and a text box you can type into.

Step 2 — Send the photo you want to edit

Tap the paperclip, +, or camera icon next to the text box, then pick **Photo** or **Gallery** from the menu. Your phone's photo library opens. Select the photo you want — just one to start, so the AI has a clear focus. The photo appears as a small preview in the chat; tap **Send** (the arrow button) to deliver it. The image uploads in a few seconds, and Meta AI now "sees" it. If you're using **meta.ai** in the browser instead, look for an upload icon — usually an arrow pointing up or a small picture frame — on the left side of the message box. Tap it, pick the photo, and send the same way. 💬 Example: One landscape photo where the sky looks flat and gray — perfect for your first edit. You'll know it worked when your selected photo shows up in the chat as a sent message.

Step 3 — Type what you want changed

In the text box right below the photo, write a short, plain sentence describing the edit. Tap send. You'll see three small dots animating while Meta AI thinks — that means it's working. Within a few seconds, sometimes up to a minute, it replies with a new image, usually with a one-sentence acknowledgment like "Here's the edited version." The new photo lands in the chat as another message bubble. If the reply is **only text** with no image, your instruction was probably unclear — try again with the wording in the example below. 💬 Example: "Make the sky look like a bright sunset with orange and pink clouds" — or simpler, "Remove the person on the right" or "Brighten the photo so it doesn't look so dark." You'll know it worked when a new image appears in the chat that looks like your photo, but with the change you asked for.

Step 4 — Look closely at the edit

Tap the new image to open it full-screen. Pinch out with two fingers to zoom in on the area you asked to change. Now look at it honestly: did the edit land where you wanted? Sometimes Meta AI nails it. Other times the change is subtle, or it shifted something you didn't ask to change — a face looks slightly different, or a building in the background got smudged. That happens, and it doesn't mean you did anything wrong. If the image is small and won't zoom in the chat, take a screenshot of your phone (side button + volume up on most phones) — that captures the full-size version for closer inspection. 💬 Example: Ask yourself — is the person really gone, or is there a faint shadow where they stood? Is the sky the color you wanted, or does it look painted on? You'll know it worked when you can decide, honestly, whether this version is good enough to keep or whether you'd like to try again.

Step 5 — Ask for a redo if it's not quite right

If the edit missed the mark, reply in the same chat with a clearer or more focused instruction. Keep the conversation going — Meta AI keeps the photo and your request in context, so you can iterate. Useful phrases: "Try again", "Make it stronger", "Keep everything else exactly the same and

only change the sky", or pointing out what went wrong, like "The person is still there — try once more". If the AI keeps misunderstanding after several tries, start a brand new chat, send just the photo again, and write your best, clearest instruction first. Long threads with many edits sometimes confuse the model. 💬 Example: First try didn't remove the person cleanly? Send: "Try again — remove the person on the right side, but keep everything else exactly the same." You'll know it worked when a new version finally looks like what you wanted — or when you've accepted that this particular edit isn't possible and you'd rather try a different photo.

Step 6 — Save the photo you like

When you're happy with an edited version, save it to your phone. Tap and hold the image in the chat — a small menu appears with options like **Save**, **Save to Photos**, or **Share**. Tap "Save" and the photo lands in your camera roll or gallery, just like one you took yourself. In some apps you'll see a small downward-arrow icon on the image — tap that instead. In **meta.ai** in the browser, long-press the image (or right-click on a computer) and choose "Save image as". Only save once you're sure — once it's in your camera roll, you can't easily get back to the older versions still sitting in the chat. 💬 Example: Save your favorite version, then close the chat. Your original photo is untouched on your phone, and the edited version is now saved next to it. You'll know it worked when you open your phone's Photos or Gallery app and see the edited photo sitting there.

Common mistakes

- **Asking for too many changes at once.** "Remove the person, fix the sky, brighten everything, and add a reflection in the water" is too much for one message. The AI may handle one or two of those, badly, while ignoring the rest. **Fix:** ask for one change at a time, and send the photo back through Step 3 each time.
- **Writing vague instructions.** "Make it better", "fix this", or "improve it" leaves the AI guessing — usually wrong. **Fix:** describe exactly what to change in plain words. "Make the sky bluer" works; "nicer sky" doesn't.
- **Not zooming in to check the result.** On a phone screen, a small image can look fine until you look closely — then you notice a smudged face or a ghostly outline where someone was removed. **Fix:** always open the image full-screen and zoom in on the edited area before saving.
- **Saving the first version without comparing.** The first attempt is rarely the best one. **Fix:** wait until you've tried once or twice and you've got the version you actually want, then save that one.

What this means for you

- **In everyday life:** vacation photos where someone walked into your shot, birthday candles that came out dim, a beach picture where the sky was gray — all fixable in under a minute with one short instruction.
- **For your business or work:** quick product photos, social media posts, or marketplace listings where the background needs cleanup. Type what you want removed, send it through, save the result. It won't replace a designer, but it gets you 80% of the way in seconds.
- **If you're just getting started:** pick a photo that's been bothering you — maybe a great group shot with a stranger in the corner — and try removing them first. It's the easiest edit to see if Meta AI is working for you.

Try it now

Open Meta AI right now in WhatsApp, Instagram, or **meta.ai** in your browser. Send one photo you've taken recently — anything from this week. Then type one short instruction, such as "Make the colors look brighter" or "Remove the trash can in the corner". The whole thing takes under two minutes, and you'll see exactly how the tool responds to your wording — which makes the next edit easier too.

How to Generate Images with Google Gemini: A Beginner's Guide

6 steps · Read online

Imagine being able to paint a picture just by describing it out loud. Whether you need a quick illustration for a work presentation or a fun graphic for a party invitation, Google Gemini lets you turn your words into custom images in a matter of seconds. This guide will walk you through the simple steps to create your own unique visuals.

Before you start

- **A Google Account:** You will need to be signed into a standard Google account. If you use Gmail or YouTube, you already have one.
- **Access to Gemini:** Open your web browser and go to `gemini.google.com`, or open the Gemini app on your mobile device.
- **Device:** This guide works on any computer, tablet, or smartphone with internet access.
- **Cost:** Image generation in Gemini is currently available for free, though usage limits may apply or change over time with AI tools.
- **Time:** This guide will take approximately 5-7 minutes to follow.

Step 1 — Open Gemini and start a new chat

First, make sure you're ready to talk to Gemini. When you open `gemini.google.com` in your web browser, you'll see the main Gemini interface, which typically features a large text box at the bottom of the screen and a "New chat" button somewhere near the top. This is your blank canvas for a fresh conversation.

When you click or tap the "New chat" button, the text box at the bottom will clear, indicating Gemini is ready for your first instruction. This prepares the AI for a new request, ensuring it doesn't get confused with previous topics.

If you don't see a "New chat" button, look for a plus sign (+) icon or a menu option like "Start new conversation" in the top-left corner.

You'll know it worked when the text input area at the bottom of the screen is empty and ready for you to type.

Step 2 — Tell Gemini to generate an image

To create an image, you need to write a **prompt** (the written instruction you give to an AI, acting like a friendly request to a helpful assistant). Gemini needs to know right away that you want a picture, not just text. Always start your prompt with a clear command for image creation. Type your command into the text box at the bottom of the screen.

After typing your command, press the "Enter" key on your keyboard or click the send icon (often a paper airplane or arrow button) next to the text box. Gemini will then process your request, and you'll typically see a "Generating..." message as it works.

If you accidentally just type a description without the command, Gemini might try to answer you with text. Just clarify in your next message, "Please make an image of that."

Example:

> "Create an image of..."

You'll know it worked when Gemini starts to show a message like "Generating image(s)..." or "I'm working on that for you."

Step 3 — Describe your subject with details

An AI **image generator** (a tool that turns written text into digital pictures) needs details to work its magic. Think of it like describing a scene to a friend who has their eyes closed. After your initial command from Step 2, continue to type in the main text box, adding specific information about what you want to see.

When you add more detail to your prompt and press Enter or the send button, Gemini will process this enriched request. It will use these clues to build a more accurate and vibrant picture, displaying the results directly in the chat window.

If the images don't quite match your vision, it often means your description was a little too vague. Try adding more adjectives or focusing on specific objects, colors, or actions.

Example:

> "Create an image of a fluffy golden retriever puppy wearing a bright red collar, sitting on a sun-drenched green lawn next to a yellow tennis ball."

You'll know it worked when Gemini presents you with usually two to four image options that directly reflect the details you provided in your prompt.

Step 4 — Choose your visual style

Gemini can create images in many different artistic styles. If you don't specify a style, it will usually default to a realistic photo. To change the mood or appearance, add style keywords to your prompt in the same text box where you've been typing. This instructs Gemini on the artistic direction for the image.

When you include a style in your prompt and submit it, Gemini will generate new images that incorporate that specific artistic look. You'll see the images appear within the chat interface, often in a grid or gallery format. This allows you to easily compare different styles or refinements.

If the style doesn't look quite right, try using a slightly different keyword (e.g., "cartoon" instead of "3D cartoon") or being more specific (e.g., "impressionist oil painting").

Example:

> "Generate a watercolour painting of a quiet beach at sunset, with soft purple and orange colors in the sky."

You'll know it worked when the new images Gemini presents clearly show the artistic style you requested, such as brushstrokes for a "watercolour painting" or distinct lines for "line art."

Step 5 — Refine your generated images

Once Gemini generates your images, you might want to make some tweaks. You don't need to start from scratch. You can simply talk to Gemini like a person, typing a new instruction in the text box that refers to the images it just made. This tells Gemini to adjust the previous output rather than starting a completely new image set.

After you send your refinement request, Gemini will consider your feedback and generate updated versions of the images, incorporating the changes you asked for. These new images will appear in the chat, usually below your request.

If Gemini seems confused by your request, try being more specific about which image you want to change ("the last one," "the first image") or what exact element needs altering ("the background," "the animal").

Example:

> "Can you make the sky in that last picture a bit more blue, and add a single sailboat in the background?"

You'll know it worked when Gemini displays new images that show the specific adjustments you asked for, like a bluer sky or a newly added sailboat.

Step 6 — Save or share your image

Once you have an image you like, you'll want to save it or share it. Look for the image within the chat window. Usually, if you click or tap on an image, it will open in a larger view, where you'll find options to download or share.

When you select the download option (often an arrow pointing down, or a "Download" button), the image will be saved to your device's downloads folder or photo gallery. If you choose to share, a menu might appear allowing you to copy the image or share it directly to other applications.

If you don't see a clear download button, try right-clicking (on a computer) or long-pressing (on a mobile device) the image itself; this usually brings up a context menu with "Save image as..." or "Download image" options.

You'll know it worked when the image file appears in your device's downloads or photo library, ready for you to use or share.

Common mistakes

- **Being too brief:** Writing "a cat" will give you a random result. Instead, give Gemini specific details.
  - **The fix:** Writing "a sleepy tabby cat curled up on a knitted blue blanket in front of a warm fireplace" gives Gemini the clues it needs to make something beautiful and specific to your idea.
- **Asking for readable text:** AI image tools sometimes struggle to spell words correctly or make them look natural inside pictures.
  - **The fix:** It is usually best to generate the image without any text, and then add your words later using a simple design app or photo editor.
- **Using negative words:** AI models respond better to what *should* be in the picture, not what *shouldn't*.
  - **The fix:** Instead of saying "a kitchen with no clutter," try saying "a clean, tidy, minimalist kitchen with polished benchtops." Always frame your request in a positive, descriptive way.

What this means for you

- **In everyday life:** You can design personalized birthday cards, create unique social media posts, or even generate custom phone wallpapers to match your mood.
- **For your business or work:** You can quickly create eye-catching graphics for slide decks, blog posts, or social media updates without needing a massive budget for stock photos or a professional designer.
- **If you're just getting started:** Don't worry about making your prompts perfect. Just play around, try different styles, and see what the AI creates. The best way to learn is by experimenting!

Try it now

Let’s make your very first image. Copy and paste the prompt below into Google Gemini right now to see what happens:

> **"Create a simple, cheerful cartoon illustration of a happy koala holding a warm cup of tea."**

Create Your First AI Image with Midjourney: A Beginner Walkthrough

6 steps · Read online

> **In short:** Midjourney turns a written description into a picture in roughly a minute. You'll need a Midjourney subscription, about ten minutes, and a clear idea of what you want to see. The finished image is yours to download and use.

In the next ten minutes, you'll go from a blank screen to a finished picture saved on your device. No art background needed — if you can describe what you want, you can do this. We'll skip the jargon and focus on the buttons you'll actually press.

Before you start

- **A Midjourney account** — sign up on the official Midjourney website using your email, a Google account, or a Discord login. Generating images requires a paid subscription; current plans and pricing are listed on Midjourney's official pricing page, so check there before you commit.
- **A web browser or the Midjourney web app** — you can work on a computer, tablet, or phone. Discord also works as an alternative entry point if you already have an account there.
- **About 10–15 minutes** of your time, plus another minute or so while the AI draws.
- **One clear idea** of a picture you'd like to see. "A cat" works as a starter; "a calico cat asleep on a stack of books, soft morning light" gives a much better result.

Step 1 — Create your Midjourney account

Open your web browser and go to the official Midjourney homepage. Look for a sign-in or sign-up button — it's usually in the top-right corner, often labeled something like **"Sign In"** or **"Get Started."** Click it, then sign in with your email, a Google account, or a Discord login.

After signing in, the page will guide you through choosing a subscription plan. Pick whichever option feels right for a first try — you can change or cancel later. Once payment is confirmed, the page reloads and drops you into your personal Midjourney workspace.

If you'd rather start in Discord, the same tool runs inside a chat window there. Either path leads to the same image-quality results, so don't worry about choosing wrong.

Example: when asked for your email, use the one you check most often — that's where receipts and account notices will go.

You'll know it worked when you see the Midjourney workspace, with a place to type your picture description and your account name visible somewhere on screen.

Step 2 — Find the create screen

Once you're in your workspace, look for the area where you can start a new image. On the Midjourney web app, this is usually a large text box near the middle of the screen, sometimes paired with a sidebar that lists past images or quick-start options. On mobile, the input area may fill the whole screen.

Click anywhere inside that text box. Your cursor should start blinking there. Don't type yet — just confirm you've found it. This box is where you'll describe your picture, so it's worth a moment to make sure you're in the right place.

If your version looks more like a Discord chat, type **/imagine** in any message field and Discord will offer to autocomplete it. Either way, you're heading to the same destination: a box waiting for your words.

If you don't see a text box at all, look for a button that says something like **"Create"** or **"Imagine,"** or shows a + symbol — clicking it usually opens the input area.

You'll know it worked when your cursor is blinking inside an empty text box and no error message is showing.

Step 3 — Write your first prompt

Now describe one picture in plain words. Imagine texting a friend who paints for you — tell them the subject, the mood, the setting, and the style. A **prompt** is just the name for this description; nothing fancy.

Be specific. *"A flower"* gives you a generic flower. *"A single red poppy in a glass jar, morning sunlight, photo-realistic"* gives you something that actually looks the way you pictured it. Short sentences are fine. So is a comma-separated list.

When your description feels right, look for a button nearby — usually a paper-plane icon, an arrow, or the word **"Generate"** — and click it once. Don't click twice; each click spends a credit and starts a fresh image, so patience saves you money.

Example: `a cozy reading nook by a window, warm lamp light, watercolor illustration, soft colors`

You'll know it worked when the button greys out, a small progress message appears, and within about a minute the first signs of your picture start forming on screen.

Step 4 — Review your generated image

When generation finishes, you'll typically see your image appear on screen. Depending on your settings and the version being used, this might be a single image or a grid of variations. Either way, take a moment to look at what Midjourney created from your description.

If you're looking at a grid and none of the images feels right, look for a button like **"Re-roll"** or a circular ↻ arrow somewhere near the grid. That tells Midjourney to try again with the same prompt and produce new attempts.

You'll know it worked when you can clearly see your generated image(s) and at least one of them makes you think *"ooh, that's nice."*

Step 5 — Upscale your favorite

If you're working with a grid, near it you'll typically find a row of small buttons labeled something like **U1, U2, U3, U4** — one under each image. The **U** stands for **upscale**, which simply means turning the small preview into a full-size, sharp picture. There are usually also **V1–V4** buttons, which create **variations** on that specific image.

Click the **U** button under the picture you like most. Within a few seconds, a single full-resolution version of that image appears. This is the one you'll keep.

If the upscaled image has a small detail you'd like changed — say the cat's eyes are the wrong color — you can run a new variation on it, or simply upscale a different image from the grid instead.

You'll know it worked when a single, larger, sharper image appears alongside or in place of the small grid preview.

Step 6 — Download the image to your device

Click the upscaled image to open it full-size. Look for a download control — this might be a small downward arrow ⬇, a button that says something like **"Save"** or **"Download,"** or an option inside a right-click or menu. One tap or click should save a PNG or JPG file straight to your device.

On a phone, the image often opens in a new tab. From there, look in your browser's menu (often three dots or a share icon) for an option to save the picture to your Photos or Gallery app. On a computer, the file usually lands in your Downloads folder.

The file is now yours. You can post it, print it, set it as your wallpaper, or send it to a friend. Midjourney's standard paid plans deliver images without a watermark, so the picture is ready to share as soon as it's saved.

You'll know it worked when the file appears in your **Downloads** folder (computer) or **Photos / Gallery** app (phone), and you can open it again from there.

Common mistakes

- **Writing a one-word prompt.** "Cat" gives you a random cat. Always include at least one detail about mood, style, or setting — Step 3 has an example you can copy and tweak.
- **Clicking Generate twice in a hurry.** Once is enough. Every extra click spends a credit and starts a fresh image. If the first result isn't ready yet, just wait — a progress indicator should be visible.
- **Asking the AI to draw readable text.** Midjourney struggles with words inside images — a sign that says "OPEN" often comes out misspelled or jumbled. Use a real text tool or a graphic editor for readable words, and let Midjourney handle the rest of the picture.

What this means for you

- **In everyday life:** your first image is great for a phone wallpaper, a birthday card, or a custom avatar — small personal touches that feel one-of-a-kind without paying a designer's fee.
- **For your business or work:** a single Midjourney picture can become a blog header, a slide background, or a quick mockup for a client conversation. It's not a replacement for a designer on real projects, but it's a fast way to test an idea visually.
- **If you're just getting started:** the easiest first move is to re-run the same prompt with one word changed each time — *"watercolor"* becomes *"oil painting"* becomes *"pencil sketch"* — and watch how the style shifts while the subject stays the same.

Try it now

Take one minute and write down a single sentence describing a picture you actually want to see. Then open Midjourney, paste it into the prompt box, and hit Generate. By the time you finish reading this paragraph, your first image will probably be ready to look at.

How to Create AI Logos with Recraft V4.1 Vector

6 steps · Read online

> **In short:** Recraft describes a two-stage AI logo workflow — generate lots of concepts on a raster model first to explore ideas, then switch to Recraft V4.1 Vector for the final mark. The vector output is an SVG file that stays crisp from a 16-pixel browser tab icon to a printed sign.

Designing a logo used to mean either learning vector design software yourself or paying a designer. Today, AI tools like Recraft can sketch dozens of ideas in seconds. The catch is the file format: a regular image — a "raster" picture made of tiny colored squares called pixels — can look fuzzy when you blow it up or shrink it down. A vector file, built from mathematical lines instead of pixels, stays sharp at any size.

This guide walks through Recraft's two-stage approach: explore concepts as raster images first, then generate the final logo in Recraft V4.1 Vector. The final file is an SVG (Scalable Vector Graphics) you can use on a tiny favicon, a website header, or a printed sign without any blurriness.

Before you start

- A free Recraft account (sign up at recraft.ai with an email address or a Google account — check the site for the current plan options, since free and paid tiers exist)
- A recent web browser (Chrome, Firefox, Safari, or Edge)
- A rough idea of your brand: name, vibe, and one or two colors you like
- About 20–30 minutes for your first attempt

Step 1 — Sign in and start a new project

Open recraft.ai in your browser and sign in with the account you just made. Look for a button that says something like "New design" or "Create" — on most creative tools it sits near the top-left of the canvas. After you click it, a blank workspace opens. You'll see a text box where you type what you want and some style controls to the side. If your screen looks crowded instead of clean, look for a "Blank canvas" option in a menu — it's the simplest starting point.

💬 Example: type a one-line description like "warm coffee shop, hand-lettered feel, soft cream and forest green" into the prompt box.

You'll know it worked when the canvas is empty and your cursor is blinking in the prompt field, ready for instructions.

Step 2 — Generate a batch of concepts on a raster model

In the style controls (often a sidebar or a row of icons above the prompt box), pick the option that produces raster-style images — sometimes labeled "Raster," "Photo," or shown with a small photo icon rather than a vector-shape icon. Recraft describes starting with raster because it's fast and easy to compare many ideas side by side. Then describe your logo in plain words and click the button that generates results — it usually looks like a small arrow, a sparkle ✨, or the word "Generate." Expect to wait about 10–20 seconds.

💬 Example prompt: "minimalist logo for a yoga studio, lotus flower, calm pastel palette, no text in the image."

You'll know it worked when several image options appear on the canvas within half a minute. If nothing appears, look for a small "Retry" or "Generate again" link — sometimes the first attempt stalls.

Step 3 — Pick a direction and copy what works

Scroll through the generated options and click the one that feels closest to what you want. Most AI image tools let you "favorite" or "select" an image — look for a heart icon, a star, or a "Use this image" button when you hover over it. You're not done yet: this is just your starting point. Read the prompt you typed in Step 2 and tweak the words that mattered most — shape, mood, colors — so your next round focuses on what you liked. This is the part of Recraft's workflow where you "lock in" the direction.

💬 Example: if you loved the lotus shape but the colors felt too cold, change "calm pastel palette" to "warm peach and sand palette" before you continue.

You'll know it worked when you can point to the one image you want to keep, and you have a tighter prompt that names its strongest traits.

Step 4 — Switch to Recraft V4.1 Vector for the final mark

Now go back to the style controls and switch from the raster setting to the vector setting. Recraft describes V4.1 Vector as the model that returns a clean SVG instead of a pixel image. SVG is just a file format where shapes are stored as mathematical instructions — like a recipe for drawing a circle — rather than a fixed grid of colored dots. Look for a label like "Vector," "V4.1 Vector," or an icon that looks like geometric shapes rather than a photograph. Make sure your refined prompt is still in the text box, then generate again.

Example: paste your tightened prompt from Step 3 — for instance, "minimalist lotus logo, warm peach and sand palette, simple shapes only, no text."

You'll know it worked when your result looks like a clean illustration with crisp edges — you can zoom in really close without it turning blurry — and the file is described somewhere in the interface as an SVG.

Step 5 — Export the SVG and verify it at small sizes

Find the export or download control — usually an icon that looks like a downward arrow ⬇, a share icon, or a button labeled "Export" or "Download." If asked to choose a format, pick SVG (not PNG or JPG — those are raster formats that go soft when scaled). Save the file somewhere easy to find, like your Desktop. Then comes the honest test: open the SVG on your phone or in any image viewer and shrink it to about 16 pixels tall — that's the size of a browser tab favicon. If your logo still reads clearly at that size, it works. If the details turn to mush, you may need to simplify the design.

Example: open the file in a browser tab, then zoom out until it looks tiny. If the lotus shape is still recognizable, you're good.

You'll know it worked when the same SVG file looks sharp both as a 16-pixel favicon and at the size of a website header — no blurriness at either end.

Step 6 — Save variations and name your files well

Before you call it done, generate one or two more versions with small tweaks — different colors, a simpler version, or one with your business name typed in. Most tools let you download each result separately. Give every file a clear name that describes what's inside — your future self will thank you when you have six SVGs on your desktop and can't tell them apart. A simple naming pattern works: `brandname-logo-style-color.svg`, something like `lotus-logo-minimal-peach.svg`.

Example: `lotusyoga-logo-minimal-peach.svg`, `lotusyoga-logo-simple-mono.svg`, `lotusyoga-logo-with-text.svg`.

You'll know it worked when you can open any file two weeks from now and instantly know which version it is, without having to open them all.

Common mistakes

- **Skipping the raster stage and going straight for vector.** The vector model is great for the final mark, but it's slower to compare ideas. Generate lots of cheap raster options first, then commit to vector once you know what you want — that two-step rhythm is the whole point of Recraft's described workflow.
- **Exporting as PNG or JPG instead of SVG.** Raster files look fine on screen but turn blurry the moment you scale them up or shrink them down. Always pick SVG for the logo you plan to use everywhere.
- **Adding too much detail.** Tiny text, thin lines, and busy patterns can vanish when the logo is shrunk to favicon size. If your design falls apart at 16 pixels, simplify it — a strong logo usually has one or two clear shapes, not ten.

What this means for you

- **In everyday life:** you can finally make a logo for your side project, bakery, garage band, or dog-walking business without paying a designer or learning Illustrator. Twenty minutes of experimenting is enough for a first version you can actually use.
- **For your business or work:** a single SVG file travels anywhere — your website header, email signature, business card, social media profile, even the tiny favicon in the browser tab — and looks identical at every size. That means one design, no redesigns.
- **If you're just getting started:** don't aim for perfect on your first try. Generate a batch, pick the one that makes you smile, then refine it. AI logo tools reward quick iteration, not overthinking.

Try it now

Open recraft.ai, type "friendly minimalist logo for a [your project name], [two colors you like]" into the prompt box, and generate a raster batch. Spend five minutes just looking at the results and noticing which shapes and colors pull you in — that's your style direction, and you'll use it as the seed for your vector version.

How to Create Custom Images with Copilot for Invites and Social Posts

6 steps · Read online

> **In short:** Copilot can turn a plain sentence into a custom image you can use on invitations, Instagram, or a personal vision board. Sign in with a Microsoft account, open the Designer tool, describe what you want in plain words, and download what comes out.

You don't need Photoshop, Canva, or any design background to get a one-of-a-kind picture for your next event, post, or vision board. Copilot's built-in image generator — Microsoft Designer — takes a sentence from you and turns it into a finished image in roughly 30 seconds. This walkthrough shows the whole flow, including what to write so the picture actually comes out the way you imagined.

Before you start

- A **free Microsoft account** (the same one you use for Outlook, OneDrive, or Xbox). If you don't have one, signing up takes about two minutes.
- A **browser** (Chrome, Edge, Safari, or Firefox) or the **Copilot mobile app**. Either works.
- Roughly **10–15 minutes** for your first picture — faster once you've done it a few times.
- The image generator is free to use, but the free plan has a **daily limit** on how many images you can create. If you want a lot more images or faster generation, look into the paid **Microsoft Copilot Pro** plan on Microsoft's official pricing page for the current rate.

Step 1 — Open Copilot and switch to image mode

Go to **copilot.microsoft.com** in your browser, or open the Copilot app on your phone. Sign in with your Microsoft account if you aren't already. Look for a chat box near the bottom of the screen that says something like **"Ask me anything..."** — and right next to it, find a small **image or picture icon** 🖼️, or a toggle labeled **"Designer"** or **"Create an image."** Tap or click that option. The chat box changes to invite an image description instead of a question, and a small note about Microsoft's content policy may pop up the first time. 💬 Example: don't type anything yet — just make sure you can see the image-creation prompt box. You'll know it worked when the chat box now asks you to **describe an image** instead of asking a question.

Step 2 — Describe the picture you actually want

Click inside the new prompt box and write a sentence or two describing your image. Be specific: include the **subject**, the **style** (watercolor, photo-realistic, cartoon, flat design), the **mood** (cheerful, calm, dreamy), and the **colors** if they matter. The more detail you give, the closer the result matches what you imagined. If your screen looks different, look for a separate **"Designer" link** in the top menu, or a button labeled something like **"Image creator"** — Microsoft moves these around as they update the site. 💬 Example: "A watercolor birthday invitation background with soft pink peonies, gold confetti, and the words 'Emma turns 30' written in elegant script at the top." You'll know it worked when you hit Enter and Copilot starts **generating four different versions** of your image — this usually takes 15–30 seconds.

Step 3 — Review the four versions and pick your favorite

When the images finish loading, you'll see a grid of **four square images**, all based on your description. Don't rush — open each one in full size by clicking it, and check the small details that matter for your project: is the text readable, are the colors right, does the mood fit the occasion? If none of them are quite right, you don't need to start over — go back to the prompt box and describe what you'd change. Copilot can run another round based on your new description. 💬 Example: "Same image but make the peonies larger, and put the text 'Emma turns 30 — Saturday June 14' in the center." You'll know it worked when at least one of the new four images is **close enough to use**, or when you've nailed down exactly which change to ask for next.

Step 4 — Polish the prompt if the text in the image looks wrong

Here's an honest limitation: image generators — Copilot included — are **still hit-or-miss at writing words inside pictures**. Names, dates, and quotes can come out misspelled, scrambled, or warped. If your invite or vision board needs accurate text (like a name or a date), the safest workflow is to **leave the text out of the prompt entirely**, then add it yourself later in a free tool like Canva, PowerPoint, or even your phone's photo editor. This is one extra step, but it guarantees your text actually says what it should. 💬 Example: for the invite, ask for "A watercolor birthday invitation background with soft pink peonies and gold confetti, no text." Then you add "Emma turns 30 — Saturday June 14" yourself. You'll know it worked when you have a beautiful background image **without any text problems to fix**.

Step 5 — Download the image you want to use

Hover your cursor over the image you like best. A small toolbar appears underneath — look for a **download icon** (a downward arrow into a tray) or a button labeled **"Download"**. Click it once. Your browser saves the image to your normal Downloads folder, usually as a PNG or JPG file. On the phone app, the download button might say **"Save"** or look like a square with an arrow — if you can't find it, take a screenshot instead, which works on every phone. 💬 Example: click the **Download** button below your favorite peony image and wait for the file to appear in your

Downloads folder. You'll know it worked when you can open the file on your computer or phone and see the finished picture, ready to drop into your design.

Step 6 — Drop the picture into your project

Open whichever tool you normally use — Canva, PowerPoint, Google Slides, Instagram, your phone's photo editor — and **import or insert** the image you just downloaded. Resize it to fit your canvas: square (1:1) for Instagram, portrait (9:16) for stories, landscape (16:9) for invitations printed on letter-size paper. If you want a different shape than the square Copilot gave you, you can either crop with the same tools or describe a different aspect ratio in Step 2 next time (some Copilot versions accept hints like "wide landscape format" or "tall poster format"). Example: in Canva, click **"Uploads"**, drag in your downloaded peony image, drop it onto a square Instagram template, and add the text "Emma turns 30 — details in bio" on top. You'll know it worked when your finished design **looks unique to you** and is ready to send, post, or print.

Common mistakes

- **Asking for text inside the image.** Image generators often misspell names and dates. Strip the text out of the prompt and add it yourself in Canva or your editor — it'll be sharper and faster.
- **Vague prompts.** "A nice background" gives you something generic. **"Watercolor peonies, pink and gold, birthday mood, square"** gives you something usable.
- **Stopping after the first round.** The first four images are a starting point, not the finish line. If they're close, ask for one specific change. The second round is usually the keeper.

What this means for you

- **In everyday life:** the next birthday invite, save-the-date, or holiday card can be made in the time it takes to brew coffee — no software to install, no design skills needed.
- **For your business or work:** social posts, internal event flyers, or vision-board slides for team workshops can be created on the fly. If you batch three or four prompts in one sitting, you have a week's worth of visuals before lunch.
- **If you're just getting started:** open Copilot today and try the birthday invite example from Step 2. Once you've made one, every future prompt is easier.

Try it now

Open **copilot.microsoft.com**, switch to image mode, and type: *"A dreamy watercolor vision board image with a calm beach at sunrise, a notebook, a cup of coffee, and soft pastel colors, no text."* Save the result you like best. That's your first custom picture — made in under five minutes.

Use Gemini on Your Phone to Pull Editable Text from Any Photo or Screenshot

5 steps · Read online

> **In short:** Open the free Gemini app, snap or attach a photo of a sign, page, or screenshot, and ask Gemini to lift the words out as copyable text. You'll need a Google account, the app on Android or iOS, and about three minutes.

You're standing in a parking lot staring at a faded sign with the rates, or you saved a screenshot of a recipe and now you want the ingredients in a real list. Gemini (Google's AI assistant app) can read the image for you and hand back the text — searchable, editable, ready to paste anywhere. This guide shows the exact taps on both Android and iPhone.

Before you start

- A Google account (the same one you might use for Gmail or YouTube). Sign-up is free if you do not have one.
- The **Gemini** app installed on your phone. Free on Android (via the Play Store) and iPhone (via the App Store). On many Android phones Gemini replaces the Google Assistant — if so, you already have it.
- A photo ready: either saved in your gallery (a screenshot, a receipt) or something you can point your camera at (a sign, a page in a book).
- Roughly **3 minutes** for the first run; faster once you know the flow.
- Honest note: on the free plan, Gemini handles a generous number of image chats per day. If you hit a limit, it will tell you, and the same feature is included in the paid plan if you ever upgrade.

Step 1 — Open Gemini and start a fresh chat

Tap the **Gemini app icon** on your home screen — a four-pointed star with pink, purple, and blue edges. The first screen is a chat page with a text box at the bottom and a few suggested prompts above it. If a previous conversation is still open, look for a small **"+"** or **pencil icon** near the text box and tap it to start fresh. If you can't see those, just force-quit the app and reopen it — a new chat loads automatically.

💬 Example: nothing to type yet — this step is just opening the app.

You'll know it worked when you see a blank chat with a blinking cursor in the text box at the bottom of the screen.

Step 2 — Add your photo (camera or gallery)

Tap the **camera icon** 📹 (or a small **picture/landscape icon** 🖼 — different phones place it slightly differently). A small menu appears offering **"Camera"** and **"Photos"** or **"Gallery"**. Pick **Camera** if you're pointing at a live sign or page; pick **Photos** if the image is already in your phone (a screenshot, a downloaded PDF page, a screenshot of a recipe). Allow camera and photo permissions the first time the app asks — the prompt only shows once. Take the photo or tap the one you want, then confirm by tapping **"Add"** or **"Use photo"**.

If your version of the app doesn't show this menu, look for a **"+"** button next to the text box and pick **Image** from the popup.

💬 Example: no text needed yet, just pick the image.

You'll know it worked when a small thumbnail of your photo appears just above the text box, and the text box is still empty waiting for your instruction.

Step 3 — Tell Gemini what to do with the image

Tap the text box and type a short instruction — or skip typing entirely by tapping the **microphone icon** 🎤 and speaking. The photo stays attached; your instruction goes below it. Gemini reads the image (the technical word is **multimodal** — meaning it can understand both pictures and words in the same chat) and replies with the text it found.

Be specific. A vague "what does this say?" works, but a clearer request gives a better answer. If the photo has a lot of clutter around the text, ask Gemini to skip the extras.

💬 Example: "Pull out all the text from this photo, line by line, exactly as written." — or — "Read the recipe in this image and list just the ingredients."

You'll know it worked when Gemini replies with the words from your photo written out as normal text (not a description like "this is an image of..."). If it gives a description instead, your prompt was too vague — rephrase with the exact wording you want.

Step 4 — Copy the text out so you can use it

Look at Gemini's reply. Long-press on the text (press and hold with your finger for about a second) until a highlight bar appears with options like **Copy**, **Select all**, and **Share**. Tap **Copy**, and the text is now on your phone's clipboard — invisible, but ready to paste. Open Notes, Messages, your email app, or Google Docs, tap into a text field, and tap **Paste**. The words from the photo appear in your new app, perfectly editable.

If long-press doesn't show a menu, look for a tiny **"Copy"** button or a ⋯ **(three-dot)** icon at the bottom-right corner of Gemini's reply message — both Android and iOS sometimes tuck the option there.

Example: paste it into Notes and you'll have a searchable copy of that faded parking sign — no more squinting at the photo.

You'll know it worked when you can paste the text into another app and the words are correct, spelled the way they were in the original image.

Step 5 — Try a quick follow-up (translate, shorten, save)

This is where it gets useful. Stay in the same chat — Gemini remembers the photo, so you don't have to attach it again. Type a new instruction and Gemini builds on the text it already pulled. Common follow-ups: translate to another language, shorten into bullet points, rewrite in simpler words, or turn it into a checklist. When the answer is what you want, tap the **share icon** on Gemini's reply to send it to another app, or long-press and copy it the same way you did in Step 4.

If the free plan ever tells you you've hit a limit for the day, the same flow keeps working on a paid plan — check Google's official pricing page for the current plan names and details.

Example: "Now translate that to Spanish." — or — "Shorten this list to the five most important points." — or — "Turn these ingredients into a grocery checklist with quantities."

You'll know it worked when Gemini's next reply is the translated, shortened, or reformatted version of the same text — proof that it's using the original photo as context, not starting over.

Common mistakes

- **Asking too vaguely.** "What does this say?" sometimes gets you a description of the image instead of the text. Always include the word *text*, *words*, or *list* and tell Gemini what to do with them.
- **Forgetting to allow camera or photo permissions.** If tapping the camera icon does nothing, open your phone's Settings, find Gemini, and switch both **Camera** and **Photos** permissions on, then try again.
- **Trying it on handwriting that's too messy.** Printed text works very well. Cursive handwriting is hit-or-miss — if Gemini misses letters, snap a flatter, more even-lit photo and try once more, or just type the few unclear words yourself.

What this means for you

- **In everyday life:** Snap a photo of a recipe card, a Wi-Fi password sticker, a museum placard, or a friend's handwritten note and turn it into text you can search, save, or share. No more typing things out by hand from a photo.
- **For your business or work:** Photograph a whiteboard after a meeting, a printed receipt, a page from a contract, or a business card at a conference and pull the text straight into an email, spreadsheet, or note. It saves a surprising amount of retyping.
- **If you're just getting started:** The easiest first move is one screenshot you already have — a recipe, an address, a quote. Attach it, ask for the text, and copy it somewhere useful. After that one try, every future photo is the same three taps.

Try it now

Open Gemini, attach any photo already in your phone's gallery — a screenshot, a receipt, a sign you snapped yesterday — and type: "Pull out all the text from this photo exactly as written." Copy the reply into your Notes app. That's the whole trick, and it's the same trick you'll use tomorrow.

How to Turn Your Saved Articles and Notes into a Personal Podcast with NotebookLM

5 steps · Read online

> **In short:** Google's NotebookLM can read up to about 50 of your own sources — links, PDFs, Google Docs, pasted text — and turn them into a roughly 10-to-20-minute audio conversation between two AI voices. It is free with a Google account, and the audio stays grounded in your sources rather than the open internet.

You've got a folder full of articles, a few PDFs you keep meaning to read, and a commute that could use better company. NotebookLM is a free Google tool built for exactly this kind of "I'll get to it later" pile. Point it at your reading list, ask for an audio overview, and two synthetic hosts will chat through the highlights while you walk, drive, or wash dishes.

Before you start

- **A Google account** — any personal Gmail address works.
- **A web browser** on a phone, tablet, or computer — NotebookLM runs in the browser, no install needed.
- **3 to 5 sources** to feed it: web links, PDFs on your computer, Google Docs from your Drive, YouTube links, or text you can copy and paste.
- **About 15 to 20 minutes total** for setup plus the first audio generation (the audio itself usually takes a few minutes to cook).
- One honest note: the audio is a *summary* of your sources, not a verbatim read-aloud. If you need every word spoken, this is not the right tool.

Step 1 — Open NotebookLM and create a fresh notebook

In your browser, go to the NotebookLM homepage and sign in with your Google account when prompted. Once you're in, look for a button that says something like **New Notebook** (it may say *Create* or show a plus-sign +). Click it. You'll see a clean workspace with a left sidebar, a center area that invites you to add a source, and a small notes section. If your screen looks different — Google updates its products often — look for a clearly empty page with wording like **Add your first source** in the middle.

Example: type a short name like "Weekend reading" or "Q3 research" into the notebook title field so you can find it later.

You'll know it worked when a new notebook opens with an empty source list and the Studio panel on the right is ready.

Step 2 — Drop in your first article or PDF

Look for a button or panel labeled **Add source** (it might say *Upload* or *Insert*). Click it and you'll see a small window offering several options: paste a website link, upload a file from your computer, pick a Google Drive document, paste raw text, or add a YouTube link. Pick whichever matches what you saved. For a typical "I've been meaning to read this" article, copy the URL from your browser and paste it in. A short loading indicator appears while NotebookLM reads the page, then the source shows up in the left-hand Sources list with a small icon and the page title.

Example: paste the URL of one long article you've been putting off.

You'll know it worked when the source appears in the list with its title — usually within a few seconds for a normal web page, longer for big PDFs.

Step 3 — Build out your reading list

Repeat the same **Add source** action for each article, PDF, or document you want covered. Aim for **3 to 10 sources** in one notebook — enough variety for a meaty conversation, not so many that the summary loses focus. Each source gets its own row in the Sources list. You can mix types freely: two web articles, one PDF report, one Google Doc of meeting notes, and one YouTube link all work together.

Example: add 4 items — a long-form article, a PDF report you downloaded last week, a Google Doc with your own notes, and one more link.

You'll know it worked when every source shows up in the Sources panel and none of them show a red error icon. If one fails to load, remove it and try again — usually a sign-in or link problem.

Step 4 — Generate the Audio Overview

Now look at the **Studio** panel on the right side of the screen. Inside it, look for an option labeled something like **Audio Overview**. Click the button that says **Generate** (the exact wording can change as Google updates the product). A small status message appears — usually "Generating audio overview..." — and NotebookLM begins writing a script between two AI hosts, then voicing it. This usually takes **2 to 5 minutes**, depending on how much material you fed it. You can leave the tab open and do something else; it'll be ready when you come back.

Example: type an optional instruction in the custom prompt box if you see one, such as "Focus on the parts about pricing strategy."

You'll know it worked when a small audio player appears in the Studio panel with a play button and a total length (typically between 8 and 25 minutes).

Step 5 — Listen, save, or try again

Press the **play button** ▶ on the audio player. Two voices start chatting through your sources — one asks questions, the other explains, they bounce ideas back and forth like a real conversation. You can pause, drag the slider, and adjust playback speed. Look for a small download or share button near the player if you want to save the file for offline listening; if you don't see one, the audio stays inside NotebookLM for that session. If the summary misses the point or sounds off, click **Regenerate** (or a similar refresh icon) — each generation is a fresh script, and quality varies.

💬 Example: while listening, jot down a question in the Notes section on the right so you can ask NotebookLM's chat about it after.

You'll know it worked when the two hosts are clearly discussing the *content* of your sources — naming the same people, companies, or ideas from your PDFs and links — instead of generic filler.

Common mistakes

- **Pasting 30 sources at once.** The summary gets thinner and more generic the more you add. Stick to a tight theme: 3 to 10 sources that actually belong together.
- **Trying to feed paywalled or login-only pages.** NotebookLM can only read what's publicly accessible on the open web. For paywalled articles, download the PDF or paste the text manually.
- **Expecting the audio to quote your sources verbatim.** It summarizes and discusses them — it does not read them aloud. If you need a literal read-aloud, this tool won't deliver.

What this means for you

- **In everyday life:** Got a stack of "read later" articles about a topic you care about — gardening, parenting, fitness — feed them in and listen on a walk. It's a friendlier way to catch up than scrolling.
- **For your business or work:** Drop in your project briefs, meeting notes, and a few industry reports. The audio overview becomes a quick briefing you can listen to on the way to a meeting.
- **If you're just getting started:** Pick one notebook, one topic, and three sources you've already saved. That's all you need to hear whether the format clicks for you.

Try it now

Open NotebookLM in your browser, create a new notebook, and paste in the link to one article you've been meaning to read. Hit the Audio Overview generate button. By the time you make a cup of coffee, you'll have your first mini-podcast.

How to Design a Simple Logo for Your Side Project with Midjourney

5 steps · Read online

> **In short:** You can use Midjourney to sketch a clean, simple logo for a side project in well under an hour, even with zero design experience. The catch: Midjourney produces pixel images, so you'll likely want one quick free tool to convert your favorite into a scalable vector before you print or stitch it on anything.

You don't need Photoshop, a design degree, or a fat wallet to get a logo you can actually use for a side project. With Midjourney and a few minutes of typing, you can sketch something clean enough for an Etsy banner, the corner of a homepage, or a T-shirt. Here's how to do it without the guesswork.

Before you start

- A paid Midjourney subscription. Midjourney no longer offers a free tier, so check the official pricing page at midjourney.com for the current plans.
- Either a browser or a Discord account. Midjourney now runs in a web app at midjourney.com, and it still works inside Discord. Pick whichever you find less confusing.
- A rough idea of your project: the name, the mood (calm, playful, bold), and one or two colors you'd be happy to live with.
- About 30–60 minutes total, including the optional vector step at the end.

Step 1 — Open Midjourney and find the prompt box

Go to midjourney.com in your browser and sign in, or open Discord and join a Midjourney-enabled server. On the web app, the main page is an empty canvas with a single text input near the bottom — that's where your prompt goes. In Discord, open one of the rooms labeled something like "general" or "newbies," and the message box at the bottom of the chat is where you'll type. In either case, the screen looks mostly empty: a big blank area on top and a single typing field at the bottom. Example: in Discord, type a forward slash in the message box and you'll see a menu. Pick /imagine from the list. You'll know it worked when the prompt box is clearly active and ready for your description.

Step 2 — Describe your logo in one clear sentence

In the prompt box, write a short description of what you want your logo to look like. Be specific about shape, color, and mood, but keep it simple — think of it like describing a logo to a friend over coffee, not writing a novel. Every word in the prompt matters, so vague prompts give vague results. After you press Enter or click Send, wait roughly 30–60 seconds while the tool thinks. Example: "minimalist line drawing of a coffee cup with steam, single dark color on white background, modern, clean, simple" You'll know it worked when a grid of four square images appears on your screen.

Step 3 — Refine with variations

Below each of the four generated images, look for a row of small buttons — something like V1, V2, V3, V4 (variations) and U1, U2, U3, U4 (upscale). Click the V button under the image you like best — V2 if you liked the top-right one, for example. The screen starts generating four new images that riff on the one you chose, keeping the parts you liked while changing the rest. If you don't see V buttons, look for words like "Vary" or "Variations" near the image; the exact labels shift a bit between the web app and Discord. Example: if you liked the third image in the first row, click V3. You'll know it worked when four new, related images appear in a fresh grid.

Step 4 — Upscale the one you love

Now click the U button (U1, U2, U3, or U4) under the image you want to use as your final logo. This asks Midjourney to make a higher-resolution version of just that one image. The other three images fade back, and a single larger version replaces your choice. Larger files look crisp on a website and don't go blurry when you stretch them across a banner. If you don't see U buttons, look for an "Upscale" label or a small magnifier icon near your chosen image. Example: if your favorite is in the bottom-right slot of the latest grid, click U4. You'll know it worked when a single, larger image appears and a download button (usually a small arrow or picture icon) shows up nearby.

Step 5 — Prepare the image for real-world use

Midjourney gives you a pixel-based image — fine for websites and social media, but for printing on shirts, mugs, or large signs you'll want a vector version (a "smart" file that stays sharp at any size, the way a road sign looks crisp from any distance). Open the free desktop app Inkscape, or any free online "image to vector" tool, and upload your saved logo. The tool traces the shapes and gives you an SVG or PDF file you can resize forever. For websites and social media, the upscaled PNG you already downloaded is usually plenty. Example: in Inkscape, select the image, choose Path → Trace Bitmap, click OK, then save as SVG. You'll know it worked when you have both a PNG (for web) and an SVG or PDF (for print) saved on your computer.

Common mistakes

- **Writing a prompt that's too long or poetic.** Long, flowery prompts confuse the tool. Stick to 10–25 words describing shape, color, and style. Fix: write it like a recipe — noun, color, mood, background.
- **Asking for tiny text or fine detail.** AI image tools struggle with spelling and thin lines, and they vanish at small sizes. Fix: leave the words out of your first logo and add them later in a free tool like Canva, or keep all shapes chunky.
- **Skipping the upscale step and using the small grid image directly.** It looks fuzzy on a phone. Fix: always click the U button before downloading, and convert to vector for any print use.

What this means for you

- **In everyday life:** a clean logo lifts your baking blog, your kid's sports team, your holiday card, or a small Discord server — anywhere a tiny mark of identity makes the thing feel finished.
- **For your business or work:** it gives your Etsy shop, freelance site, or weekend app a polished first impression without paying a designer for the very first version. You can always hire a designer later to refine the one you like.
- **If you're just starting:** skip words entirely on your first try. A simple shape with one strong color looks more professional than a busy mark with letters you can't read. And remember to check Midjourney's current terms on the official site before you sell anything featuring the logo — the rules around commercial use can change.

Try it now

Open midjourney.com (or the Discord bot), type this exact prompt, and see what comes back: "minimalist mountain icon, geometric, single deep teal color on white background, modern, clean, simple". Pick the one you like best, upscale it, and save the result to your desktop. You just made your first logo.

Learning & getting started

Practice a new language or presentation with Claude for instant feedback

7 steps · Read online

Ever wished you had a patient tutor or a friendly colleague on standby to help you practice? Claude, an AI assistant, can be just that. This guide will show you how to use Claude to simulate conversations for language learning or get immediate, constructive feedback on your presentation ideas and wording, helping you build confidence.

Before you start

- You'll need a web browser (like Chrome, Edge, Safari, or Firefox).
- An internet connection.
- A free Claude account (signing up is quick, usually just an email and password or Google login).
- This guide assumes you're using the free web version of Claude. While there are paid plans with higher usage limits, the free version is perfect for getting started with this task.
- The whole process should take about 5 to 10 minutes to get through the steps, plus however long you want to practice!

Step 1 — Open the Claude website

First, you'll need to open Claude in your web browser. This is where you'll interact with the AI assistant, which is like a digital helper that can understand and generate human-like text. Type `claude.ai` into the address bar at the top of your browser window and press Enter.

You'll see the Claude homepage, possibly asking you to log in or sign up. If you're not logged in, either create a new account or sign in with your existing details. What happens next is you'll be taken to the main chat interface. If it looks different, look for a "Log In" or "Sign Up" button, usually in the top right corner. You'll know it worked when you see a chat window, ready for you to type a message.

Step 2 — Start a new conversation

Now that you're in Claude, let's begin a fresh interaction. Look for a button that says "Start a new chat" or simply "New Chat", often found on the left side of the screen, or sometimes as a plus (+) icon. Click this button to clear any previous conversations and start fresh.

You'll see a blank chat area, with a text box at the bottom that says something like "Message Claude...". This is where you'll type your instructions, known as a 'prompt', to tell Claude what you want it to do. If it looks different, try clicking any existing chat entry to open it, then look again for the "New Chat" option, or a similar icon to begin a fresh conversation. You'll know it worked when the main chat window is clear, and the message box at the bottom is empty, waiting for your input.

Step 3 — Set the scene for language practice

Let's tell Claude to become your language tutor. In the message box at the bottom, you'll type an instruction that sets Claude's role. This is called *role-playing*, where the AI pretends to be someone specific to help you. Be sure to specify the language and your learning goal.

You'll see Claude process your request, and then it will respond, confirming its role and asking how you'd like to start. If it looks different, and Claude seems confused, try rephrasing your instruction more simply. You'll know it worked when Claude replies, saying something like "Hello! I'm ready to be your French tutor. How can we begin?"

Example: `Act as a friendly French tutor for a beginner. Ask me simple questions about my daily life in French, and gently correct any grammar or vocabulary mistakes I make, explaining why.`

Step 4 — Begin your language practice

With Claude in its tutor role, it's time to start speaking (or rather, typing!) the new language. Claude will have likely asked you a question in the target language. Type your answer directly into the message box, doing your best to use the language you're learning.

You'll see Claude read your response, and then it will provide feedback. This might include corrections, explanations, and perhaps another question to keep the conversation going. If it looks different, and Claude just gives a generic response, remind it of its role: "Please remember, you are my French tutor." You'll know it worked when Claude gives you specific feedback on your language use and continues the conversation.

Step 5 — (Alternative) Prepare for presentation feedback

If you're practicing a presentation, let's reset Claude's role. Click "New Chat" (from Step 2) to start fresh. This time, you'll instruct Claude to act as a presentation coach. Be clear about the kind of feedback you need, such as clarity, tone, or structure.

You'll see Claude acknowledge your request and usually ask for the presentation content you want reviewed. If it looks different, and Claude seems to be waiting for more instruction, tell it, "I will now give you parts of my presentation script." You'll know it worked when Claude replies

with something like, "Understood! I'm ready to act as your presentation coach. Please share your script, and tell me what kind of feedback you're looking for."

Example: `You are my presentation coach. I need honest, constructive feedback on my script's clarity, conciseness, and how engaging it sounds. Imagine I'm presenting to a general audience. I'll share parts of my script one by one.`

Step 6 — Deliver parts of your presentation

Now, paste the first part of your presentation script into the message box. It's often best to give it section by section (like your introduction or a key point) rather than the whole thing at once, so Claude can give focused feedback.

You'll see Claude analyze the text, then provide its feedback based on your instructions. This feedback will highlight areas for improvement or suggest alternative phrasing. If it looks different, and Claude just summarizes your text, re-emphasise: "Please provide feedback on clarity and conciseness, not just a summary." You'll know it worked when Claude offers specific suggestions or praises effective parts of your script.

Step 7 — Refine and get more feedback

After Claude gives its feedback, you can continue the conversation to refine your work. You might ask: "How could I make that introduction more impactful?" or "Can you suggest a stronger closing sentence for this paragraph?" You can also paste revised sections to see if the changes improved things. This back-and-forth interaction allows you to iterate and improve.

You'll see Claude respond to your specific questions, offering new suggestions or confirming that your revisions are effective. This helps you hone your skills. If it looks different, and Claude repeats earlier feedback, try being more specific with your follow-up questions or rephrase what you're asking. You'll know it worked when Claude provides fresh, actionable advice that helps you move forward with your practice.

Common mistakes

- **Vague prompts:** If you don't tell Claude exactly what role to play or what kind of feedback you need (e.g., "Just help me learn French"), the responses might be too general. Always be specific about the language, your level, or the type of presentation and audience.
- **Expecting human-like intuition:** While Claude is smart, it's an AI. It can't read your mind, understand unspoken nuances, or see your body language during a presentation. Stick to text-based interactions and explicitly state your goals.
- **Overloading the AI:** Pasting a whole 20-minute presentation script at once might overwhelm Claude, leading to less detailed feedback or hitting free usage limits. Break it down into smaller, manageable chunks for better, more focused responses.

What this means for you

- **In everyday life:** Claude can be your personal language coach, helping you prepare for travel or converse with friends in a new tongue. It's a no-pressure way to improve your skills on your own schedule.
- **For your business or work:** Use Claude to refine sales pitches, team meeting updates, or client proposals. Getting clear, concise feedback can make your communication more effective and help you present with more confidence.
- **If you're just getting started:** The easiest first move is to pick one goal – either practicing a simple conversation in a new language or getting feedback on just the introduction of a presentation – and stick to that. You'll be surprised how quickly you can make progress.

Try it now

Open Claude, start a new chat, and try asking it to act as an English proofreader. Then, type a sentence you're unsure about and ask it for suggestions to improve clarity. See how it helps!

Build a personal study-buddy chatbot in 15 minutes using Hugging Face Spaces (no code)

7 steps · Read online

Imagine having a patient AI tutor that never gets tired, remembers what you've learned, and quizzes you just before you forget. With Hugging Face Spaces, you can build that tutor in under 15 minutes — no coding required. Perfect for students, lifelong learners, or anyone who wants to remember more with less effort.

Before you start

- A **free Hugging Face account** (sign up at huggingface.co)
- A **computer or phone with internet**
- A **quiet 15 minutes** (you'll be guided step by step)
- No prior coding experience needed

> ⚠️ This guide uses the free tier of Hugging Face. Some advanced features may require a paid plan, but the core steps work for free.

Step 1 — Sign in to Hugging Face and open Spaces

Go to huggingface.co and tap **Sign Up** in the top-right corner if you don't have an account. After signing in, click **Spaces** in the top menu (it's next to "Models" and "Datasets"). This is where you'll create your study-buddy.

You'll see a list of public Spaces created by others. Look for a blue button that says **"New Space"** in the top-right — tap it.

💬 Example:

> You can copy this into your browser: https://huggingface.co/spaces

You'll know it worked when you see a page with a form asking for your Space name and license.

Step 2 — Name your Space and choose settings

On the "Create new Space" page, fill in these details:

- **Space name**: Type something like `my-study-buddy` (use lowercase letters and hyphens only).
- **Visibility**: Keep it **Public** so you can share it with friends or revisit it later.
- **Hardware**: Choose **CPU** — it's free and fast enough for a chatbot.
- **SDK**: Select **Gradio** (it's the most beginner-friendly option and works great for chatbots).

Click **Create Space**.

💬 Example:

> Space name: `my-study-buddy` Visibility: Public Hardware: CPU SDK: Gradio

You'll know it worked when you land on a new page titled "Your Space is being created..." and you see a progress bar.

Step 3 — Wait for your Space to build

Your Space is now being set up in the cloud. This usually takes 1–2 minutes. You'll see a spinning icon and a message like "Your Space is being created..."

While you wait, think about what you want your study-buddy to help with — for example, biology terms, math formulas, or language vocabulary.

You'll know it worked when the spinning icon disappears and you see a green banner saying "Your Space is ready!"

> ⚠️ If the build fails, refresh the page. If it keeps failing, check your Space name uses only lowercase letters and hyphens.

Step 4 — Open the Space editor

Once your Space is ready, you'll see a page with tabs like **Files**, **Settings**, and **App**. Click the **Files** tab.

You'll see a file called `app.py`. Click it to open the editor.

💬 Example:

> You'll see a single line: `def main():`

You'll know it worked when you see a simple code editor with a few lines of starter code.

Step 5 — Replace the starter code with a study-buddy template

Delete everything in the editor and paste this exact code:

```
import gradio as gr

def respond(message, history):
    # Simple study-buddy response
    if "define" in message.lower():
        return "A **study buddy** is an AI assistant designed to help you learn and remember information through conversation, quizzes, and
    elif "quiz me" in message.lower():
        return "Quick quiz: What is the capital of the US? (Type your answer)"
    elif "bye" in message.lower():
        return "Good luck with your studies! Remember: spaced repetition is key to long-term memory."
    else:
        return "I'm here to help you study! Try asking: 'Define study buddy' or 'Quiz me'."

demo = gr.ChatInterface(
    respond,
    title="📚 My Study Buddy",
    description="Ask me to define terms, quiz you, or explain concepts — I'll adapt to your learning pace.",
    theme="soft"
)

demo.launch()
```

Click **Commit changes to main** (a green button at the bottom).

💬 Example:

> You can copy the code above and paste it directly into the editor.

You'll know it worked when the page refreshes and you see your chat interface appear below the editor.

Step 6 — Test your study-buddy in the chat

In the chat panel that just appeared, type:

> Define study buddy

You should see the AI explain what a study buddy is.

Then try:

> Quiz me

The AI will ask you a question. Type your answer and press Enter.

💬 Example:

> Type: "Define study buddy" Expected response: "A study buddy is an AI assistant designed to help you learn..."

You'll know it worked when the chat responds naturally to your messages and keeps the conversation going.

> ⚠️ If the chat doesn't respond, refresh the page. If it still doesn't work, check that you pasted the code exactly as shown.

Step 7 — Customize your study-buddy (optional)

Want to teach it about a specific subject? Replace the `respond` function with your own rules. For example, to teach it about photosynthesis:

```
def respond(message, history):
    if "photosynthesis" in message.lower():
        return "Photosynthesis is the process by which green plants use sunlight to make food from carbon dioxide and water. In simple terms
    elif "quiz me" in message.lower():
        return "Quick quiz: What gas do plants release during photosynthesis? (Type your answer)"
    else:
        return "Ask me about photosynthesis or type 'quiz me' to test your knowledge!"
```

Click **Commit changes to main** again.

💬 Example:

> Replace the entire `respond` function with the code above.

You'll know it worked when the chat now responds to questions about photosynthesis.

> ⚠️ This feature is beginner-friendly, but you'll need to write simple rules. If you're not comfortable editing code, stick with the original version.

Common mistakes

- **Mistake**: Typing the wrong Space name with uppercase letters or spaces. **Fix**: Use only lowercase letters and hyphens, like `my-study-buddy`.
- **Mistake**: Choosing GPU hardware instead of CPU. **Fix**: CPU is free and fast enough for a chatbot. GPU is for advanced users.
- **Mistake**: Not waiting for the Space to build. **Fix**: Wait until you see the green "Your Space is ready!" banner before proceeding.

What this means for you

- **In everyday life**: Build a personal AI tutor for any subject — biology, history, languages, or even your work training. Share it with classmates or revisit it whenever you need a study session.
- **For your business or work**: Use it to create onboarding quizzes or internal knowledge checkers. No need to hire a trainer — your AI can do it.
- **If you're just getting started**: This is your first step into AI-powered learning. Once you're comfortable, try adding more subjects or even voice chat.

Try it now

Open your new study-buddy chatbot at: **https://huggingface.co/spaces/[your-username]/my-study-buddy**

Replace `[your-username]` with your actual Hugging Face username. Ask it to define a term or quiz you — and see how it adapts to your learning style.

Build a simple weekly study plan with Meta AI in 10 minutes

7 steps · Read online

Imagine opening your calendar on Sunday night and seeing every exam, quiz and assignment laid out in color-coded blocks. You'll know exactly what to study each day and when to take breaks. Meta AI can build that plan for you in under ten minutes, even if you're starting from scratch.

Before you start

- A phone, tablet or computer with the **Meta AI app** (iOS/Android) or **Meta website** open in a browser.
- A free **Meta account** (sign-up takes 2 minutes).
- Your exam dates, quiz dates and assignment deadlines written down or in your calendar.
- Roughly **10 minutes** of quiet time.

Step 1 — Open Meta AI and start a fresh chat

Open the **Meta AI app** on your phone or go to meta.ai in any browser. On the home screen you'll see a search bar and a list of recent chats. Tap the **plus icon** (a circle with a pencil inside) or the "New chat" button to begin a fresh conversation. A blank chat screen appears with a blinking cursor ready for your first message.

You'll know it worked when you see a clean chat window and the prompt "Ask Meta AI anything..." at the bottom.

Example: *Just open the app or website.*

Step 2 — Tell Meta AI your courses and dates

In the chat box, type a clear message that lists each course, its exam date, any quizzes, and assignment deadlines. Use simple bullet points so the AI can read them easily.

For example: "Here are my courses and dates:

- Math exam 15 June
- History essay due 20 June
- Chemistry quiz 12 June
- English speech 25 June"

After you press send, Meta AI will reply with a short summary confirming it understood everything.

You'll know it worked when you see a reply that repeats your courses and dates back to you.

Example: *Math exam 15 June, History essay due 20 June, Chemistry quiz 12 June, English speech 25 June*

Step 3 — Ask for a weekly study plan

On the same chat line, type: "Create a 7-day study plan that spreads the work evenly and includes short breaks."

Meta AI will reply with a color-coded weekly calendar. Each day shows which topics to study and when to take 10-minute breaks. If the reply feels too long, ask for a shorter version by typing "Make it more compact."

You'll know it worked when you see a calendar with dates, study blocks, and break reminders.

Example: *Create a 7-day study plan that spreads the work evenly and includes short breaks.*

Step 4 — Adjust the plan to fit your life

Look at the plan Meta AI gave you. If an evening block feels too late, reply: "Move the Math block to 4 pm each day instead."

If you already have a part-time job on Wednesdays, say: "Remove Wednesday evening and shift those hours to Thursday morning."

Meta AI will update the calendar instantly. If the button you expect isn't there, look for a small "Edit" or "Regenerate" link near the plan.

You'll know it worked when the calendar updates and the new times appear.

Example: *Move the Math block to 4 pm each day instead.*

Step 5 — Ask for topic breakdowns

Pick one course from the plan and ask Meta AI to break it into smaller topics. For example: “For Math, list the exact chapters and topics I should study each day.”

Meta AI will reply with a detailed checklist. If the list feels overwhelming, ask for “only the top three topics per day.” If you prefer flashcards, say “turn the list into flashcard questions.”

You’ll know it worked when you see a clear checklist or flashcard list for that day.

Example: *For Math, list the exact chapters and topics I should study each day.*

Step 6 — Save or export the plan

At the bottom of the chat, look for a small **share icon** (an arrow pointing up or three connected dots). Tap it and choose “Copy” to save the plan to your phone’s notes or “Save as image” to add it to your calendar app. If you don’t see the icon, look for a menu with “Export” or “Share.”

You’ll know it worked when you can paste the plan into another app or save it as an image.

Example: *Tap the share icon and choose “Copy.”*

Step 7 — Set phone reminders

Open your phone’s calendar or reminder app and paste the copied plan into tomorrow’s events. Set each study block as a calendar event with a 10-minute alert. If you prefer voice reminders, ask Meta AI to “turn the plan into voice reminders” and follow the link it provides.

You’ll know it worked when tomorrow’s events appear in your calendar with study blocks and alerts.

Example: *Paste the plan into tomorrow’s calendar events.*

Common mistakes

- **Too vague a prompt:** Writing “Plan my study” gives a messy reply. Fix: list your courses and dates clearly.
- **Ignoring breaks:** Meta AI may pack blocks too tightly. Fix: reply “add a 10-minute break every 50 minutes.”
- **Not editing:** The first plan might clash with your job or sport. Fix: ask Meta AI to shift or remove blocks until it fits.

What this means for you

- **In everyday life:** Use this plan to stay on track without last-minute cramming. You’ll finish each day knowing exactly what you accomplished.
- **For your business or work:** The same technique works for project deadlines and client deliverables—just swap “exam” for “milestone.”
- **If you're just getting started:** Start with one small course and build the habit of reviewing the plan each morning.

Try it now

Open Meta AI, paste your courses and dates, and ask for a weekly study plan. You’ll have a color-coded calendar in under two minutes.

Getting Started with Mistral's Le Chat: Your Private AI Guide

6 steps · Read online

Want to try an AI assistant but feel a bit uneasy about where your personal information might end up? This guide is for you. You'll learn how to use Le Chat (which is French for "The Cat," a playful nod to "chatting"), a friendly, privacy-focused AI assistant from Europe. By the end, you'll be able to use it confidently for everything from sorting lists to drafting emails, knowing your conversations are handled with care.

Before you start

- **A device with internet access:** This guide works whether you're using a computer, tablet, or smartphone.
- **A modern web browser:** You'll need an up-to-date browser like Chrome, Safari, Firefox, or Edge.
- **An email address:** To create your free Le Chat account.
- **Rough time commitment:** Allow about 5-7 minutes to follow these steps and complete your first chat.
- **Cost:** Le Chat offers a generous free tier that's perfect for everyday use. All features in this guide are available on the free plan. For very high-volume or advanced use, a paid plan might be offered, but it's not needed here.

Step 1 — Create your free account

To begin using Le Chat, you first need to set up a personal profile. Because its creator, Mistral, is based in Europe, Le Chat operates under strict European privacy laws, which can offer peace of mind if you're conscious of your digital footprint.

Open your web browser and search for "Le Chat Mistral AI" or navigate directly to their official website. Once the main page loads, look for a prominent button or link that says "Sign up," "Create account," or "Try Le Chat" – this is usually found in the top-right corner or the center of the page.

Example: Click the "Sign up" button. On the next screen, enter your email address in the designated field, then create a strong, secure password. Click the "Submit" or "Continue" button. You will then receive a verification email from Le Chat; open this email and click the link inside to activate your account. You'll know it worked when a confirmation message appears in your browser, and you are automatically redirected to the Le Chat chat interface, ready to explore.

Step 2 — Adjust your privacy settings

Before you type your first message, let's make sure your account is configured to keep your conversations as private as possible. By default, many **AI tools** (which are clever computer programs that can understand and generate human-like language) use your interactions to help train and improve their underlying **AI model** (the brain of the AI). Le Chat gives you clear control over this.

On the Le Chat screen, look for a settings icon, often shaped like a small **gear** ( ), or perhaps your **profile picture/initials**, typically located in the bottom-left or top-right corner. Clicking this will reveal a menu or panel with various options.

Example: From the menu that appears, select "Settings" or "Data Privacy." You will then likely see an option labeled something like "Store Chat History" or "Use my data for training." Toggle this option **off** (so it shows as gray or inactive, rather than blue or green). This tells Le Chat not to use your specific conversations to further train its AI. You'll know it worked when the toggle switch visually indicates it's turned off, usually by changing color or position, and a small notification might briefly appear confirming "Settings saved" or "Preferences updated."

Step 3 — Write your first prompt

Now that your account is set up with privacy in mind, you're ready to chat! At the bottom of the Le Chat screen, you'll see a large text box, often labeled "Message Le Chat" or "Type your prompt here." This is where you'll type your **prompt** (which is simply the instruction or question you give to the AI, telling it what to do). Let's start with a practical task: organizing a messy list. This is an easy way to see how the AI understands and structures information.

Click directly into the text box at the bottom of the screen. Your cursor will appear, indicating it's ready for you to type.

Example: Copy and paste the following text into the box, then press the "Send" button (often a paper airplane icon or similar) located on the right side of the text box:

> "Please organize this messy shopping list into neat categories like 'Fresh Produce', 'Dairy', and 'Pantry': apples, milk, carrots, cheese, pasta, bananas, butter, eggs, bread, coffee, chicken." You'll know it worked when Le Chat quickly processes your request (you might see a brief "typing" indicator) and then displays an organized list categorised as you asked, appearing as a new message above your prompt.

Step 4 — Refine the results

The beauty of AI assistants like Le Chat is that they're designed for conversation. You don't have to accept the first answer as final. If the output isn't exactly what you wanted, you can simply ask Le Chat to make adjustments, building on its previous response. This is because it remembers

the **context window** (think of this as the AI's short-term memory during your current conversation, much like you remember what you just said to a friend).

After receiving the organized shopping list from the previous step, the chat interface remains open with your last prompt and Le Chat's response visible. Simply type your follow-up instruction into the *same text box* at the bottom of the screen, just like you would continue a text message conversation.

Example: Type this follow-up message into the chat box and send it:

> "Thank you! Now, can you rewrite that list but put a little checkbox symbol [ ] next to each item so I can print it out and tick them off when I go shopping?" You'll know it worked when Le Chat quickly generates a new version of the shopping list, keeping the categories but adding a `[ ]` symbol next to each item, showing it understood your refinement and applied it to the previous output.

Step 5 — Start a fresh conversation

As you chat with Le Chat, it remembers everything within the current conversation, building on what you've discussed so far. This "short-term memory" is great for refining a single task or continuing a topic, but if you want to start a completely new, unrelated topic, it's best to begin a fresh chat. This prevents the AI from getting confused by old information or trying to connect new ideas to a previous, unrelated conversation.

On the Le Chat interface, look for a button or icon that allows you to start a new chat. This is typically labeled "New Chat," "Start new conversation," or represented by a **plus symbol** (+) or a **speech bubble icon**, often found near the top-left of the chat interface or within a side menu. Click this button to clear the current conversation.

You'll know it worked when the current conversation scrolls away (but remains accessible in your chat history, usually displayed as a list on the left-hand side), and the main chat area becomes empty and ready for a brand-new prompt, with the text box cleared and waiting for input.

Step 6 — Brainstorm ideas for a party

Le Chat isn't just for organizing lists; it can be a fantastic brainstorming partner. Let's use your new, empty chat window to generate some creative ideas. This demonstrates the AI's ability to think broadly and provide inspiration based on a simple request, showing its creative side.

With your new, empty chat screen visible, click into the text box at the bottom. The cursor will appear, indicating it's ready for your input. This time, we'll ask for some creative suggestions rather than an organizational task.

Example: Type this prompt into the box and send it using the paper airplane or "Send" button:

> "I'm planning a small outdoor birthday party for a friend who loves vintage sci-fi movies and board games. Give me five unique theme ideas for the party, including a brief description for each." You'll know it worked when Le Chat quickly responds with a list of five distinct and creative party theme ideas, each with a short description, showing its ability to generate imaginative content based on your specific interests.

Common mistakes

- **Treating it like a rigid search engine:** Don't just type keywords like "meal plan." Le Chat thrives on natural language.
  - **Fix:** Speak to Le Chat as if you're talking to a helpful assistant. Use full sentences like, "Could you please help me plan three simple dinners using chicken for this week, suitable for a family?"
- **Forgetting the context window limit:** Le Chat, like all AI, has a "short-term memory" called its **context window**. If you chat for hours about ten different topics in one go, it might start to forget the early details.
  - **Fix:** If you want to start a brand-new task that's unrelated to the current discussion, use the "New Chat" button (as shown in Step 5) to give the AI a fresh slate and prevent confusion.
- **Not being specific enough in your prompts:** If Le Chat's answer isn't quite right, the mistake often lies in your instruction, not the AI.
  - **Fix:** Instead of just saying "make it better," tell it *how* to improve. For example, "That's a good start, but can you make the tone more formal and shorten this to two paragraphs?"

What this means for you

- **In everyday life:** You can use Le Chat to quickly draft polite emails to service providers, brainstorm weekend activity ideas for the family, summarize long articles from the web, or even help you plan your shopping trips – all with the added peace of mind that your privacy is a priority. It's like having a clever, discreet personal assistant available 24/7.
- **For your business or work:** You can draft professional business correspondence, generate marketing copy ideas for social media, organize meeting notes into clear summaries, or quickly outline presentations. The strong European privacy standards mean your sensitive business ideas and communications are handled with an extra layer of protection.
- **If you're just getting started:** You don't need to be a technology expert. Le Chat is designed to be as simple and intuitive as sending a text message to a friend, but with the power of AI assistance ready to help you with a wide range of tasks. Just ask it naturally!

Try it now

Open your Le Chat window right now (or start a new chat if you're continuing from the last step) and type this simple sentence to see how natural the conversation feels:

"Give me three quick, positive suggestions for a weekend morning routine that doesn't involve looking at my phone, and explains why each is good for me."

How to Translate and Understand Other Languages with AI

6 steps · Read online

By the end of this guide, you'll be able to turn any piece of foreign text—whether a quick chat, a restaurant menu, or a PDF report—into clear English (or your preferred language) with an AI tool. You'll also learn how to keep the meaning intact, ensuring your translations are accurate and natural-sounding. This guide is for anyone who's never used artificial intelligence for translation before.

Before you start

- **An internet-connected device:** This could be a mobile phone, tablet, or computer.
- **A free account on an AI chat platform:** You'll need access to a text-generating AI like ChatGPT, Claude, or Gemini. The steps in this guide work similarly across these tools.
- **A piece of text to translate:** Have the foreign text ready to copy-paste, or displayed on your screen so you can type it out.
- **Time commitment:** This guide takes about 5-10 minutes to follow along for the first time.

Step 1 — Open your AI chat tool

First, open the website or app for your chosen AI chat platform. This is the place where you usually "talk" to the AI, typing messages and receiving replies. Once it's loaded, you'll typically see your past conversations on one side and a main chat area. To start fresh, look for a button that says **"New chat"** or a simple **plus-icon (+)**. Tapping or clicking this will clear your screen.

What you see: A blank main chat area, often with a message box at the bottom ready for your input. **What happens next:** The AI prepares a clean, new conversation thread, meaning it won't be influenced by any previous topics you've discussed. **If it looks different:** On a mobile app, the "New chat" button might be in a corner or accessible through a menu icon (three horizontal lines). On a computer, it's often in a sidebar on the left. **Example:** Click the "New chat" button. **You'll know it worked when** you see a blank message box and no previous conversation messages are visible.

Step 2 — Ask the AI to translate, adding context

In the blank message box, type a clear instruction (this is called a **prompt** – it's the specific request you give the AI) telling the AI what you need. After your instruction, paste the foreign text you want to translate. Adding details like "keep the tone friendly" or "preserve the formal style" is crucial because it tells the AI not just *what* to translate, but *how* to translate it, helping it retain the original meaning and feel.

What you see: The cursor blinking in the message box, waiting for your input. **What happens next:** As soon as you hit Enter or click the "Send" button, the AI will begin processing your prompt and generating its translation. **If it looks different:** The "Send" button might be an arrow icon, a paper airplane symbol, or simply the Enter key on your keyboard. **Example:**

```
Please translate the following Italian text into English and keep a relaxed, conversational tone:

Ciao! Come stai? Spero che la giornata vada bene. Ci vediamo stasera?
```

You'll know it worked when the AI starts typing its translated response below your message, usually appearing word by word.

Step 3 — Review the AI's initial translation

Once the AI finishes typing, carefully read the translation it provides. While AI is very good, it can sometimes miss **nuance** – subtle differences in meaning, tone, or style that depend on the specific situation or culture. Scan for any words or phrases that seem a little off, don't quite make sense, or don't match the tone you expected. Don't worry if it's not perfect yet; the next step will fix it.

What you see: The AI's full translated text displayed on your screen. **What happens next:** You mentally compare the AI's output with your understanding of the original text, looking for any potential improvements. **If it looks different:** The AI might offer multiple versions, or ask a clarifying question. If it does, read its question and answer it in the message box. **Example:** If the AI translated "Ciao! Come stai?" as "Hello! How are you?" **You'll know it worked when** you have identified any parts of the translation that might need further adjustment, based on your original instruction or intent.

Step 4 — Refine the translation with specific feedback

If something in the initial translation didn't quite hit the mark, don't just accept it. Reply to the AI in the same conversation, giving it clear, specific feedback. You can highlight a particular word, ask for a different tone, or even provide additional context that the AI didn't have before. This helps the AI learn and give you a more accurate and contextually appropriate translation.

What you see: The previous AI translation and a new blank message box below it. **What happens next:** The AI will use your feedback to revise its previous answer, often providing a much better translation that addresses your specific concerns. **If it looks different:** The AI might offer an explanation for its initial choice before providing a revised version. **Example:** Following on from the previous step, you might type:

```
Could you make the greeting even more casual, like "Hey there! How's it going?" for "Ciao! Come stai?" And ensure the entire message keeps t
```

You'll know it worked when the AI provides a new version that makes more sense, sounds more natural, or perfectly matches the tone you intended.

Step 5 — Translate longer documents using “chunking”

For larger documents like multi-page PDFs or long articles, consider breaking the text into smaller sections rather than pasting everything at once. This method is called **"chunking"**. Modern AI tools can handle much larger texts than before (many support tens of thousands of words in a single message), but chunking still helps maintain accuracy, especially for complex or technical texts, and makes it easier to review each section.

What you see: Your larger document open on one side of your screen, ready for you to select and copy text. **What happens next:** You copy a manageable section of text, then repeat steps 2-4 for each "chunk," ensuring each part is translated and refined before moving to the next. **If it looks different:** The exact method for copying large text varies by program; you might need to drag your mouse, use keyboard shortcuts (Ctrl+A then Ctrl+C on Windows, Cmd+A then Cmd+C on Mac), or use a "Select All" option. **Example:**

```
Translate the following section of a legal document into plain English, maintaining a formal and precise tone, without using any legal jargo
```

```
[Paste the first few paragraphs of your legal document here]
```

You'll know it worked when the AI successfully translates that section without error or truncation, providing a complete response to that specific chunk.

Step 6 — Save your translated text in a useful format

Once you're satisfied with a translated section (or the entire document, if you were chunking), select the AI's answer and copy it. Then, open a document program on your device (like Microsoft Word, Google Docs, Apple Pages, or even a simple Notepad application) and paste the text there. Remember to save your new document with a clear, descriptive name so you can easily find it later.

What you see: The AI's translated text highlighted on your screen, ready to be copied. **What happens next:** The text is stored in your device's clipboard, ready to be pasted into another application. **If it looks different:** Most AI tools have a "Copy" icon next to their response, or you can highlight the text and right-click to choose "Copy" (or use keyboard shortcuts: Ctrl+C on Windows/Linux, Cmd+C on Mac). **Example:** Highlight the AI's translation, press **Ctrl + C** (or **Cmd + C** on Mac), open a new Google Doc, and press **Ctrl + V** (or **Cmd + V** on Mac) to paste. **You'll know it worked when** you see the translated text appear in your chosen document program, ready for you to save.

Common mistakes

- **Forgetting to add context:** Simply asking "Translate this" without specifying the tone or purpose can lead to a translation that feels inappropriate or robotic.
  - **Fix:** Always include a phrase like "keep a friendly tone," "translate for a business email," or "explain in simple terms for a beginner."
- **Pasting too much text at once:** While modern AI tools have large capacity, trying to translate an entire book in one go can sometimes result in the AI missing details or giving incomplete answers.
  - **Fix:** For very long texts, consider breaking them into smaller sections and translate them one by one, keeping them in the same chat conversation to maintain context.
- **Accepting the first translation blindly:** Even advanced AI can misinterpret something, especially if the original text is highly specialized or uses slang.
  - **Fix:** Always read through the AI's translation and, if anything feels off, ask the AI to rephrase or clarify it with more specific instructions.

What this means for you

- **In everyday life:** You can instantly understand a foreign menu, decode a message from a friend overseas, or get the gist of a news article in another language. No more guessing what "Risotto ai funghi" means!
- **For your business or work:** You can quickly translate client emails, understand product specifications, or review documents from international partners, saving you time and potentially avoiding the need for a professional translator for initial drafts.
- **If you're just getting started:** The easiest first step is to open your AI chat, paste a short sentence in another language (like "Hola, ¿cómo estás?"), and ask for a simple translation into English. You'll see the AI do the heavy lifting in seconds.

Try it now

Take a piece of foreign text you have handy—maybe a label on a food item, a line from a song, or a snippet from an article online. Open your preferred AI chat tool, type the prompt "Please translate the following into English and keep it casual:" and then paste your text. Hit send, and watch the AI instantly provide its translation!

Use ChatGPT to Practice a Foreign Language Without Feeling Awkward

6 steps · Read online

> **In short:** ChatGPT can become a patient conversation partner for almost any language — it speaks at your level, corrects you kindly, and never gets bored. You'll need a free ChatGPT account and the app or website; you'll learn the most by speaking out loud.

Practicing a new language in front of another human can feel awkward — especially when you're worried about sounding silly. A chatbot removes that pressure. There's no raised eyebrow when you mix up "ser" and "estar," and you can ask the same question five different ways without feeling like you're wasting anyone's time. This guide walks you through setting up a calm, private practice session you can repeat whenever you have a few spare minutes.

Before you start

- **A free ChatGPT account** is enough to begin. Voice practice and longer conversations may need a paid plan — if a feature asks you to upgrade, that's the moment to know.
- **The ChatGPT app on your phone** is the easiest way to talk out loud. The website on a laptop also works, especially if you'd rather type.
- **The language you want to practice.** If you only have a vague idea ("maybe Spanish"), that's fine too — pick something specific in Step 1.
- **Headphones with a microphone**, if you're in a shared space and don't want others to hear you practicing.
- **Roughly 15 minutes** for the first setup. After that, every practice session takes as little as five minutes.

Step 1 — Open a new chat and name it for the language

Open the **ChatGPT app** on your phone, or visit chatgpt.com on a browser, and tap the **"New chat"** button. On the phone app this is usually a pencil-and-paper icon at the top-right; on the website it's a button at the top-left. A blank conversation opens with a message box at the bottom that says "Ask anything" or "Message ChatGPT". Now rename the chat so you can find it again later. Tap the **chat title at the very top of the screen** — it usually shows today's date by default — and a small text field appears. Type your chosen name and tap **Done**, **Save**, or the checkmark to confirm. If tapping the title doesn't do anything, look for a **pencil icon** ✏️ next to it, or open the **three-dot menu** ⋯ at the top and choose "Rename". Naming the chat keeps all your practice in one tidy place. 💬 Example: name the chat *"Spanish practice — beginner"*. You'll know it worked when the top of the screen shows your chosen name instead of today's date, and the chat is ready for your first message.

Step 2 — Tell ChatGPT exactly how you want to practice

Tap the message box at the bottom of the screen and type a short "briefing" message explaining what you want. Mention your level, the language, the kind of situation you want to rehearse, and how you'd like corrections — the more specific you are, the more useful the reply will be. ChatGPT answers well when you give it a small "job description," almost like briefing a friendly tutor. Tap the **up-arrow send button** (or press Enter on a computer) when you're ready. If the first reply comes back in English instead of your target language, simply remind it — *"Please reply only in Spanish from now on."* 💬 Example: *"Please reply only in simple Spanish sentences. I'm a complete beginner. After each of my messages, correct any mistakes in a kind bullet list, then ask me a follow-up question to keep the conversation going. Keep your replies short."* You'll know it worked when ChatGPT's first reply is in your chosen language, uses short sentences, and ends with a question (like *"¿Cómo te llamas?"*).

Step 3 — Send your first reply and let it correct you

Tap the message box at the bottom of the chat and type your answer to ChatGPT's question in your own words — even a rough guess is fine. Type slowly if it helps. After you tap send, ChatGPT will usually (a) gently correct any spelling or grammar slip, (b) give you a smooth model version of the sentence, and (c) ask you a new question on the same topic to keep things moving. This is the heart of practice: writing forces you to recall vocabulary, and seeing ChatGPT's correction turns the mistake into a real lesson. If it doesn't correct you, add to your original instructions from Step 2: *"Please always correct me, even on small mistakes."* 💬 Example: if ChatGPT asked *"¿De dónde eres?"*, type *"Yo soy de México"* — or even the imperfect *"Yo es de México"* — so ChatGPT has something to correct, which is the whole point. You'll know it worked when you see your own message followed by a short, kind correction and another question waiting for you at the bottom of the screen.

Step 4 — Turn on voice mode to practice speaking out loud

Speaking is the part most learners skip because it feels awkward — and it's the part that builds real fluency. In the ChatGPT app, look for a **microphone icon** 🎤 or a **headphones icon** near the message box, usually at the bottom-right. Tap it, and a "Listening..." indicator appears while the keyboard slides out of the way. If you don't see a microphone, open the side menu (look for a button like ☰ or your profile picture) and look for a setting called **"Voice mode"** or **"Voice chat"**. Speak one short sentence, pause for a moment, and let ChatGPT respond out loud. Voice mode is included on the free plan with daily usage limits; the more natural-sounding **Advanced Voice** (longer replies, smoother turn-taking) usually needs a **paid Plus plan**. If voice isn't available right now, read ChatGPT's written replies out loud instead — that still counts. 💬 Example: say out loud *"Buenos días, ¿cómo estás?"* and wait for ChatGPT to answer in voice. You'll know it worked when ChatGPT replies out loud in Spanish (or your chosen language) and you can keep talking without typing.

Step 5 — Drill the words you keep getting wrong

After a few exchanges, you'll start noticing patterns — the same word you keep misspelling, the same tense that always confuses you. Type a message in the chat box asking ChatGPT to focus on those exact trouble spots. This is where a chatbot beats a textbook: it can run the same drill over and over without getting tired, and it can build the drill around *your* mistakes, not a generic list. You can also ask it to slow down, give you more examples, or quiz you with little flashcards inside the same chat. 💬 Example: *"I keep mixing up 'ser' and 'estar.' Can you give me 5 more practice sentences and tell me which one to use, with a short explanation?"* You'll know it worked when ChatGPT returns a short drill, you answer it, and the explanations match your actual confusion rather than a textbook chapter.

Step 6 — Save your progress and find the chat again tomorrow

Before you close the chat, type one last message asking ChatGPT for a short summary of what you learned and what to review next time — this turns the conversation into a personal study note you can re-read. To find the chat again later, open ChatGPT and look at the **list of past conversations**. On a laptop browser, that list usually lives in a **left-hand column (sometimes called the "sidebar")**. On the phone app, tap the **menu icon ☰** at the top-left, or tap your **profile picture**, to open the same list — you'll see the chat title you chose in Step 1. Tap it and the conversation reopens exactly where you left off. Most apps also offer a **share or bookmark icon** (often a downward arrow or a little flag) if you'd like to save the summary into your phone's notes app. Consistency beats long sessions: five minutes a day for two weeks will take you further than one long hour once a month. 💬 Example: *"Please summarize today's lesson in 5 bullet points I can review before our next session, and list 3 things we should practice again."* You'll know it worked when you have a short, readable recap at the bottom of the chat and you can find the same conversation again in seconds.

Common mistakes

- **Skipping the setup instructions.** If you just type *"Let's practice Spanish"* with no detail, ChatGPT will speak like a textbook. Take two minutes in Step 2 to set the rules — that small effort changes everything.
- **Stopping at the first mistake.** It's tempting to give up when ChatGPT corrects you. The whole point is the retry. Read the correction, copy the right version out loud, and send it back in your own words.
- **Treating it like a quiz instead of a chat.** If you only ask "translate this word" drills, you practice vocabulary but not conversation. Mix in real little exchanges — ordering coffee, introducing yourself, asking for directions. That's where fluency actually grows.
- **Expecting ChatGPT to grade your accent.** ChatGPT can correct your grammar and word choice, but even the paid **Advanced Voice** mode (a Plus plan and above) won't reliably judge how you pronounce things. For real pronunciation feedback, record yourself with a voice-memo app and listen back, or practice with a human tutor.

What this means for you

- **In everyday life:** You can warm up before a real trip or a phone call with a relative who speaks another language — five minutes of "ordering coffee" role-play in ChatGPT is surprisingly effective at calming nerves.
- **For your business or work:** If you ever speak with international colleagues or clients, this is a private rehearsal space. Practice a short self-introduction, a polite disagreement, or small talk before the real meeting — no risk of freezing up on the spot.
- **If you're just getting started:** Pick *one* language you're curious about, set up the chat with very simple beginner rules, and commit to one short session a day for a week. That's all it takes to feel less awkward.

Try it now

Open ChatGPT right now, name a new chat *"Spanish practice — beginner"*, and send this exact message: *"Please reply only in simple Spanish. I am a complete beginner. Correct my mistakes gently after each message, then ask me one easy question. Start now."* Reply to its first question — even a guess counts — and you're practicing.

How to Build a 30-Day Plan to Learn Almost Any New Skill from Scratch

6 steps · Read online

> **In short:** Tell an AI assistant the skill you want, how much time you really have, and your starting point — then ask for a day-by-day 30-day plan. Review it, cut anything too ambitious, and you'll have a roadmap you can actually follow.

Learning a new skill can feel like staring at a mountain. The trick isn't motivation — it's a realistic day-by-day plan that fits your real life. In the next few minutes, you'll use an AI assistant (ChatGPT, Gemini, Claude, Copilot, Le Chat, or any other) as your personal planning partner, and walk away with a 30-day roadmap tailored to your free time. The AI can draft the plan, but it can't open the app for you on day 14 when you're tired — that's still your job.

Before you start

- An account on any AI assistant you like — ChatGPT, Gemini, Claude, Copilot, and Le Chat all work for this. A free plan is enough.
- A rough idea of the skill you want to learn, written as one specific sentence ("beginner conversational Spanish" beats "Spanish").
- An honest estimate of your real free time per day. Most people overestimate by 2–3× — write down the realistic number, not the hopeful one.
- About 15 minutes total.

Step 1 — Pick one skill and write it down in one sentence

Open a notes app, or grab a piece of paper. Write the skill you want to learn as one short, specific sentence. The more specific you are, the better your plan will be. "Get fit" is too vague; "be able to jog 3 km without stopping" is concrete enough to plan around.

You'll see the sentence sitting there in front of you. Read it back and ask: *Can I picture what success looks like in 30 days?* If not, tighten it. If your skill is broad (like "learn to code"), pick the smallest useful version of it (like "build a simple personal website in HTML and CSS").

If you change your mind halfway through, that's fine — but starting with one clear sentence keeps the AI focused instead of guessing what you want.

Example: "I want to learn beginner conversational Spanish so I can order food and have small talk on a trip to Mexico."

You'll know it worked when you can read your sentence out loud and immediately picture what 'done' looks like at day 30.

Step 2 — Write down your real time and your starting point

Before you open the AI, get honest about three numbers: how many minutes per day you can actually commit, how many days per week you'll likely skip (most weeks have at least one off day), and where you're starting from (complete beginner, or some experience).

You'll see a clear, simple list when you're done — something like "20 minutes a day, 5 days a week, total beginner." This step matters because AI assistants are optimists — they'll happily suggest a 2-hour daily plan that you'll quietly abandon by day 4.

If you're not sure how much time you have, look at your actual week and find the smallest realistic slot — early morning, lunch break, or after dinner. Twenty focused minutes beats two hours you'll never find.

Example: "20 minutes a day on weekdays, weekends flexible, complete beginner, no paid courses."

You'll know it worked when you could defend every number to a skeptical friend without flinching.

Step 3 — Open your AI assistant and send the planning prompt

Open the AI assistant you chose and look for the chat or message box — usually a large text area at the bottom or center of the screen. Paste a single message that includes your skill, your time, your level, and the structure you want.

After you press send, you'll see the AI start typing a response. Most assistants produce a table or day-by-day list within a few seconds. If you get a wall of generic advice instead, your prompt was too vague — add more specifics and try again.

If your tool has a "Projects," "Gems," or "Custom GPT" feature, you can save this prompt for later — but a regular chat is enough for now. Different assistants label the send button differently ("Send," "Ask," or a paper-plane icon), so look for the button at the bottom-right of the message box.

Example: "I want to learn beginner conversational Spanish for a Mexico trip in 30 days. I have 20 minutes a day on weekdays, weekends are flexible, and I'm a complete beginner. Please build me a day-by-day plan that: 1) fits in 20 minutes on weekdays and up to 60 minutes on weekends, 2) focuses on speaking and listening first, 3) uses free resources only, and 4) includes a small checkpoint at the end of each week. Format as a numbered list, one entry per day."

You'll know it worked when the response shows day 1, day 2, day 3... all the way to day 30, with a specific task for each.

Step 4 — Check the plan for realism and cut anything too ambitious

Read through the plan with a critical eye. Highlight anything that asks for more time than you said you had, anything that jumps ahead before you've built the basics, and any day that requires a paid resource you don't actually have. Your job here is to make the plan smaller, not bigger — AI tends to pack each day full.

You'll see your edited plan sitting next to the original. Most people find they need to cut 20–30% of what the AI suggested. That's not a failure of the tool; it's the whole point of this step.

If a section feels off (too advanced, wrong focus), tell the AI directly and ask it to redo just that part. For example: "Week 3 feels too hard — rewrite it for someone who only started last week."

💬 Example: "Days 18–22 ask for 90 minutes each. I only have 60 minutes on weekends. Please rewrite that week with shorter sessions and remove the paid app recommendation."

You'll know it worked when you can glance at any single day and think, "Yes, I can actually do that in the time I have."

Step 5 — Save the plan somewhere you'll see it every day

Copy the finished plan into a place you'll actually look at — a note on your phone, a printed page on the fridge, a calendar reminder, or a simple spreadsheet. The best plan in the world does nothing if you forget it exists by Tuesday.

You'll see the plan sitting where your future self can find it. If you use a calendar app, block the time slots now while you're thinking about it. If you prefer paper, print it — there's something about a physical page that makes people follow through.

If your AI assistant has a "Projects" or "Memory" feature, you can also save the conversation so you can ask follow-up questions later ("I'm stuck on Day 12 — help me rework it"). Otherwise, bookmark the chat in your browser.

💬 Example: "I'll paste the plan into my calendar as 20-minute events on weekdays, and a 60-minute block on Saturday morning."

You'll know it worked when, two days from now, you don't have to search for the plan — it's already in front of you.

Step 6 — Check in with the AI at the end of each week

After every 5–7 days, come back to the same chat (or start a new one with the same prompt) and tell the AI what actually happened. Be honest: what you skipped, what worked, what felt boring or too hard. Ask it to adjust the next week based on real life, not the original plan.

You'll see the assistant suggest a revised week — usually tightening what worked and dropping what didn't. This is the step that turns a static plan into a living one, and it's the difference between people who finish and people who quietly quit at day 9.

If you've fallen behind, don't restart from day 1 — ask the AI to compress the plan instead. Frame it as a shift forward, not a failure.

💬 Example: "I'm on day 8. I skipped day 5 and day 6 because of work. Days 1–4 went well, but the 30-minute listening exercise felt too long. Can you adjust the rest of the plan so I still finish in 30 days total, and shorten the listening parts to 15 minutes?"

You'll know it worked when each new week feels easier than the last — because it actually fits your life.

Common mistakes

- **Picking a skill that's too broad.** "Learn photography" gives you a useless plan. Tighten it to something specific, like "learn to take sharp, well-lit photos in natural light with my phone."
- **Lying about your free time.** If you tell the AI you have 90 minutes a day but you really have 20, you'll quit by day 5. Use the honest number from Step 2.
- **Treating the first plan as final.** The first response is a draft, not a contract. Always cut and adjust before you commit.
- **Forgetting to check in.** A plan you never revise becomes a plan you abandon. Set a weekly reminder to update it with the AI.

What this means for you

- **In everyday life:** You can use this exact method for cooking, drawing, a new language, an instrument, basic coding, or any hobby you've been putting off. The real skill here is the prompt, not the topic.
- **For your business or work:** The same prompt structure works for picking up a new tool at work — Excel pivot tables, a CRM, basic SQL, a design tool. You'll stand out by having a real plan instead of "I'll figure it out eventually."
- **If you're just getting started:** Pick one small skill you've actually wanted for at least a month — not a New Year's resolution from last week — and try this prompt tonight. Small, specific, and already wanted is the winning combination.

Try it now

Pick one skill, open your favorite AI assistant, and paste this: "I want to learn [SKILL] over 30 days. I have [X minutes] per day, [Y days] per week. I'm a [beginner / some experience]. Build me a day-by-day plan using free resources, with weekly checkpoints. Keep each daily task realistic for my time." Save the response, then cut one thing that feels too ambitious. You'll have a real plan in under ten minutes.

How to turn your messy study notes into practice quiz questions

6 steps · Read online

> **In short:** You can paste your messy notes into any AI assistant — ChatGPT, Gemini, Claude, Copilot, or similar — with a clear prompt, and get back flashcards, multiple-choice questions, and short-answer prompts in seconds. No special app or paid plan is required for the basic version.

Reading your own notes is one of the slowest ways to study. Quizzing yourself works much better — but writing questions by hand takes ages. The good news is that any modern AI assistant can turn your messy class or textbook notes into a ready-made practice set in well under a minute. This guide walks you through it step by step, even if you've never typed a single prompt before.

Before you start

- **An AI assistant account.** Any major one will do: ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or similar. The free tier is fine for this task.
- **A device with internet** — phone, tablet, or computer.
- **Your messy notes** in a form the AI can read. Typed notes from a notes app or a Google Doc work best. Photos of typed pages usually work too; photos of cramped handwriting often don't — type those out first.
- **About 10 minutes** for your first run. After that, it's a two-minute habit.

One honest note up front: AI assistants sometimes invent confident-sounding but wrong answers, especially with dates, numbers, and names. Treat the questions as a fast first draft you check, not a finished test. (We cover that in Step 4.)

Step 1 — Open a fresh chat in your AI assistant

Open the website or app of the AI assistant you picked (for example, chatgpt.com, gemini.google.com, or claude.ai). Look for a button that says **"New chat"** or **"New conversation"** — usually near the top-left of the screen or in a sidebar. Click or tap it. A blank chat box appears with a blinking cursor and a send button on the right.

If you don't see "New chat", look for a **pencil icon** ✏️, a **plus icon** ➕, or a **"+"** next to "Chats" — every assistant uses one of these for the same job. If a sign-in page pops up first, just enter your account details and continue.

💬 Example: no typing in this step — you're just opening a clean chat.

You'll know it worked when the screen shows an empty message box and any previous conversation is gone.

Step 2 — Paste your notes into the chat box

Click inside the empty message box at the bottom of the screen. Copy your notes from wherever they live — a notes app, a Google Doc, a Word file, or even an email to yourself — and paste them in. On a phone, tap and hold, then choose **Paste** from the small menu that appears. On a computer, the shortcut is **Ctrl+V** (Windows) or **Cmd+V** (Mac). The text will show up in the chat box, usually as one long block.

If your notes are very long (many thousands of words), paste them in chunks instead — one section at a time. Most assistants work best when each message stays reasonable in size. If your notes are on paper, type the key points into the chat box manually before moving on; don't try to paste a photo of your handwriting.

💬 Example (a short chunk you might paste): "Cell biology chapter 7. Mitochondria = powerhouse of the cell, makes ATP. Double membrane: outer is smooth, inner folded into cristae. ATP made via the Krebs cycle in the matrix and the electron transport chain on the cristae. Chloroplasts are found only in plant cells and do photosynthesis. Both mitochondria and chloroplasts have their own DNA — endosymbiosis theory."

You'll know it worked when your notes appear as readable text inside the chat box and the send button becomes active.

Step 3 — Ask for the quiz format you want

Below the pasted notes, still inside the same message, type a short request for the kind of quiz you want. The clearer you are, the better the output — don't just write "make a quiz", say what kind of questions and how many. Press the **send button** — usually an arrow icon ➤ on the right of the chat box, sometimes the Enter key on a keyboard — to submit.

If the assistant's first reply ignores your quiz request and does something else instead — like summarizing the notes — send a quick follow-up message restating what you want. Most assistants respond well to a clear, second nudge.

💬 Example (a prompt you can copy and adapt): "Turn the notes above into: (1) 5 multiple-choice questions with 4 options each, (2) 5 short-answer questions, (3) 5 flashcard-style Q&A pairs. Put the correct answer on a separate line for each one."

You'll know it worked when the assistant replies with a clearly numbered or labeled list of questions, each with options, an answer line, or a Q&A pair.

Step 4 — Quickly check the questions against your notes

Read the questions the assistant just produced. Keep your original notes open in another tab, a second window, or on paper beside you. Spot-check about five questions at random: does the answer actually match what your notes say? Pay extra attention to numbers, dates, names, and "who said what" — those are the spots where AI assistants most often slip up.

If a question is wrong, underline or circle it. If several are wrong, copy the bad ones back into the chat and ask the assistant to fix them. This is normal — think of the assistant as a fast first draft, and you as the editor.

Example (a quick fix prompt): "Question 3 says the answer is 1972, but my notes say 1927. Please correct it and double-check the rest for similar mistakes."

You'll know it worked when every question on your list has an answer that matches your source notes — or when you've flagged the ones that don't.

Step 5 — Save the questions somewhere you can find them

Highlight the assistant's reply with your cursor (or press and hold on a phone), then **copy** it. Open any notes app you already use — Apple Notes, Google Keep, Notion, a Word doc, even a plain text file on your desktop — and **paste** it in. Give the entry a clear title with the date and the topic, like *"Bio Chapter 7 practice — 12 March"*.

If your assistant offers a **share**, **download**, or **export** button — some show a paperclip icon, a download arrow, or three dots ⋯ — you can also use that to send the questions to your email or save them as a file. If you don't see that option, copy-paste works on every device.

Example (a paste title): "History Unit 4 practice questions — 12 March 2026"

You'll know it worked when you can close the AI tab and reopen your notes later to find the full question set, neatly stored under a recognizable title.

Step 6 — Use the questions to actually test yourself

Put your notes away. Open the saved quiz. For each multiple-choice question, cover the answer line with your hand and pick an answer out loud or in your head. For short-answer and flashcard questions, try to answer in one or two sentences **before** you peek. Mark the ones you got wrong — those are the spots to re-read tonight.

This last step is the whole point: the AI made the quiz so you don't have to, but the learning only happens when you wrestle with the questions yourself. Two passes through the same set usually beat one long review.

Example (a self-test routine): "10 questions, cover the answers, answer first, then check — repeat for any I missed."

You'll know it worked when you have a small list of weak spots to revisit and you feel more confident about the rest.

Common mistakes

- **Pasting a photo of your own handwriting and expecting it to work.** Most assistants can read typed text in images but stumble on cramped handwriting. Fix: type or dictate the key points first, then paste the typed version.
- **Writing a vague prompt like "make a quiz from this".** You'll usually get a few questions, but not the format you wanted. Fix: say exactly what you need — multiple-choice, short-answer, or flashcards — and how many of each.
- **Trusting every answer without checking.** AI assistants sometimes invent plausible-sounding facts that aren't in your notes at all. Fix: keep your notes open and spot-check five questions per set, especially any with dates, numbers, or proper names.

What this means for you

- **In everyday life:** if you're a student, this turns an hour of note review into a 10-minute habit you can run the night before any quiz — and the questions are reusable the next day.
- **For your work:** if you're learning a new skill, certification, or tool for your job, the same trick works on training videos, manuals, and slide decks. Paste the highlights, get a quiz, test yourself.
- **If you're just getting started:** the easiest first move is to take the single page of notes you already have open right now, paste in one paragraph, and ask for three multiple-choice questions. That's it — the rest of the habit builds from there.

Try it now

Open a fresh chat in your AI assistant, paste the last page of notes you studied, and ask for **three multiple-choice questions with the answer on a separate line**. Ninety seconds from now you'll have a tiny quiz you can test yourself on right away.

Boost Your Learning: Get Quick Explanations and Study Guides with Copilot

6 steps · Read online

By the end of this guide, you'll be able to turn a confusing topic into a clear summary or a ready-to-use study guide in under five minutes — even if you've never used an AI assistant before. This is for students, curious learners, and anyone who wants a friendlier way to tackle tricky subjects.

Before you start

- **A free Microsoft account** (the same one you'd use for Outlook or OneDrive). If you don't have one, you can create one at Microsoft's sign-up page — it only takes a minute.
- **A device with a browser** (computer, tablet, or phone). You can also use the Copilot mobile app if you prefer — search for "Microsoft Copilot" in your app store.
- **The free plan works fine** for this guide. Some advanced features (longer conversations, image generation) may need a paid Microsoft 365 Copilot subscription, but the core explanation and study guide features are available on the free tier.
- **Rough total time:** about 5 minutes for your first summary, and around 10 minutes to build a fuller study guide.

> One honest note: Copilot can explain and summarize, but it can't replace your own reading, lectures, or assessments. Think of it as a study buddy that helps you understand — not a shortcut that does the learning for you.

Step 1 — Open Copilot in your browser

Go to **copilot.microsoft.com** in any modern browser such as Chrome, Edge, Firefox, or Safari. You'll land on a clean page with a big text box in the middle and a friendly greeting at the top. Below the box you'll usually see a few suggestion chips — these are just quick prompts to show you what's possible. **What you see:** a chat-style page with a message input field, a send button (often a paper airplane icon), and options to attach images or turn the microphone on. **What happens next:** the page is ready for you to type your first message. **If it looks different:** some people see Copilot inside the Bing search page instead — that's fine, the chat box works the same way. On a phone, you can also download the free Microsoft Copilot app from your app store; the layout is similar. **You'll know it worked when** you can see a blinking cursor inside the message box, ready for you to type. Example: type `Hi Copilot, can you help me understand a tricky topic?` and press Enter — you'll see a short, friendly reply appear in seconds.

Step 2 — Ask for a plain-language explanation of a topic

Click into the message box and type your topic, then ask Copilot to explain it like you're a beginner. The trick is to name the subject clearly and add the words **"explain simply"** or **"in simple terms"** so the answer comes back friendly rather than textbook-heavy. **What you see:** your message appears in the box, and a short "thinking" animation runs while Copilot drafts a reply. **What happens next:** a few paragraphs of clear explanation appear below, usually with a heading at the top. **If it looks different:** some browsers show Copilot inside a side panel — you can still type into it the same way. If nothing seems to happen, check your internet connection or refresh the page. **You'll know it worked when** you see a written explanation in plain English (no confusing jargon), with key terms spelled out. Example: `Explain how photosynthesis works in simple terms, like I'm 12 years old.`

Step 3 — Ask follow-up questions to go deeper

Read the explanation, then ask a follow-up question in the same chat to clear up anything that's still fuzzy. You don't need to start a new conversation — Copilot remembers what you've already asked within your current session. This is how you turn a single answer into a mini-lesson. **What you see:** the chat history scrolls up, and a new empty message box sits at the bottom. **What happens next:** Copilot answers your new question while keeping the earlier explanation in mind, so the reply feels connected. **If it looks different:** if the chat forgets what you said earlier (it can happen after a long break or new session), just paste a short summary of the topic back in: "Earlier we were discussing photosynthesis. Now..." **You'll know it worked when** the new answer clearly builds on the previous one, instead of repeating basic facts from scratch. Example: `Can you give me a real-world example of photosynthesis that I see in everyday life?`

Step 4 — Turn the explanation into a study guide

Now ask Copilot to package what you've learned into a study guide you can review later. Be specific about the format you want — bullet points, questions, a short quiz — so the result is genuinely useful and not just another wall of text. **What you see:** a longer reply appears, usually with headings like "Key Points" or "Practice Questions". **What happens next:** the study guide is ready to read on screen right away. **If it looks different:** if the format isn't quite right, simply ask again: "Can you redo that as a numbered list?" or "Make it shorter." **You'll know it worked when** you have a tidy, scannable summary with headings, bullet points, or questions you can actually revise from. Example: `Great, now turn this into a short study guide with 5 key points and 5 practice questions I can test myself with.`

Step 5 — Copy your study guide somewhere safe

Click and drag to highlight the study guide text, then copy it with **Ctrl + C** on a computer (or **Cmd + C** on a Mac), or long-press the text on a phone and choose "Copy". Paste it into a notes app, a Word document, a Google Doc, or even an email to yourself. This step matters because fresh chats don't carry old answers over. **What you see:** the text becomes highlighted in blue, and after copying you can paste it anywhere you like. **What happens next:** the study guide is saved outside of Copilot, ready for you to read on the bus, print out, or share with a study buddy. **If it looks different:** if copying doesn't work, look for a small "Copy" icon that sometimes appears in the corner of Copilot's replies. **You'll know it worked when** you can open your notes app and see your study guide sitting there, ready to read. 💬 Example: paste it into a note titled `Photosynthesis — quick study guide` so you can find it easily next time.

Step 6 — (Optional) Ask for a one-line memory hook

Finish by asking for a single sentence or short rhyme you can repeat in your head — these tiny memory hooks (sometimes called mnemonics) make recall much easier during a test or conversation. It's a small step, but learners swear by it. **What you see:** a short, snappy reply appears at the bottom of the chat. **What happens next:** you can write that line on a sticky note, a flashcard, or the top of your study guide. **If it looks different:** if Copilot gives a long answer, reply with "Shorter please — just one sentence." **You'll know it worked when** you have a memorable phrase or sentence you'd feel confident saying out loud. 💬 Example: `Can you give me one short sentence I can remember to recall how photosynthesis works?`

Common mistakes

- **Mistake:** Asking a vague question like "Tell me about history." **Fix:** Add the specific topic, time period, or angle you care about — for example, "Explain the main causes of World War One in simple terms."
- **Mistake:** Treating the first answer as the final word. **Fix:** Always ask at least one follow-up question. The second and third answers are usually where the real understanding clicks.
- **Mistake:** Forgetting to save the study guide. **Fix:** Copilot's chats may not carry over between sessions on the free plan, so copy anything important into a notes app straight away, before you close the tab.

What this means for you

- **In everyday life:** Next time you read a confusing news article or hear a tricky concept at a dinner party, open Copilot and ask for a plain-English explanation. Five minutes later you'll actually understand what's going on.
- **For your studies or work:** You can build a personalized revision guide for almost any subject — history, science, software, even a new work process — without paying for a tutor.
- **If you're just getting started:** The easiest first move is to pick one topic you've always wondered about, open Copilot, and type: "Explain [topic] in simple terms." That's it — your learning habit starts there.

Try it now

Open **copilot.microsoft.com** right now, type `Explain how the immune system fights a cold in simple terms`, and read the reply. That's your first study guide — and it took less than a minute.

How to use Qwen Chat to create personalized study flashcards for any subject

6 steps · Read online

By the end of this guide, you'll be able to paste in your study notes, ask Qwen (Alibaba's AI assistant) to turn them into question-and-answer flashcards, and then copy them into whatever quiz tool you already use — so revising takes a fraction of the usual time. This is for anyone studying for an exam, picking up a new skill, or just trying to remember things better.

Before you start

- **An account at chat.qwen.ai.** You can sign in with a Google account, a GitHub account, or an email address — pick whichever you already use. A free plan is available and is enough for this task.
- **A device with a browser** (laptop, tablet or phone). Qwen Chat is web-based, so no install is required.
- **Some study material to work from.** A few paragraphs of notes, a chapter summary, or even a short topic description is plenty. The more specific your notes, the better the cards.
- **Roughly 10–15 minutes.** Most of that time is reading and picking the cards you want to keep.

Step 1 — Open Qwen Chat and start a fresh conversation

Open your browser, go to **chat.qwen.ai**, and sign in if you aren't already. Once you're on the main chat screen, look for a button like **"New Chat"** (often a small + icon or a pencil symbol near the top or left side of the page). Click it. A blank conversation box opens, ready for your message. Starting a new chat each time keeps earlier topics from confusing the AI, which is important because a long, mixed conversation tends to make answers less focused. **If it looks different:** if you can't find a "New Chat" button, just type your request into whatever text box is visible — Qwen will treat it as the start of a new thread on its end. **You'll know it worked when** the screen shows an empty message field with no previous conversation visible.

Step 2 — Tell Qwen what subject you're studying and paste in your notes

Type a short intro that says (a) the topic, (b) roughly what level you're at, and (c) the notes you want turned into cards. Then press **Enter** or tap the send arrow. Qwen will read your message and start processing it. This first message is where you set the "frame" — the more clearly you describe the subject, the more on-target the flashcards will be. **If it looks different:** if your notes are very long, consider pasting just one section at a time so Qwen doesn't have to skim anything important. **You'll know it worked when** Qwen replies by acknowledging the topic and confirming what it's about to do.

Example: *"Hi Qwen — I'm studying for a Year 9 biology exam on cell structure. Here's my notes: [paste your notes here, or describe: 'The main topics are animal cells vs plant cells, the function of the nucleus, mitochondria, ribosomes, cell membrane, and how osmosis works.' Please make 12 flashcards covering these."*

Step 3 — Ask Qwen to format each flashcard clearly

In the same conversation, send a second message asking Qwen to lay out every flashcard in a consistent format — usually **"Question: … / Answer: …"** on separate lines, with a number in front of each one. Press send and wait a few seconds. Qwen will re-format its previous answers into a clean list you can easily copy. A consistent layout saves you time later because you can paste the whole block into a spreadsheet or a flashcard app without re-typing anything. **If it looks different:** if Qwen gives you a paragraph instead of a list, just reply with *"Please list them numbered, with Question and Answer on separate lines."* — it will happily redo them. **You'll know it worked when** the response is a numbered list where each item has a clear "Question" line and an "Answer" line.

Example: *"Great, please rewrite those as a numbered list, with each one formatted exactly like this: 'Card 1 — Question: [question] | Answer: [answer]'. Use a vertical bar to separate Q and A so I can copy them straight into a spreadsheet."*

Step 4 — Ask for more cards, harder cards, or a quick quiz check

Now send another short message asking Qwen to either (a) generate **more cards** on the same topic, (b) make the existing cards **harder** (for example, "give me application-level questions, not just definitions"), or (c) **quiz you** by hiding the answers and asking the questions one at a time. This is where Qwen really shines as a study partner — you can keep refining the same set until it matches your exam style. **If it looks different:** the "quiz me" option is just a chat request; there's no special button. **You'll know it worked when** Qwen either produces a longer list, asks you a question, or offers a "harder" version of the same cards.

Example: *"Please add 6 more cards, but make them harder — focus on application and 'why' questions rather than simple definitions. Then quiz me on the first 5 cards one at a time."*

Step 5 — Copy the flashcard list out of Qwen

Scroll back up to the formatted list and **select the text** with your mouse, then copy it (**Ctrl+C** on Windows/Linux, **Cmd+C** on Mac, or long-press on mobile). Open the place you want to study from — Google Sheets, Notion, a notes app, or a flashcard app like Anki or Quizlet — and paste it in. Qwen doesn't store your cards in a built-in flashcard deck, so this copy step is the bridge between the AI's output and your study routine. **If it looks different:** on mobile, you may need to tap-and-hold inside Qwen's response to bring up a "Copy" or "Select all" option. **You'll know it worked when** the flashcards appear as editable text in your chosen app, ready to import or tweak.

Step 6 — Test yourself, then come back and refine weak cards

Spend a few minutes quizzing yourself with the cards you just saved. Notice which ones you got wrong, felt fuzzy on, or found too easy. Then **return to your Qwen chat** and ask it to focus on exactly those weak spots — for example, *"Give me 5 more cards that drill the difference between osmosis and diffusion, because I keep mixing them up."* This feedback loop is what makes AI-generated flashcards genuinely useful: the set evolves with your understanding, instead of staying static like a textbook chapter. **If it looks different:** if Qwen has lost the thread of your earlier messages, just paste in your original topic description again as a refresher. **You'll know it worked when** you can identify at least 3 cards that you now answer confidently, and a short list of topics you still want to drill.

💬 Example: *"I got these wrong: [list 2-3 topics]. Please make 4 more cards on those specific points, and include a tiny hint on the back of each one."*

Common mistakes

- **Pasting a 50-page textbook and expecting perfect cards.** Qwen works best when you give it focused chunks of notes or a clear topic summary. **Fix:** paste one section at a time, or describe the key concepts in your own words first.
- **Skipping the "format" step.** If you don't ask for a numbered "Question / Answer" layout, you'll get a paragraph that's painful to copy into anything. **Fix:** always send the second message asking for the exact format you want.
- **Treating AI cards as automatically correct.** Qwen can occasionally get a fact subtly wrong, especially for niche subjects. **Fix:** spot-check the first 5 cards against your actual textbook or syllabus before you commit to the whole set.

What this means for you

- **In everyday life:** any time you need to remember a list — countries and capitals, foreign vocabulary, recipe ratios — you can spin up a quick flashcard set in minutes instead of handwriting them.
- **For your study or exams:** you get a personalized revision deck that adapts to *your* notes, not a generic one from the internet, and you can ask for harder questions as the exam gets closer.
- **If you're just getting started:** don't overthink it — pick one small topic, run through the six steps above once, and you'll have the hang of it by the second subject.

Try it now

Open **chat.qwen.ai**, start a new chat, and paste in this exact message to make your first set of cards right now:

> *"Please create 8 flashcards covering the basics of [your subject — e.g. 'the solar system', 'French food vocabulary', 'intro to spreadsheets']. Format each one as: 'Card N — Question: ... | Answer: ...'. Keep answers under 20 words."*

Two minutes from now you'll have a deck you can copy into your favorite notes app and start quizzing yourself tonight.

Turn a Reading into a Study Quiz with GitHub Copilot

6 steps · Read online

> **In short:** You can use GitHub Copilot's chat to turn any short reading into a list of practice questions, then ask for the answers and verify each one yourself against the original text. Copilot is built for coding, so this works best with shorter readings (about one page) and a free or paid GitHub account with Copilot access.

GitHub Copilot is best known as a coding assistant — picture a very fast pair-programmer who sits beside you and suggests the next line. But the same chat panel also reads, summarizes, and writes questions about plain text. That makes it handy for students, book clubs, or anyone who wants to test whether they really understood an article. The two habits that make this work: keep the reading short, and double-check every answer before you trust the quiz. One honest limit up front — Copilot can draft questions about a text, but it cannot grade open-ended essay responses or replace doing the reading yourself.

Before you start

- A free GitHub account at github.com. Sign up if you don't have one — the same account works for the website.
- A web browser (Chrome, Edge, Firefox, or Safari).
- A reading you want to study. Pick something between one paragraph and roughly one printed page. Longer texts work, but smaller chunks give cleaner quizzes.
- Copilot Chat access. It is available on a free tier with daily message limits; heavier use may require a paid Copilot plan — see the current pricing page for the latest details.
- About 10 minutes.

Step 1 — Open Copilot Chat on github.com

Sign in to github.com, then look at the top-right corner of the page for a small chat-bubble icon or a "Copilot" label. If you don't see it there, click your profile picture in the same corner and scan the menu that drops down — Copilot is usually listed near "Settings" or "Your repositories." Clicking it opens a chat panel, normally as a side window on the right or as a pop-up near the bottom of the screen. You'll see a blank text box at the bottom labeled something like "Ask Copilot" or "Message Copilot."

Example: don't type anything yet — just confirm the chat is open and you can click into the box.

You'll know it worked when the chat panel is visible, the text box is empty, and your cursor blinks inside it.

Step 2 — Paste your reading into the chat

Open the reading you want to study — a news article, a chapter section, a Wikipedia entry, anything text-based — and copy the whole passage to your clipboard. On Windows this is usually Ctrl+C; on a Mac, Cmd+C. Click into the Copilot text box, then paste with Ctrl+V or Cmd+V. The text appears inside the box, often wrapping over several lines. If the reading is very long, paste only one section at a time — Copilot handles shorter chunks much better and gives you tighter questions. If nothing pastes, right-click inside the box and choose Paste, or use the browser's Edit menu.

Example: paste this short paragraph — *"The Amazon rainforest produces about six percent of the world's oxygen and stores billions of tons of carbon. It also hosts ten percent of all known species. Fires there release stored carbon and reduce the forest's ability to absorb more."*

You'll know it worked when the full reading sits inside the chat box and you can still see a blinking cursor at the end.

Step 3 — Ask Copilot to make a quiz

Now you ask for what you want. In plain English, type a short instruction that tells Copilot the question type, the count, and the difficulty mix. Vague asks produce vague quizzes, so be specific. Press Enter or click the send icon — usually a paper-airplane shape or an arrow inside the text box — to submit. Copilot usually replies within a few seconds with a numbered list of questions, each followed by four options labeled A, B, C, and D. If the reply looks too long or off-topic, click "Stop" if you see it, then send a clearer prompt.

Example: *"Read the passage I pasted. Write 5 multiple-choice questions about it, with 4 answer choices each. Make 2 easy recall and 3 about meaning or application. Don't give me the answer key yet."*

You'll know it worked when Copilot responds with a numbered list of questions, each with four lettered options.

Step 4 — Ask for the answer key

In the same chat, send a follow-up message asking for the answers. Tell Copilot to put each correct letter on its own line and to quote a short supporting phrase from the original text. This second pass is where verification matters — Copilot can guess, and you'll check it in the next step. If the chat shows a "regenerate" or "try again" button, ignore it for now and just send your own text request instead.

Example: *"Now give me the answer key. For each question, write the correct letter (A/B/C/D) and quote one short phrase from the original passage that supports it. Don't paraphrase — quote directly."*

You'll know it worked when Copilot replies with a clean list — for example, "1. B — 'six percent of the world's oxygen'" — one entry per question.

Step 5 — Check every answer against the original reading

This is the step most people skip, and it's the one that turns a fun experiment into a real study tool. Open the original reading again in another browser tab or window. For each answer Copilot gave, find the exact phrase it quoted and read the surrounding sentence to confirm the answer is actually right and the quote really supports it. If a quote is missing, doesn't match the text, or the answer feels off, delete that question or rewrite it yourself. Treat every answer as a draft until you've seen the supporting phrase with your own eyes.

Example: if the answer says "B — 'ten percent of all species'" but the original paragraph says the rainforest *hosts* (not produces) ten percent of species, you've caught a subtle mistake — correct the question before using it.

You'll know it worked when every answer on your list is backed by a real phrase in the original text and every question still makes sense to a fresh reader.

Step 6 — Save or share the quiz

Once the quiz is verified, copy the questions and answers into wherever you actually study — a notes app, a document, or a flashcard tool. If you want to share with a class or study group, paste just the questions into a shared doc and keep the answers hidden (for example, on a second page, in a separate file, or behind a spoiler tag). In Copilot Chat, you may see a "copy" or "share" icon near the reply; if not, just drag your mouse across the text and copy it out the usual way. Then close the chat — your verified quiz now lives outside it.

Example: paste your verified list into a note titled "Quiz — Amazon rainforest (WIP)" so you can come back and revise it later.

You'll know it worked when the finished quiz lives somewhere outside the chat and you can re-open it without Copilot.

Common mistakes

- **Pasting a whole book chapter at once.** Copilot loses track of long inputs and the questions become vague or invented. Stick to one section at a time — a page or less — and run a new chat for each chunk.
- **Trusting the answers without checking.** Copilot sometimes invents a quote that sounds right but isn't in the text. Treat every answer as a draft until you've actually found the supporting phrase in the original.
- **Asking for too many questions at once.** Asking for 20 questions usually gives you weaker questions than asking for 5. Smaller, sharper sets are easier to verify and easier to actually use.

What this means for you

- **If you're studying on your own:** turn each article you read for class into a 5-question quiz the same day. Doing this within an hour of reading cements the material much better than re-reading alone.
- **If you teach or run training:** build a quick quiz at the end of any handout. It takes about ten minutes per reading, and the answers stay yours to verify before sharing.
- **If you just want to remember what you read:** even a 3-question quiz, done once, beats highlighting the same paragraph three times.

Try it now

Pick one short article you saved recently — anything between 200 and 500 words. Open github.com, start a Copilot chat, paste the article, and ask for 4 multiple-choice questions. Then ask for the answer key. Spend two minutes checking each answer against the original. You'll have a verified quiz before your coffee gets cold.

How to Create a Simple Weekly Study Plan with Any AI Assistant

6 steps · Read online

> **In short:** You can turn a messy list of classes, deadlines, and life commitments into a realistic day-by-day plan in about 15 minutes. Paste your real workload into any AI assistant and ask it to build a balanced week, then tweak it as you go.

You don't need a fancy app or a productivity course. With about fifteen minutes and any AI assistant — ChatGPT, Gemini, Claude, Microsoft Copilot, Le Chat, or DeepSeek — you can hand it your real week and get back a draft timetable you can actually follow.

Before you start

- An account with any major AI assistant. **Free plans work fine** for this; you only need a paid plan if you want the longest chats or image features, neither of which matter here.
- A rough list of what's on your plate this week: classes, labs, due dates, exams, shifts, family stuff. A messy notes-app dump is fine — neatness doesn't matter yet.
- About **15–20 minutes** the first time. After that, around 5 minutes per week to refresh the plan.

Step 1 — Gather your real workload in one place

Open your phone's Notes app, a blank document, or a piece of paper. Spend five minutes dumping everything that touches your time this week: every class, every due date, every shift, every appointment, plus any "I should really find time for…" thing (laundry, gym, calling home). Don't sort it yet — just get it out of your head. If you have a syllabus PDF or a screenshot of your course page, keep it handy, but plain text works too. You'll know it worked when you have a single messy list of roughly 10 to 30 items and a rough sense of which ones are urgent.

Step 2 — Open a fresh chat in your AI assistant

Open any major AI assistant — ChatGPT, Gemini, Claude, Microsoft Copilot, Mistral's Le Chat, or DeepSeek — in your browser or app. Start a **brand-new conversation** so the chat starts empty; an old thread carries context that may confuse the planner, and you want a clean slate. Click once into the message box so the cursor blinks. If your app drops you into an old thread by default, look for a **"New chat"** or + button — usually top-left, sometimes labeled "New conversation". You'll know it worked when you see an empty input box with no previous messages above it.

Example: no typing here yet — just confirm the empty chat is open and ready.

Step 3 — Paste the workload and ask for a balanced week

Click into that empty box and paste your messy list, then add the request on a new line. A **prompt** is just the instruction you give to an AI — think of it as telling a friend what's going on. The exact wording matters less than being specific about your time window and your real constraints. After you press send, the assistant will start typing out a draft schedule, usually in under a minute. You'll know it worked when you see a day-by-day or block-by-block grid appear, with each of your tasks placed into a morning, afternoon or evening slot. If you get generic study tips instead of an actual plan, your prompt was too vague — tighten it and try again.

Example:

Here's everything I have to do this week (Monday to Sunday). Some classes are fixed times, the rest is flexible. Build me a realistic day-by-day plan with blocks for studying, breaks, meals, sleep by 11pm, and one hour of exercise. Flag anything you couldn't fit. Mon 9am Biology 101 lecture (fixed); Tue 6pm cafe shift until 10pm (fixed); Wed Psychology essay due 11:59pm; Thu 7pm soccer practice; Fri nothing fixed; Sat family lunch 1-3pm; Sun Statistics problem set due.

Step 4 — Ask the assistant to sanity-check its own plan

This is the step most people skip, and it's the difference between a plan that works and one you abandon by Tuesday. After the draft appears, send a short follow-up in the same chat and ask the AI to be its own critic. Expect a reply that points out overloaded days, missing buffers, or tasks you forgot. You'll know it worked when you get back 2 to 5 honest "this won't work because…" notes, plus a revised v2. If the assistant just says "looks great!", press harder — tell it you're a real person with limited energy and ask it to mark the top three risk points. Honest pushback is more useful than confident rubbish.

Example:

Now read that plan back like a strict friend. Tell me which days are unrealistic, where I forgot breaks, and what you would cut first if I only had four hours of real focus time per day. Give me a v2.

Step 5 — Copy the plan somewhere you'll actually open

AI chats aren't where plans live — they're where plans are **born**. Pick one place you'll check daily: Google Calendar, Apple Calendar, Outlook, a paper wall planner, or a to-do app like Todoist or TickTick. Manually transfer the blocks, or simply copy the schedule text across. You'll know it worked when the plan shows up on the screen you actually look at first thing in the morning. If your phone calendar feels like too much friction, the back of an envelope on the fridge works too — the point is that the plan leaves the chat and enters your real life.

Step 6 — Book a 10-minute Sunday reset for next week

Set a recurring 10-minute slot on Sunday evening called "Plan next week". When the reminder pops up, open a new chat, paste last week's plan plus any new assignments, and ask for a refreshed draft. Repeating this turns the first awkward 15-minute session into a 5-minute habit. You'll know it worked when Sunday planning feels boring rather than stressful. If you skip a week, just restart — there's no streak to protect, only a calmer Monday.

Common mistakes

- **Asking for "a study plan" with no detail.** The AI will give you generic advice because it has nothing specific to work with. **Fix:** always paste your actual list of classes, deadlines, and fixed commitments before asking for the plan.
- **Treating the first draft as final.** AI drafts are optimistic — they ignore tiredness, hunger, and the time it takes to actually start a task. They also can't send you reminders or open your calendar for you. **Fix:** always run Step 4 and ask the assistant to mark overload honestly.
- **Leaving the plan inside the chat.** If you don't transfer it, you won't follow it. **Fix:** pick one external home — calendar, paper, or a to-do app — and commit to reading it from there every morning.

What this means for you

- **In everyday life:** A clear weekly plan means fewer late-night panic sessions and more energy for the things you actually enjoy, because you're not constantly re-deciding what to do next.
- **For your business or work:** Even if your job is part-time or a side gig, the same habit works — feed in shifts, deadlines, and fixed commitments, and you'll stop double-booking yourself and your boss.
- **If you're just getting started:** Don't try to optimize anything else this week. Just do Steps 1 and 3 with the 💬 example prompt, paste in your real list, and see what comes back.

Try it now

Open your Notes app, write down every class, deadline, shift, and appointment you have in the next seven days (yes, even the small ones), then paste that list — exactly as messy as it is — into any AI assistant with the prompt from Step 3. You'll have a draft plan in under a minute.

Turn Your Messy Class Notes into Study Flashcards with Cohere

6 steps · Read online

> **In short:** Cohere's chat can turn a pasted block of lecture notes into a Q&A flashcard set in under a minute. You'll need a free Cohere account, your notes in plain text, and a few minutes to copy and study.

Lecture notes are useful, but they don't test you. Flashcards do — they make you recall the answer before you see it. The painful part is making them. This guide walks you through using Cohere's free chat to skip the typing and get straight to quizzing yourself, so an evening of studying actually sticks.

Before you start

- A free Cohere account (sign up with an email at cohere.com — no credit card needed)
- Your notes in plain text — copy them from your laptop, Google Doc, or note app
- About 10 minutes total; the actual chat takes under a minute
- A notes app, document, or flashcard tool to paste the finished cards into

Step 1 — Create a free Cohere account and open the chat

Go to **cohere.com** in your browser and click the **Sign up** or **Get started** button (look for it near the top-right of the page). Enter your email, set a password, and confirm the verification email. After you confirm, you land inside the Cohere dashboard. Look for a chat box or a labeled chat/playground area — it's usually front and center. If you see a row of model options above the chat, leave the default one selected; any of them can turn notes into Q&A pairs. Example: a one-line email like "yourname@example.com" works fine to create the account. You'll know it worked when you see a chat box with a text field at the bottom and a blinking cursor ready for your message.

Step 2 — Open your notes and copy a small section

Open your class notes in a separate window or tab — Google Docs, Word, Notion, or a plain text file all work. Find one topic or sub-section, such as "The Three Branches of the US Government" or "Mitosis — phases and outcome." Highlight just that section and copy it (Ctrl+C on Windows, Cmd+C on Mac). Use a chunk of about 200–500 words at a time; longer chunks mean the AI may skip details or summarize too loosely. Example: copy only the bullet points about "phases of mitosis," not the whole biology chapter. You'll know it worked when the notes are in your clipboard and you can paste them anywhere with Ctrl+V or Cmd+V.

Step 3 — Paste your notes into the chat and ask for flashcards

Click inside the Cohere chat box and paste your notes (Ctrl+V or Cmd+V). On the next line, add a short instruction telling the AI what you want. A simple, clear instruction works best — say how many cards you want, the format, and whether you want answers hidden on the back. Hit **Enter** or click the **Send** button (it usually looks like a paper airplane or arrow). Example: "Turn these notes into 10 study flashcards. Format each one as: Q: [question] / A: [answer]. Cover the main definitions, dates, and examples." You'll know it worked when the chat starts typing a response and a numbered or bulleted list of Q&A pairs appears on screen.

Step 4 — Check the cards and ask for fixes if needed

Give the response a quick read. Each flashcard should have one clear question, one short answer, and no facts that weren't in your notes. If cards are too long, too vague, or skip details, ask for a redo. Be specific about what to change. Cohere doesn't browse the web in chat, so it should only use what you pasted — but always double-check before you trust a card. Example: "The answer to card 3 is too long. Make every answer under 15 words, and remove any dates I didn't include in the notes." You'll know it worked when every card has a clean question, a short answer you can verify in your notes, and nothing has been added from outside.

Step 5 — Copy the flashcards into your study app or a document

Highlight the entire flashcard block in Cohere's reply (click at the start, then Shift+click at the end, or press Ctrl+A / Cmd+A inside the chat panel if the chat allows it). Copy it, then paste it into a doc, a notes app, or a flashcard tool like Anki or Quizlet. Keep the Q: / A: format so each card is easy to import. If you want a printable set, paste into Google Docs and use a table. Example: paste into a new Google Doc titled "Bio 101 — Mitosis flashcards" so you can find it again later. You'll know it worked when the cards are saved in a place you can open on your phone or laptop, and the question and answer for each card are clearly separated.

Step 6 — Study the cards the same day

Open your saved flashcards and quiz yourself: read the question, say the answer out loud or in your head, then flip to check. Sort cards into "got it" and "missed it" piles. Re-quiz the missed ones tomorrow. Flashcards only work if you come back to them — making them with AI saves time, but the learning still happens when you test yourself. Example: spend 5 minutes quizzing tonight, then 5 more tomorrow morning on the cards you missed. You'll know it worked when you can answer at least 8 out of 10 cards from memory without looking.

Common mistakes

- **Pasting a whole textbook chapter at once.** The AI will skip details or lose focus. **Fix:** paste one sub-topic at a time, around 200–500 words per request.
- **Skipping the fact-check step.** Cohere can rephrase or summarize your notes, but it won't catch every typo in your source — and the AI is summarizing, not generating from scratch, so any ambiguity in your notes can show up in the cards. **Fix:** spend a minute comparing each card to your original notes before you save.
- **Saving the cards in the chat only.** If you close the chat or the window, the cards can be hard to find again. **Fix:** copy the finished block into a doc, notes app, or flashcard tool right away.

What this means for you

- **In everyday life:** if you're studying for a test, a work certification, or even a new language, the same trick works — paste text, ask for Q&A cards, save them, quiz yourself.
- **For your business or work:** if you're onboarding to a new tool or learning a new process, you can paste documentation into Cohere and get self-quiz cards to learn it faster.
- **If you're just getting started:** the easiest first move is to take one page of notes you already have and try this on a single topic tonight.

Try it now

Open your most recent class notes, copy one sub-topic about 200 words long, paste it into Cohere's chat, and ask for "10 flashcards in Q: / A: format." In two minutes you'll have a small set of cards ready to study — and a system you can reuse for every subject from now on.

AI Made Simple: A Gentle First Guide for Beginners

5 steps · Read online

If you've ever felt a bit daunted by new technology, this guide is for you. You'll learn how to confidently use artificial intelligence (AI) to help you with daily tasks, spark creative ideas, and even assist with writing – all while understanding how to keep your personal information safe.

Before you start

- **An internet-connected device:** A computer or tablet offers the easiest experience for typing and seeing the responses clearly. Smartphones work too, but the screen might feel a bit small.
- **A free AI account:** You will need to sign up for a free AI tool like ChatGPT (from OpenAI) or Gemini (from Google) using your email address. This guide focuses on free versions.
- **Stable internet connection:** Ensure you have a reliable internet connection to access the AI.
- **Time:** Allow about 5-10 minutes to follow this guide and complete your first interaction.

Step 1 — Open your AI chat and sign in

Your first step is to open your chosen AI program. Think of these as helpful digital assistants that you can text back and forth with, just like sending a message to a friend. These programs are often called **chatbots** (which is just a fancy name for a computer program designed to have conversations with people).

Open your web browser (like Safari, Chrome, or Edge) and go to the website for ChatGPT or Gemini. Look for a "Log in" or "Sign up" button. If you've never used it before, choose "Sign up" and follow the simple on-screen prompts to create a free account using your email address.

What you see: You'll typically see a clean login page asking for your email and password, or options to sign in with your Google or Apple account. **What happens next:** After successfully logging in or signing up, you'll be taken to the AI's main chat interface, which usually has a large, empty text box at the bottom of the screen. **If it looks different:** If you see a "Download App" option, that's fine too! The steps in the app will be very similar. If you're struggling to find the sign-in, look for words like "Get Started" or "Join for Free". **You'll know it worked when:** You see a clear, simple screen with a large empty box at the bottom, ready for you to type your message.

Example: Type **chatgpt.com** or **gemini.google.com** into your browser's address bar. Then, click "Sign up" and follow the prompts to create your free account.

Step 2 — Find the chat box and say hello

Once you're logged in, you'll be on the main chat screen. This is where you'll interact with the AI. It's usually a very straightforward screen with most of the space dedicated to the conversation area, and a typing box at the bottom. The message you type into the AI is called a **prompt** (the instruction or question you give the AI so it knows how to help you).

To begin, simply click on the large empty text box at the bottom of your screen. This will make your keyboard appear if you're on a tablet or phone, or make a cursor blink if you're on a computer, showing it's ready for you to type. Then, type a simple greeting.

What you see: A large, often white or light-colored, empty box at the bottom of the screen, usually labeled "Message" or "Type your prompt here". **What happens next:** As you type, your message will appear in the box. When you press Enter (or tap a Send button, often an arrow icon), your message will move into the main conversation area, and the AI will begin to process your request. **If it looks different:** The typing box might be at the top of the screen or on the side. Look for the largest area where you can type, usually with a small arrow or "send" button nearby. **You'll know it worked when:** Your greeting (e.g., "Hello!") appears in the main chat area, and the AI responds with its own greeting or acknowledgement.

Example: Type **Hello, how are you today?** into the chat box and press Enter.

Step 3 — Write your first clear question (your 'prompt')

The biggest hurdle for beginners is knowing what to type to the AI. You don't need to use special computer code or keywords. Simply type in full, natural sentences, just as if you were talking to a friendly person who wants to help. The clearer your question (your **prompt**), the better the AI can understand and assist you.

Think about what you want to achieve, then phrase it as a polite request or question. It's like asking a librarian for help: the more detail you give, the better they can guide you. After typing your request, press Enter or click the Send button (often a paper airplane or arrow icon).

What you see: Your previous conversation with the AI, and the empty chat box at the bottom of the screen. **What happens next:** The AI will "think" for a few moments (you might see a typing animation or a loading spinner), then it will generate and display its response directly in the chat window, just above your last message. **If it looks different:** Some AI tools might offer voice input options (a microphone icon). You can

try those too, but typing is usually more precise for beginners. If the AI seems confused, try rephrasing your question more simply. **You'll know it worked when:** The AI generates a complete answer to your question, appearing as new text in the conversation history.

Example: Instead of typing short, blunt terms like "scone recipe easy", type: **Could you please share a simple recipe for traditional scones that uses basic ingredients I likely have in my pantry?**

Step 4 — Ask the AI to refine its answer

The AI is incredibly flexible and patient. If the first answer it gives you is too long, too complicated, or not quite right, you can ask it to change the result. You don't need to start a new conversation; just type a follow-up request in the same chat. You can ask it to make things simpler, shorter, or even more fun.

This process of giving feedback and asking for adjustments is called **iterating** (which just means trying again, but with improvements, based on what you learned from the last attempt). Don't be afraid to ask for exactly what you need. Type your request into the same chat box you used before, and press Enter or the Send button.

What you see: The AI's previous answer displayed in the chat, and the empty chat box waiting for your next instruction. **What happens next:** The AI will take your feedback and generate a revised version of its previous answer, incorporating your new instructions. It will display this new answer in the chat below your request. **If it looks different:** Sometimes the AI might apologize if it didn't understand. If this happens, try rephrasing your request for changes even more simply and directly. **You'll know it worked when:** The AI provides a new, updated version of its original answer that better matches what you asked for (e.g., a shorter recipe or a different tone).

Example: If the recipe it gave you looks too detailed, you can type: **That recipe looks good, but could you please simplify it into just five clear, numbered steps?**

Step 5 — Understand what AI can't do (and stay safe)

While AI is a wonderful helper, you must treat it like a polite stranger at a library information desk. They are happy to help you find information, but you would never hand them your bank card or tell them your medical history. The AI lives *only* within its chat window in most cases; it cannot handle sensitive personal data the way a doctor or bank can. (Some newer AI tools can browse the web or interact with applications, but they still operate mainly within their chat interface for everyday use.)

Never type your home address, phone number, passwords, bank details, or private medical worries into an AI. Its primary purpose is to generate text, not to handle sensitive personal data. Protecting your privacy is essential, and the AI cannot truly keep your information private in the way a doctor or bank can.

What you see: Your ongoing conversation with the AI. No changes on screen will indicate a privacy risk, so this step relies on your awareness. **What happens next:** If you stick to safe topics, nothing negative will happen. The AI will continue to provide helpful, general information. If you accidentally share personal data, the AI might even warn you, but it's best not to rely on that. **If it looks different:** Some AI tools have optional privacy settings in your account menu (e.g., "Data & Privacy"). You can explore these, but the core rule of not sharing personal details always applies. **You'll know it worked when:** You feel confident using the AI without worrying about accidentally giving away private information, knowing its limitations regarding sensitive data.

Example:

- **Safe prompt: What are some gentle exercises to help improve balance at home?**
- **Unsafe prompt: My name is Margaret, I live at 12 Main Street, and my doctor just prescribed me this specific medication for my heart…**

Common mistakes

- **Expecting perfect accuracy:** Sometimes AI makes things up with absolute confidence. Tech experts call this a **hallucination** (when an AI confidently invents a false fact or detail). Always double-check important details—like travel times, opening hours, or historical dates—on an official website or trusted source.
- **Using it for medical, financial, or legal decisions:** AI is brilliant for brainstorming dinner ideas or drafting emails, but it is not a doctor, accountant, or lawyer. Always consult real, qualified professionals for medical, legal, or financial advice.
- **Being too vague in your prompt:** If your first attempt at a question doesn't give a good answer, it's often because the AI didn't have enough detail to work with. **Fix:** Instead of "write about dogs", try "Write a short, heartwarming story about a Golden Retriever who helps an elderly neighbor in a small American town."

What this means for you

- **In everyday life:** You can easily draft tricky emails to businesses, write warm RSVP notes for family gatherings, or get creative ideas for using up leftovers in your kitchen. It's like having a thoughtful assistant available 24/7.

- **For your business or work:** You can quickly organize newsletter articles, write clear meeting agendas, or brainstorm fundraising ideas for your local club. It can help you save time on routine writing tasks.
- **If you're just getting started:** The simplest way to build confidence is to play. Do not worry about saying the "wrong" thing; the AI is infinitely patient and will never judge you. Start with simple questions you know the answer to, just to see how it works.

Try it now

Open your chosen AI tool right now, type this exact sentence into the chat box, and press Enter:

Can you please write a warm RSVP email to my friend Sarah, letting her know I would love to attend her afternoon tea party next Tuesday?

How to pick the right free AI app when you're just getting started

5 steps · Read online

> **In short:** All five major free AI apps — ChatGPT, Gemini, Claude, Copilot, and Le Chat — can handle the three tasks beginners try most: writing help, answering questions, and coding assistance. They work differently under the hood, but for everyday use the real difference comes down to integrations, conversation style, and how each one fits into your existing tools.

The idea that one AI assistant is simply "the best" is a myth. Each tool is built by a different company, trained in its own way, and optimized for its own strengths. For a beginner, this actually works in your favor — it means there is a good fit for almost every need. This guide walks you through comparing the free tiers of the five most widely available assistants so you can choose with confidence, not guesswork.

Before you start

- An internet connection and any device (phone, tablet, or computer).
- At least one free account — you do not need to pay for anything to follow this guide.
- About 15 minutes to test a few tools side by side.

Step 1 — Decide what you actually want to do

Before looking at any app, name your goal in one sentence. The three tasks beginners try most are:

- **Writing and drafting** — emails, social posts, summaries, brainstorming ideas.
- **Answering questions and explaining things** — looking something up, learning a new topic.
- **Getting help with code** — understanding a programming concept or writing a simple script.

Most free AI apps handle all three reasonably well. The differences show up in *how well* they handle each one, and whether they fit into tools you already use.

💬 Example: "I want to draft a short email to my manager explaining why a project deadline needs to move." You'll know it worked when you can state your goal clearly — that sentence is your comparison checklist.

Step 2 — Compare writing and creative tools

For drafting, editing, and generating ideas, each app takes a different approach.

ChatGPT (OpenAI) and **Gemini** (Google) both produce quick, well-formatted drafts across many formats. They are good all-around choices for everyday writing tasks.

Claude (Anthropic) tends to give longer, more detailed responses and often catches nuance in longer pieces. If you write a lot of reports, emails, or documents, it may feel like a thoughtful editor rather than a speed-draft tool.

Copilot (Microsoft) connects directly to Microsoft Word and Outlook. If you already pay for Microsoft 365, you can generate and edit text inside those apps without switching windows.

Le Chat (Mistral AI) handles multilingual writing naturally — a useful advantage if you work across languages or need content in a language other than English.

All five will produce usable drafts. The question is whether one of these fits more naturally into your existing workflow.

💬 Example prompt to try in two tools: "Write a 150-word product description for a small handmade candle business." You'll know it worked when you can compare the results side by side and spot the style differences — one may be more energetic, another more formal.

Step 3 — Compare question-answering and knowledge tools

When you need an answer rather than a draft, the tools differ most on two points: how current the information is, and how deeply they explain.

Gemini draws on Google's search infrastructure for real-time information. If your question involves recent events, current prices, or anything time-sensitive, it may have an edge here on the free tier.

ChatGPT and **Claude** have knowledge cutoff dates — they do not browse the live web the same way. They explain concepts in depth and handle follow-up questions particularly well, making them strong for learning and research conversations.

Copilot uses the same underlying model as ChatGPT but displays search results alongside its answer, which some users find helpful for verifying factual claims.

Le Chat handles questions in multiple languages fluidly and can switch between them mid-conversation.

If your primary need is quick, up-to-date answers, try Gemini for that task. For learning and exploring a topic over several questions, ChatGPT or Claude may feel more natural.

Example prompt: "Explain what a budget spreadsheet is to someone who has never made one." You'll know it worked when the explanation uses everyday language and builds on ideas step by step — not when it throws jargon at you.

Step 4 — Compare coding and learning tools

If you are curious about how code works or want help writing a simple script, here is how the tools line up.

ChatGPT and **Copilot** both generate code snippets with plain-language explanations. Copilot's advantage shows up if you write code inside Microsoft development tools — the integration is built in.

Claude tends to explain code in more detail and can walk through longer programs step by step, which some beginners find helpful.

Le Chat supports a wide range of programming languages and can switch between them in the same conversation.

Gemini integrates with Google tools and works well for web-related code and simpler scripts.

None of the free tiers are full development environments — they will not run your code for you. But for understanding, learning, and writing small scripts, all five are worth trying.

Example prompt: "Write a simple Python script that asks for your name and prints a greeting." You'll know it worked when you can copy the result, paste it into a free online Python editor, and run it without errors.

Step 5 — Make your decision in three questions

After testing all five tools with the same prompt, answer three quick questions to narrow your choice:

1. Which response felt most natural to read?
2. Did any tool integrate with an app you already use (Microsoft 365, Google Workspace)?
3. Did any tool answer in the language or style you preferred?

Write the answers down. If you are a writer or content creator, you will likely lean toward ChatGPT or Gemini for their versatility. If you already use Microsoft tools every day, Copilot removes friction by living where you already work. If you value thoughtful, detailed responses, Claude may feel like the best fit. If you work across languages or want a tool that switches easily between them, Le Chat is worth a closer look.

Example: After comparing results from Steps 2 through 4, note which tool you would open first if you had a real task due tomorrow. You'll know it worked when you feel a clear pull toward one tool for one specific reason — that is a sign the fit is right, not just familiarity.

Common mistakes

- **Assuming one tool beats all others.** Every comparison you read online reflects one person's use case. Your needs may be different. The tools are genuinely close in quality for most everyday tasks.
- **Missing tools you already have access to.** If you already pay for Microsoft 365, Copilot may already be on your computer. If you have a Google account, Gemini is a free web app. Check your existing subscriptions before signing up for something new.
- **Sticking with the first tool you try.** First impressions vary. A tool that felt awkward at first may click once you adjust how you phrase your prompts. If a tool almost worked, try rephrasing before switching.

What this means for you

- **In everyday life:** A good AI assistant can cut the time you spend staring at a blank page, wrestling with an unfamiliar topic, or figuring out where to start on a task. The right tool makes that feel easy rather than frustrating.
- **For your business or work:** If you handle any written communication, internal documents, or basic research as part of your job, even a free AI assistant can take the first draft off your plate. The best fit is the one that connects to the tools you already use.
- **If you are just getting started:** Pick one tool and commit to using it for three real tasks this week. At the end of the week, ask yourself whether it felt natural. If yes, keep going. If not, try the next one on this list.

Try it now

Open two of the five tools in separate tabs or windows right now. Paste the same prompt into both:

> "Give me three ideas for a weekend project I could finish in two hours with no special skills or tools."

Compare the results side by side. Notice the differences in tone, length, and creativity. That feeling of "this one speaks my language" is your signal. Trust it — and open that one first next time you have a task to tackle.

How to study and learn faster with AI – a beginner's step-by-step guide

5 steps · Read online

By the end of this guide, you'll be able to ask an AI to break down a tough topic, create a handy summary, and build a short quiz – so you can master new material quicker than before. This guide is written for anyone who wants to study smarter, not harder, by using a free digital learning assistant.

Before you start

- **An active internet connection** and a web browser (like Chrome, Firefox, Safari, or Edge).
- **A free OpenAI ChatGPT account.** You'll use this popular AI tool as our example. If you don't have one, you'll create one in Step 1.
- **A device** such as a computer, laptop, tablet, or smartphone. The steps are written for a web browser experience, but the concepts apply to mobile apps too.
- **A way to take notes** (either pen and paper or a digital note-taking app) to write down your quiz answers.
- **Rough total time:** Allow about 15-20 minutes to go through all the steps and try them out with your own topic.

Step 1 — Set up your AI chat

First, open your web browser and go to the AI chat website. For this guide, we'll use ChatGPT, a **large language model (LLM)** – which is essentially a powerful computer program trained on a huge amount of text, allowing it to understand and generate human-like language. You'll need to sign in, or create a free account if you don't have one yet.

You'll know it worked when you see a blank text box ready for your message, usually at the bottom of the screen.

Example: Visit `https://chat.openai.com`. On the first page, click the "Sign up" button. Follow the prompts to enter your email, create a password, and verify your email address. You may also need to provide your name and date of birth. After completing the short setup, you'll land on the main chat interface. If it looks different, look for "Log In" or "Sign Up" usually in the top-right corner or center of the screen.

Step 2 — Ask for a plain-language explanation

Now that you're in the chat, think of a subject you want to understand better – maybe "photosynthesis" from biology or "supply-chain logistics" for a business course. Type a clear instruction, called a **prompt**, into the text box that tells the AI exactly what you need. Make sure to specify the level of detail, like asking it to explain it "simply" or "like I'm a beginner."

You'll know it worked when the AI starts typing out an explanation in response to your message, appearing just above your prompt.

Example: Type into the message box:

```
Explain photosynthesis in simple terms, like I'm a high-school student with no science background.
```

Then, press the 'Enter' key on your keyboard or click the 'Send message' button (often a paper plane icon) next to the text box.

Step 3 — Request a bullet-point summary

After getting a basic description, it's helpful to condense the main ideas. This makes it easier to remember the core concepts. In the same chat, ask the AI to summarize the key points into a short, easy-to-read list. This helps you grasp the structure of the topic at a glance, like creating a quick cheat sheet for yourself.

You'll know it worked when the AI presents a concise, numbered or bulleted list directly below your request, making the key information easy to skim.

Example: Type into the message box:

```
Summarize the main steps of photosynthesis in 5 bullet points.
```

If it looks different, and the summary isn't quite right, try being more specific, for example: "Can you re-summarize that for me, focusing only on the cause and effect, in three short sentences?"

Step 4 — Create a quick practice quiz

Testing yourself is one of the fastest ways to figure out what you've actually learned and what you still need to review. Ask the AI to generate a few questions based on the summary it just provided. This step turns passive reading into active recall, which is a proven technique for strengthening your memory and understanding. You can choose the type of quiz, like multiple-choice or short answer.

You'll know it worked when the AI displays a short list of questions, each with several options to choose from, or a clear question with space for your answer.

Example: Type into the message box:

```
Make three multiple-choice questions about the bullet-point summary of photosynthesis. Include one correct answer and three distractors for
```

If it looks different, and the questions aren't in the format you wanted, rephrase your prompt to specify, for instance: "Instead, can you make three open-ended questions?"

Step 5 — Check your answers and get clarification

Now that you have your quiz, write down your answers on a piece of paper or in your notes app. Once you're done, type your answers back into the chat. Then, ask the AI to tell you which ones were correct and, importantly, why any wrong answers were incorrect. This feedback loop helps to correct misunderstandings and deepens your knowledge, converting a guess into solid understanding.

You'll know it worked when the AI clearly identifies which of your answers were correct and provides explanations for any you got wrong, or expands on the right ones.

Example: Type into the message box:

```
Here are my answers to the photosynthesis quiz: 1-B, 2-A, 3-C. Which are correct, and can you explain any mistakes?
```

If it looks different, and the AI seems confused, make sure your answers are clearly formatted (e.g., "Question 1: A, Question 2: B") so it can easily understand what you're checking.

Common mistakes

- **Being too vague with your prompt:** Asking "Explain photosynthesis" without specifying the audience can give you a textbook-dense reply that's hard to digest.
  - **The fix:** Always add context like "in simple language," "for a beginner," or "like I'm in primary school."
- **Skipping the quiz or checking answers:** It's tempting to just read the summary and move on. Without actively testing yourself, you won't truly know which gaps remain in your understanding.
  - **The fix:** Treat the quiz and answer check as a critical part of the learning process. It forces your brain to retrieve information, which solidifies memory.
- **Relying solely on one AI answer:** AI tools can occasionally **hallucinate** (confidently make up facts or information that isn't true). While usually accurate, it's not infallible.
  - **The fix:** If something feels off or incredibly important, copy the same question into a new chat, or search a reliable source (like a textbook or reputable website) to verify the information.

What this means for you

- **In everyday life:** Want to understand your mortgage statement, a complex news article, or a new cooking recipe? Use the exact same steps – ask for a plain explanation, a short list of key points, then a quick quiz to lock in your understanding. It's like having a personal expert on standby.
- **For your business or work:** Turn a complex report or industry jargon into a one-page briefing for your team, and follow it with a short quiz to ensure everyone is on the same page about new procedures or project goals.
- **If you're just getting started:** The easiest first move is to pick a current study topic, ask the AI for a simple explanation, and then see how the summary looks. You'll instantly feel the benefit of having a personal tutor on demand.

Try it now

Open your AI chat tool, pick the subject you're studying this week, and type the first prompt from Step 2. In the next two minutes, you'll have a clear, beginner-friendly explanation right in front of you – the essential first step towards faster, more effective learning.

Using Perplexity AI for Quick Study Summaries and Exam Prep

5 steps · Read online

Imagine having a tireless study assistant that doesn't just summarize your textbook, but actually shows you exactly where it found the answers. This guide will show you how to use Perplexity AI to create trustworthy study guides, clear explainers, and practice quizzes for any subject, helping you get ready for assessments with confidence. Perplexity is an **AI-powered search engine** (a smart online tool that uses artificial intelligence to find information) that searches the live internet to find up-to-date information, unlike some traditional tools that rely only on what they learned during training.

Before you start

- **A device with an internet connection:** You'll need a computer, tablet, or smartphone to access Perplexity AI.
- **A free Perplexity account:** You can easily sign up on their website with your email address. The basic, free plan provides access to the core features needed for this guide, including a limited number of "Copilot" queries per day, which offers more interactive help.
- **Rough total time:** This guide should take you about 10-15 minutes to follow along with and complete your first summary and quiz.

Step 1 — Open Perplexity AI and get ready to search

Your first step is to open the Perplexity AI website or launch the app on your device. This will take you to the main interface where you can start asking questions.

When you open Perplexity, you'll see a clean homepage with a prominent search bar in the center of the screen, often displaying a helpful prompt like "Ask anything...". After the page loads, you'll be ready to type your first query. If you're using the mobile app, the layout will be similar, but the search bar might be closer to the bottom. You'll know it worked when the Perplexity homepage or search interface is clearly visible on your screen.

Step 2 — Choose your search focus

Before typing your study question, you can guide Perplexity on where to search by using its "Focus" feature. This is like telling a librarian which section of the library you're interested in.

You'll usually see a button labeled "Focus" or a dropdown menu near the main search bar, sometimes with the default "All" selected. Tapping or clicking this will open a list of options such as "Academic", "Writing", "YouTube", and "Reddit". When you select one, a small icon or text will appear next to the search bar indicating your choice. If it looks different, look for a similar option in the bottom-left corner of the screen (on desktop) or within a "Settings" menu (on mobile apps). You'll know it worked when your chosen focus (e.g., "Academic") is displayed next to the search bar. Example: Click on the **Focus** button and select **Academic** if you are preparing an essay for a history or science topic.

Step 3 — Craft a clear question (prompt)

To get the best study summaries, you need to write a clear **prompt** (the instruction you give to an AI). Think of it like giving precise instructions to a diligent research assistant.

Type your detailed question into the search bar. For study purposes, tell Perplexity what topic you want to learn and exactly how you want the information organized (e.g., "as a comparison table," "in bullet points"). After typing, press the 'Enter' key or click the 'Search' button. Perplexity will then begin to process your request, often showing a "Generating" message. If the AI doesn't understand, it might ask a clarifying question, in which case you can rephrase your prompt. You'll know it worked when Perplexity starts typing out a summary or explanation based on your specific request. Example: "I am a biology student preparing for an exam. Summarize the key differences between mitosis and meiosis. Present this as a simple comparison table followed by bullet points explaining the main phases of each."

Step 4 — Review and verify sources (citations)

One of the main challenges with AI is **hallucination** (when an AI confidently makes up incorrect information). Perplexity helps you beat this by providing **citations** (direct links to the original website where it found the information).

Inside the summary Perplexity generates, you will see small numbers in brackets (like [1], [2]) next to facts. These are clickable links. Click on one of these numbers to open the original source website in a new browser tab. After the new tab loads, quickly scan the original source to verify that the AI has accurately summarized the information. If it looks different, ensure your browser isn't blocking pop-ups, or try right-clicking the number to open the link in a new tab manually. You'll know it worked when you can see the original article or document that Perplexity used for that specific piece of information.

Step 5 — Generate a self-test quiz

The best way to prepare for an exam is through **active recall** (testing your memory rather than just rereading your notes). You can ask Perplexity to turn the summary it just made into a quick multiple-choice quiz.

Once your summary is complete, type a follow-up question into the same chat window, asking Perplexity to create a quiz. The AI will then read its previous response and generate questions based on that content. You'll see the questions appear in your chat, often with multiple-choice options. If it provides answers immediately, simply ask it to hide them. You'll know it worked when you see a list of numbered questions, usually with options (A, B, C, D), ready for you to test your knowledge. 💬 Example: "Based on the summary of cell division you just created, write a 5-question multiple-choice quiz. Hide the answers at the very bottom of your response so I cannot see them while I test myself."

Common mistakes

- **Blindly trusting the summary:** Even the best AI can misinterpret or make up information.
  - **Fix:** Always click at least one or two of the bracketed citations ([1], [2]) to cross-reference the information with the original source. It's like checking the references in a textbook.
- **Being too vague in your prompt:** If you just ask for "American history," you'll get a very generic, long wall of text.
  - **Fix:** Be specific about the era, topic, or event you are studying. Include keywords like "summarize," "explain," or "compare" to guide the AI.
- **Forgetting to export your best summaries:** Perplexity saves your chat history, but it's wise to keep important study notes elsewhere.
  - **Fix:** After a useful summary or quiz is generated, copy and paste the text into a document on your computer (like Google Docs or Microsoft Word) or a note-taking app to save it for later review.

What this means for you

- **In everyday life:** You can quickly research any hobby, household project, or local news story with reliable, verified facts in seconds, getting accurate answers to your curious questions.
- **For your business or work:** You can efficiently research industry trends, market competitors, or new regulations, knowing exactly where the data came from and that it's up-to-date.
- **If you're just getting started:** Use Perplexity as a super-fast, reliable starting point for any research task to get your head around difficult topics without getting lost in endless search results.

Try it now

Open Perplexity AI and ask it: "What are the main stages of photosynthesis, presented as a bulleted list with simple explanations for each stage?" Make sure to click at least one citation to verify a fact.

Use Alexa+ to Pick Up Common Phrases in Another Language

5 steps · Read online

By the end of this guide, you'll know how to ask your Amazon Echo device for a translated phrase, hear it spoken out loud, and have a short back-and-forth to try it for yourself. This is for complete beginners who want a low-pressure way to dip into a new language at home. Honest note up front: **Alexa+** (Amazon's AI-powered upgrade to the standard Alexa assistant) isn't a dedicated language tutor. It won't grade your accent or run a structured lesson, but it IS great for quick translations, examples and casual practice you can fit into a normal day.

Before you start

- An **Amazon Echo device** (such as an Echo Dot or Echo Show) that's already set up and connected to your home Wi-Fi.
- **Alexa+ enabled on your account.** Amazon currently offers Alexa+ to a limited set of users. If you're not sure whether you have it, ask your device *"Alexa, what's new with you?"* — a more chatty, natural-sounding reply suggests Alexa+ is on; a more scripted reply suggests the older version.
- A **language in mind** to practice. This guide uses **Spanish** as the example, but the same approach works for French, German, Italian, Japanese and many others.
- A quiet room so the microphone can pick you up clearly, and around **10 minutes** for a first practice session.

Step 1 — Wake Alexa+ and ask for a phrase

Say the wake word out loud — **"Alexa…"** — followed by your request in plain English. There's no special "language mode" to open; you just talk to Alexa+ the way you'd talk to a helpful friend. The light ring on top of your Echo turns blue (or a small listening animation appears on an Echo Show screen) to show it's heard you. **If it looks different:** on Echo Show devices, a thin Alexa bar may slide up from the bottom of the screen with what you said written on it. If your device only answers in the old, more scripted Alexa style, the Alexa+ upgrade may not yet be active on your account — check Amazon's help pages for the latest rollout details. 💬 Example: say *"Alexa, how do I say 'where is the bathroom?' in Spanish?"* You'll know it worked when Alexa+ speaks the translated phrase out loud and usually adds a short English explanation alongside it.

Step 2 — Ask for an example sentence and a slower version

Once you've heard the phrase, ask Alexa+ to put it into a short example sentence, and then ask it to slow down so you can really hear how each word sounds. This is where the conversational style of Alexa+ helps — you can keep asking follow-up questions in plain English without starting over. **If it looks different:** if Alexa+ speaks too fast on the first try, simply ask again and add *"slower, please"*. If it gives a very short answer and stops, just keep the conversation going with another question; there's no need to begin again from scratch. 💬 Example: say *"Can you use that in a short example sentence? And could you say it more slowly?"* You'll know it worked when Alexa+ repeats the phrase at a noticeably slower pace and gives you a sample sentence you could imagine using in a real situation.

Step 3 — Repeat the phrase out loud and ask for the spelling

Now it's your turn. Say the phrase back out loud, even if it feels a little awkward at first. Then ask Alexa+ to spell it out letter by letter, so you can scribble it down in a notebook if you'd like to remember it later. **If it looks different:** on Echo Show devices, the phrase may already appear on screen alongside what you said — handy if you're a visual learner. If Alexa+ doesn't catch your voice clearly, move a step closer to the device and try again; background noise (kitchen fans, loud music) can confuse the microphone. 💬 Example: say *"Alexa, can you spell that phrase for me? I'd like to write it down."* You'll know it worked when Alexa+ reads the phrase one letter at a time and you can copy it into a notebook without having to guess.

Step 4 — Try a short back-and-forth before going all-in

Before jumping into a full conversation in another language, ask Alexa+ to walk you through how a real exchange would go — first in English, then in Spanish, side by side. This is a low-stakes way to rehearse a scenario like ordering coffee or asking for directions. **If it looks different:** if Alexa+ gives you a single sentence and stops, prompt it with *"what would they say back to me?"* — it usually continues the mini-dialogue. Be aware that Alexa+ isn't built as a role-play tutor, so the conversation may feel a little bumpy; treat it as practice, not a polished lesson, and don't worry about mistakes. 💬 Example: say *"Could we walk through ordering a coffee in Spanish? Show me what I would say, then what the barista would say back."* You'll know it worked when Alexa+ gives you both sides of a small exchange, line by line, in both languages, and you can repeat each line back to yourself.

Step 5 — Wrap up and save your favorite phrases

When you're ready to stop, simply say *"Alexa, stop"* or *"Alexa, that's all, thanks"*. Then grab a small notebook (or your phone's notes app) and write down two or three phrases you'd like to remember. Alexa+ doesn't currently keep a personal "phrases I've learned" list for you, so writing things down yourself is the most reliable way to remember them. **If it looks different:** if you'd like a refresher tomorrow, ask Alexa+ to **set a reminder** for a specific time — it will ping you out loud when the moment arrives. A paid Alexa+ plan, if your account has one, may add extra features over time, but the basic asking-and-listening flow above is what most people can use today. 💬 Example: say *"Thanks! Please set a*

reminder for tomorrow at 6pm to practice my Spanish phrases again." You'll know it worked when Alexa+ confirms the reminder out loud and the light ring returns to its usual idle color.

Common mistakes

- **Treating Alexa+ like a structured language course.** Alexa+ doesn't run lessons with levels, quizzes or progress tracking — it's a general AI helper. Use it for quick look-ups and casual practice, not a full curriculum.
- **Speaking too quietly or mumbling.** The far-field microphone works well, but it struggles in noisy kitchens. Stand about a meter away, face the device, and pause for a full second after each sentence.
- **Expecting perfect pronunciation feedback.** Alexa+ will happily repeat phrases slowly, but it won't grade your accent the way a human tutor would. Use it for rhythm and confidence, not as the final word on how you sound.

What this means for you

- **In everyday life:** a five-minute session each morning while the kettle boils is enough to pick up a handful of useful phrases in a week — no app switching, no screens, just talking to a small speaker on your bench.
- **For your business or work:** if you have clients or colleagues who speak another language, a quick daily session with Alexa+ can help you learn polite greetings and key questions before a trip or meeting, without needing to install a new app.
- **If you're just getting started:** the very easiest first move is to spend two minutes learning just three phrases — hello, thank you, and goodbye — in whichever language you've always been curious about.

Try it now

In the next two minutes, walk over to your Echo device and say: **"Alexa, how do I say 'thank you very much' in Japanese, and could you say it slowly?"** Listen to the reply, repeat it once out loud, and you're already on your way.

The main AI assistants explained simply (and what each is best at)

4 steps · Read online

If you've ever felt overwhelmed by the growing number of AI assistants on the market, this guide is for you. We'll help you spot the right one for your everyday tasks by exploring their unique strengths. You'll learn how to identify which AI is best suited for different situations, from writing help to quick research.

Before you start

- **Internet connection:** You'll need a stable internet connection and a web browser (like Chrome, Firefox, or Safari) to access most of these tools. Many also offer free mobile apps for your smartphone or tablet.
- **Free account:** Most of the assistants we mention offer a "free tier" or trial, which lets you try them out without spending any money. You'll usually need to create a free account with an email address.
- **Device:** You can follow along using a computer, laptop, tablet, or smartphone. The screens might look slightly different, but the core steps remain the same.
- **Time commitment:** This guide will take roughly 5-7 minutes to read and understand. You can spend an extra 5-10 minutes trying out one or two of the assistants for yourself after reading.

Step 1 — Meet the friendly all-rounders (ChatGPT and Gemini)

These assistants are like versatile pocket-knives, great for a wide range of tasks. Both run on a **large language model (LLM)** – think of this as the AI's "brain" that lets it understand your questions and generate human-like text responses. Your action here is to mentally consider how these tools work.

When you open ChatGPT (from OpenAI) or Gemini (from Google), you'll typically see a clean chat window, often with a large text box at the bottom where you type your questions or "prompts." After you type your request and press Enter, the AI will start generating text in response. If the screen looks different, don't worry – most AI chat tools will have a clear input area marked "Ask me anything," "Type your message," or similar.

Example: "Give me a step-by-step plan to organize my garage this weekend, including how to categorise items."

You'll know it worked when the AI starts typing out its response to your prompt, offering suggestions for your garage organization. While these tools can draft plans, they can't physically tidy your garage for you!

Step 2 — Use the work-and-writing specialists (Claude and Copilot)

If your main need is professional writing or office-based tasks, these assistants are designed to shine. Your action in this step is to imagine using one of these for a work-related task.

For Claude (from Anthropic), you'll find a similar chat interface to the all-rounders, usually with a clear text input box. When you enter a prompt asking for help with writing or summarizing, Claude will generate a response, often known for its thoughtful and natural tone. If you're using Copilot (from Microsoft), you won't typically open a separate app; instead, you'll see a small "Copilot" icon or button appear directly within Microsoft Office apps like Word, Excel, or Outlook. Tapping this often opens a side panel where you can chat with the AI about the document you're working on.

Example: "Read this draft email to my landlord and help me make the tone more polite but firm."

You'll know it worked when Claude gives you a refined version of your email or Copilot, within your Office app, offers suggestions directly relevant to your document. Note that Copilot's deeper integration might require a specific Microsoft 365 subscription. These tools can polish your words, but they can't actually send the email or write the entire report from scratch without your input.

Step 3 — Try the search and social discoverers (Perplexity and Grok)

When you need facts, summaries from the web, or insights into the latest trends, these assistants act like research helpers. Your action for this step is to understand how these tools gather and present information differently from a standard search engine.

If you open Perplexity, you'll typically see a straightforward search bar, much like a regular internet search engine. After you type your question, instead of just a list of links, Perplexity will present a clear, summarized answer based on its web search, complete with footnotes or links to the sources it used. For Grok (from xAI), which is linked to the X platform (formerly Twitter), you might access it via the X app or a dedicated Grok interface. It will provide real-time summaries of news, public opinion, or trending topics as they happen on the platform.

Example: "What are the latest updates on the local public transport extensions, and where can I read the official announcements?"

You'll know it worked when Perplexity provides a concise answer with clickable source links, or Grok presents a summary of current discussions and news related to your query. While these tools are great for summarizing information, they don't perform actual investigative journalism or verify the ultimate truth of every piece of content they find.

Step 4 — Explore international options (Qwen and DeepSeek)

Beyond the most well-known names, there are powerful AI assistants developed by companies around the world, each with unique strengths. Your action here is to consider these tools for tasks that might benefit from specialized capabilities like language translation or complex reasoning.

When you access Qwen (from Alibaba) or DeepSeek, you'll generally encounter a familiar chat interface. For Qwen, after you input text, it excels at translating and understanding cultural nuances, particularly across Asian languages. DeepSeek often stands out for its strong abilities in logical problem-solving, mathematics, and coding assistance. The screen usually features a clear chat history and a text input area at the bottom.

💬 Example: "Translate this business greeting into polite conversational Japanese and explain any cultural etiquette I should keep in mind."

You'll know it worked when the AI provides a translation that feels culturally appropriate, or delivers a well-reasoned answer to a complex, logical problem. These assistants can provide sophisticated language and reasoning help, but they cannot replace a human expert for critical cross-cultural communication or advanced coding projects.

Common mistakes

- **Assuming perfect accuracy:** AI can sometimes "hallucinate" – meaning it confidently generates incorrect or made-up facts. Always double-check critical information like dates, legal details, medical advice, or statistics, as the AI cannot guarantee truth.
- **Stopping after the first answer:** If the AI's first response isn't quite right, don't give up! Treat the assistant like a conversation partner. Guide it by saying, "That's a bit too formal – can you make it sound friendlier?" or "Can you expand on point number three?"
- **Using overly vague prompts:** Giving the AI too little detail can lead to generic or unhelpful answers. Instead of "Write a story," try "Write a short, uplifting story about a shy koala who learns to sing, set in a misty rainforest, for children aged 5-7."

What this means for you

- **In everyday life:** Use Gemini to find a recipe from the ingredients you already have in your fridge, or consider Claude to craft a heartfelt birthday card with just the right tone.
- **For your business or work:** Let Copilot summarize long meeting notes directly inside your document, or have ChatGPT draft social media posts to free up your creative time for other tasks.
- **If you're just getting started:** Pick one assistant that seems most appealing, spend a few minutes exploring its interface and trying a simple prompt, then gradually add more tools to your digital toolkit as you become comfortable.

Try it now

Open a free account with **ChatGPT** or **Gemini** in your web browser right now. Look for the large text box at the bottom of the screen, then copy and paste this prompt to see how the assistant responds: *“Write a friendly three-sentence reminder for my neighbor about the upcoming rubbish collection day.”* You'll know it worked when the AI generates a response directly in the chat window.

How to use Meta AI inside WhatsApp – step-by-step for beginners

6 steps · Read online

By the end of this guide, you'll be able to ask Meta AI for answers, recipe ideas, or quick drafts right from your WhatsApp conversation, even if you've never used AI before. This is perfect for anyone who wants to quickly tap into AI without switching between apps.

Before you start

- You need a WhatsApp account set up on your phone (either an iPhone running iOS or an Android device).
- Your WhatsApp app should be updated to its latest version, as this often brings new features like the Meta AI button.
- Meta AI is generally a free feature within WhatsApp, so you won't need a paid plan for this guide.
- You should be comfortable tapping a few icons and typing messages, no coding or special skills are needed.
- This guide will take roughly 5-10 minutes to follow all the steps and try out your first AI chat.

Step 1 — Find the Meta AI chat

To begin, open your WhatsApp application on your phone. Look for a circular button, usually green with a small sparkle or star icon, often located at the bottom-right or top-right of your main Chats screen. This is the quick way to start a new chat with Meta AI. If you don't see this sparkle button, you can also tap the **Search** bar at the top of your screen. **What you see:** A green circle with a sparkle icon, or the "Search" bar. **What happens next:** Tapping the sparkle button usually opens a direct chat with Meta AI. If using the search bar, it will show results as you type. **If it looks different:** If the sparkle button isn't there, simply tap the **Search** bar at the top of your Chats screen and proceed with the example. Apps can look a bit different depending on your phone or recent updates. 💬 Example: Type "Meta AI" into the search bar. **You'll know it worked when:** You see "Meta AI" appear in your search results or as a chat option, often with a distinctive robot-like avatar or a similar sparkle icon next to its name.

Step 2 — Start your first conversation

Once you've found "Meta AI" (either through the sparkle button or search), tap on its name to open the chat window. This is just like starting a conversation with a friend. You'll likely see a welcome message explaining what Meta AI can do. To activate the AI, you need to send it your first message. **What you see:** The Meta AI chat screen with a welcome message from the AI, and a message typing box at the bottom. **What happens next:** After sending your message, Meta AI will process it using a **Large Language Model (LLM)** – think of an LLM as a very smart, digital brain that understands and generates human-like text. It will then type a response back to you. **If it looks different:** Instead of a "Start chat" button, you might just be able to type directly into the message box. This is normal. 💬 Example: Type "Hi Meta AI, what can you do?" and press the send button (usually a paper plane icon). **You'll know it worked when:** Meta AI sends a friendly greeting or an explanation of its capabilities in reply to your message.

Step 3 — Ask a question or give a task

Now that you've started the conversation, it's time to ask Meta AI for something specific. This instruction you give is called a **prompt**. The clearer and more detailed your prompt, the better and more useful Meta AI's response will be. Think of it as telling a helpful assistant exactly what you need. **What you see:** Your prompt appears in the chat, and a "typing..." indicator usually shows that Meta AI is working on its response. **What happens next:** Meta AI will generate an answer, idea, or draft based on your prompt and display it in the chat. This process usually takes a few seconds. **If it looks different:** Sometimes the AI replies almost instantly, without a "typing..." indicator. This is fine – it just means it's processed your request very quickly. 💬 Example: Type "What's a simple, healthy breakfast idea I can make in under 5 minutes?" and send it. **You'll know it worked when:** Meta AI's detailed response appears in the chat window, directly addressing your question.

Step 4 — Use the AI's reply

Once Meta AI has generated its response, you have several options for how to use the information. You can copy the text, forward it to another chat, or simply read it and continue the conversation. For now, let's try copying the response so you can paste it anywhere you like. **What you see:** Meta AI's reply is displayed in the chat. Next to or below the message, you might see small icons like an arrow to forward, or you can often tap and hold the message for more options. **What happens next:** The selected text will be copied to your phone's clipboard, ready for you to paste into a note, an email, or another WhatsApp chat. **If it looks different:** If you don't see a specific "copy" button, try tapping and holding Meta AI's message. A small menu should pop up with options, including "Copy". 💬 Example: Tap and hold Meta AI's breakfast recipe, then choose "Copy" from the menu that appears. **You'll know it worked when:** You can switch to another app (like your notes app) and successfully paste the recipe text there. Remember, Meta AI can draft the message, but it can't send it to other people for you.

Step 5 — Keep the conversation flowing

Meta AI remembers what you've discussed earlier in the same chat. This is called its **context window** – think of it like the AI's short-term memory for your current conversation. You can ask follow-up questions or refine your previous request without having to repeat all the details. **What you see:** Your new follow-up question in the chat, appearing below Meta AI's last reply. **What happens next:** Meta AI will consider your

new question in light of everything you've talked about so far in this chat. It will then provide a response that builds on the previous information. **If it looks different:** If the AI seems to "forget" an earlier detail, it might be due to the context window reaching its limit, or your follow-up prompt might be too vague. Just try rephrasing your question with a little more detail. 💬 Example: After getting the breakfast idea, type "Can you suggest some healthy lunch ideas too, keeping the 5-minute prep time in mind?" and send. **You'll know it worked when:** Meta AI provides new lunch ideas that are specifically quick and healthy, without you having to re-explain those requirements.

Step 6 — Start a fresh chat for a new topic

If you want to discuss something completely different that isn't related to your current conversation, it's a good idea to start a brand new chat with Meta AI. This clears its "memory" and gives you a fresh start, ensuring the AI focuses entirely on your new topic. **What you see:** Your main WhatsApp "Chats" screen again. **What happens next:** Tapping the sparkle icon will open a new, blank chat window with Meta AI, ready for a completely unrelated conversation. **If it looks different:** If you can't find the sparkle icon, you might need to go back to the WhatsApp "Chats" screen, then tap the "new chat" icon (usually a speech bubble or a pen-and-paper icon) and select "Meta AI" from your contacts list. 💬 Example: Go back to your main WhatsApp "Chats" screen, then tap the green sparkle icon or the "new chat" icon to begin a fresh conversation. **You'll know it worked when:** The new Meta AI chat opens with a fresh welcome message, indicating a clean slate for your next topic.

Common mistakes

- **Not finding the Meta AI chat:** Sometimes the sparkle button isn't immediately visible, or you might overlook the search bar.
  - **The fix:** Always try the **Search** bar at the top of your WhatsApp Chats screen and type "Meta AI". It will appear in your contacts. Also, ensure your WhatsApp app is fully updated, as new features roll out in versions.
- **Using vague prompts:** Asking things like "Help me with work" is too broad, and Meta AI won't know where to start.
  - **The fix:** Be specific! Provide context, keywords, or examples. For instance, "Draft a short email to my team about our meeting on Friday" is much better.
- **Expecting unlimited memory:** While Meta AI remembers your current chat, it doesn't remember previous chats or things from days ago.
  - **The fix:** If you're switching to a completely new topic or continuing a discussion from a long time ago, start a new chat (see Step 6) and provide all the necessary context in your first prompt.

What this means for you

- **In everyday life:** Need a quick travel tip, a joke for a group chat, or a summarized news headline? Just type it into WhatsApp and Meta AI does the heavy lifting, giving you instant information or creative sparks without leaving your messaging app.
- **For your business or work:** Draft a short client email, get a bullet-point outline for a meeting agenda, or brainstorm taglines for a project without needing to open a separate writing tool. It's a quick assistant right in your pocket.
- **If you're just getting started:** The easiest first move is simply sending a friendly "Hello Meta AI" and seeing its response. This requires no extra apps or sign-ups beyond your existing WhatsApp account.

Try it now

In the next two minutes, open WhatsApp, find "Meta AI" using the sparkle button or the search bar, send the message "What's a fun weekend activity for two people near the beach?" and glance at the answer. You've just taken your first step into AI-powered chatting!

A Beginner's Guide to Using ChatGPT Voice for Hands-Free Tasks

6 steps · Read online

By the end of this guide, you'll be able to start a spoken conversation with ChatGPT and get it to help you cook, drive, or tidy up – all without lifting a finger. This is perfect for anyone who wants a little extra help while multitasking and keeping their hands free.

Before you start

- **ChatGPT app installed:** You'll need the official ChatGPT app on your smartphone or tablet (available for iPhone and Android devices).
- **OpenAI account:** Make sure you're signed in with your OpenAI account. This feature is available on both free and paid plans, though paid users might get access to newer voice models sooner.
- **Quiet spot:** For the best experience, choose a place with minimal background noise so your device can hear your voice clearly.
- **Rough time commitment:** This entire process should take about 5-10 minutes to set up and try your first conversation.

Step 1 — Start the voice conversation

Open the ChatGPT app on your device. Once the main chat screen loads, you'll see a text input box at the bottom. To activate the voice mode, **tap the voice icon** located in the bottom-right corner, next to where you'd normally type your message. This is typically represented by a waveform or sound wave symbol. Tapping it switches the app into *real-time voice* – meaning you can speak to the AI and hear its reply instantly, much like a phone call.

If it looks different: The icon might sometimes move slightly due to app updates, but it's typically in the bottom-right area. Look for a waveform symbol or a microphone icon that suggests voice interaction.

You'll know it worked when a new screen appears with a dynamic animation (often a colorful wave or line) that responds to sound, indicating it's ready to listen.

Step 2 — Choose a voice you like

When you use the voice mode for the very first time, the app will usually present a list of sample voices for ChatGPT to use. **Tap on each option** to hear a short demo of the different tones and accents available, then **select your favorite** by tapping the "Choose" button (or a similar selection option). This voice will be used for all future spoken interactions.

If it looks different: If you don't see the voice selection screen, it means a default voice is already set. You can change this later by going to the ChatGPT app's main "Settings" menu (usually found by tapping your profile picture in the top-left or bottom-left corner of the main chat screen) and looking for a "Speech" or "Voice" option under the ChatGPT settings.

You'll know it worked when the screen transitions from the voice selection menu to the main listening indicator, confirming your chosen voice has been saved.

Step 3 — Give your first spoken instruction

The screen will now display a simple listening indicator – typically a wave animation or a microphone icon – letting you know ChatGPT is ready for your input. **Speak your first *prompt* naturally**, as if you were talking to a friend. A *prompt* is simply the instruction, question, or statement you give to the AI. There's no need to press any button while you talk; the app automatically detects when you've finished speaking based on pauses.

Example: "Hi! I'm planning a weekend trip to the Blue Ridge Mountains. Can you suggest some scenic bushwalks that are good for beginners?"

If it looks different: If the listening indicator doesn't seem to react to your voice, check your phone's microphone permissions for the ChatGPT app in your device's main settings. You might also need to speak a little louder or clearer if you're in a noisy environment.

You'll know it worked when the dynamic animation responds to your voice as you speak, and then after a brief pause, you'll see your words appear as text on the screen, ready for ChatGPT to process.

Step 4 — Hear the reply and respond

After you finish speaking and your words appear on screen, the ChatGPT app will pause briefly to process your request. Then, in your chosen voice, it will **speak the answer back to you**. Listen carefully to the response. Once the AI has finished its reply, the listening indicator will reappear automatically. Simply **speak your follow-up question or comment** immediately – you don't need to tap the screen or press any buttons to continue the conversation.

Example: "Those sound great! What's the best way to get to the Wentworth Falls track if I'm coming from Chicago?"

If it looks different: If the AI cuts off or mishears you, just say "Could you please repeat that?" or "I wasn't finished, sorry!" and continue. Sometimes, a strong accent or background noise can make it difficult for the AI to understand everything perfectly.

You'll know it worked when the AI responds directly to your follow-up, seamlessly continuing the dialogue as if you were talking to a person.

Step 5 — Control your spoken conversation

While you're chatting with ChatGPT, you have a few ways to manage the conversation. If you need it to **pause** while it's speaking, you can simply say "Stop" or "Hold on." To manually **end the entire voice conversation** and return to the text chat screen, look for a prominent 'X' icon or a 'Stop' button that usually appears on the screen during the voice session. Tapping this will close the voice interface.

Example: (While ChatGPT is explaining) "ChatGPT, hold on a moment please."

If it looks different: The exact look and placement of the 'X' or 'Stop' button might vary slightly depending on your device or app version, but it's typically easy to spot. If you can't find it, you can always close the app entirely and reopen it to reset the session.

You'll know it worked when the AI pauses its reply immediately, or when the voice interface closes and you return to the standard text chat screen.

Step 6 — Go fully hands-free

Now that you know how to start and control a voice conversation, you can truly go hands-free! **Place your phone on a kitchen bench**, slip it into a pocket, or set it on the car console. The voice mode is designed to keep listening for your commands and responses, even if the screen dims. You can now use ChatGPT as a constant companion for various tasks without needing to interact with the screen.

Example: "Okay ChatGPT, tell me a short mystery story, but pause every two minutes so I can guess the suspect."

If it looks different: If the AI stops responding when the screen dims, check your phone's power-saving settings; some modes can restrict background app activity. Ensure you have a stable internet connection, as dropped signals can interrupt the conversation.

You'll know it worked when you can continue to ask questions, receive answers, and have a natural back-and-forth with ChatGPT, all while your hands are busy with other activities.

Common mistakes

- **Pausing too long mid-sentence:** The AI may think you're finished and start responding, which can interrupt your thoughts.
  - **Fix:** If this happens, simply say "Hold on, I wasn't done yet" and continue speaking. The AI is usually forgiving and will pick up where you left off.
- **Background noise:** Loud noises like a running tap, frying pans, or traffic can confuse your phone's microphone, making it hard for ChatGPT to understand you.
  - **Fix:** Using earbuds with a built-in microphone helps a lot, as they focus on your voice and filter out ambient sounds.
- **Using keyword strings:** You don't need to speak like a search query (e.g., "Recipe pumpkin chicken fast"). This isn't how conversational AI works.
  - **Fix:** Talk naturally, using full sentences and common language. ChatGPT's language model is designed to understand everyday speech patterns, making your interactions much smoother.

What this means for you

- **In everyday life:** Turn mundane chores like washing dishes, tidying up, or gardening into moments for learning. Ask for recipe ideas, quick facts, shopping list additions, or even play an interactive game – all without needing to touch your phone.
- **For your business or work:** Dictate initial drafts of emails, practice a presentation out loud, or simply think through a complex problem while you're commuting or taking a walk. It can act as a sounding board, helping you organize your thoughts.
- **If you're just getting started:** Try using the voice feature while you're cooking. It keeps the screen clean of sticky fingers and makes the kitchen a more interactive and helpful space. ChatGPT can read out ingredients, convert measurements, or even suggest substitutions on the fly.

Try it now

Grab your phone, open the ChatGPT app, and tap the voice icon in the bottom-right corner. Once the listening screen appears, say:

> "Hi! I have leftover roast chicken, half a pumpkin and some spinach. What quick dinner can I make?"

You'll hear the AI's suggestion right away – that's your first step into having a hands-free conversation with AI!

Getting started with DeepSeek for everyday help

6 steps · Read online

Imagine having a brilliant, patient assistant who doesn't just give you quick answers, but actually sits down and thinks through your trickiest problems step-by-step. By the end of this guide, you will know how to use DeepSeek, a powerful AI tool, to organize your daily life and solve complex tasks, giving you more clarity and saving you time.

Before you start

- **An internet connection:** You'll need to be connected to the internet.
- **A computer or mobile device:** You can access DeepSeek through a web browser on your computer or by downloading their official app on your smartphone or tablet.
- **An email address or phone number:** You will need one of these to create your free DeepSeek account.
- **Rough total time:** This guide should take about 6-8 minutes to follow through completely.

Step 1 — Set up your free account

Your first action is to visit the official DeepSeek website in your web browser or download the DeepSeek app from your device's app store. Once on the site or in the app, you'll see a landing page. Look for a button like **Sign Up** or **Log In**, usually located in the top-right corner or center of the screen. Tap or click this to begin. After choosing to create an account (e.g., with email, Google, or mobile number), you'll be guided through providing your details and verifying them. If it looks different: The buttons might say "Register" or "Create Account". Just pick the option that lets you set up a new account. You'll know it worked when you see a clean, simple chat screen appear, usually with a text box at the bottom ready for your messages.

Step 2 — Turn on "Deep Think"

DeepSeek has a special superpower called **Deep Think**, which is its *reasoning model*. Think of a *reasoning model* like a thoughtful chef who talks through a complex recipe out loud, planning each step before they start chopping, rather than just handing you a plate of food instantly. This mode helps DeepSeek slow down, plan its thoughts, and makes it incredibly accurate for puzzles, math, or tricky organization tasks. On your chat screen, look for a toggle switch or button near the text input box that says **Deep Think**. Click or tap it to turn it on. If it looks different: The toggle might be labeled "Reasoning Mode", "Plan & Execute", or might be hidden under a small settings icon next to the chat box. Look for any option that implies a more thoughtful, step-by-step approach. You'll know it worked when the toggle switch changes color, moves to an "on" position, or shows a clear indicator that Deep Think is active.

Step 3 — Write your first prompt

Now it's time to write your first *prompt*. A *prompt* is simply the instructions or questions you type to the AI, like giving clear directions to a helpful assistant. Because DeepSeek excels at reasoning, you can ask it complex questions that involve budgets, scheduling, or making difficult decisions. Type your instructions into the text box at the bottom of the screen, ensuring you provide as much detail as possible to help the AI understand your needs. If it looks different: The text box might have placeholder text like "Ask me anything..." or "Type your message here." Example: Copy and paste this prompt into the chat box to try it out:

> "I have a busy Saturday ahead. I need to take the dog to the vet at 9:00am (takes 1 hour), buy groceries (takes 45 minutes), and drop off a birthday present at my sister's house, which is a 20-minute drive from the vet. I also need to be home by 1:00pm for a delivery. Can you organize the perfect, stress-free schedule for me, including realistic travel times?" You'll know it worked when your text appears clearly in the input box, ready to be sent.

Step 4 — Send your prompt and observe the 'thinking'

With your detailed prompt ready in the text box, your next action is to send it to DeepSeek. Look for a "send" button, usually an arrow, paper plane icon, or a button labeled "Send" or "Enter", typically located next to your text box. Click or tap this button. What happens next is a small moment of anticipation. DeepSeek will start to process your request, and you'll often see an indicator showing it's "thinking" or "generating" a response. This process might take a few seconds, especially with Deep Think enabled, as it meticulously plans its answer. If it looks different: The send button might be a simple text label that says "Go" or "Submit." Just click the clear action that makes your message go to the AI. You'll know it worked when your prompt disappears from the input box and appears in the chat history above, and you see an animation or text like "DeepSeek is thinking..." appear.

Step 5 — Review DeepSeek's thought process

Once DeepSeek has finished processing your request, it will display its "thought process" before giving you the final answer. This is a unique and valuable feature. You'll see a small section appear, often collapsible, labeled something like "Thought Process" or "Steps Taken". Your action is to click on this section to expand it. What happens next is fascinating: you'll see a detailed breakdown of how the AI approached your problem,

organizing the timeline, double-checking travel distances, and ensuring everything fits your criteria. If it looks different: Look for a small arrow, a "Show details" button, or "Expand" next to the thought process summary. This might also be represented by a small icon that you click. You'll know it worked when the thought process section opens up, revealing several bullet points or numbered steps explaining the AI's internal logic.

Step 6 — Read the final answer

After reviewing DeepSeek's planning steps, scroll down past the expanded thought process. Your next action is simply to read the final, polished answer or solution that DeepSeek has generated based on its detailed reasoning. What you see is the practical outcome of its "thinking" mode – in our example, a perfectly organized, stress-free schedule for your Saturday. The answer will usually be clearly presented, often using bold text, bullet points, or numbered lists to make it easy to follow. If it looks different: The formatting of the final answer might vary slightly, but the core information will be there. It might be a solid block of text if your prompt was very simple. You'll know it worked when you see the complete, well-structured schedule, or the answer to your question, displayed clearly below DeepSeek's thought process.

Common mistakes

- **Impatience with "Deep Think":** Because Deep Think mode is working hard behind the scenes to calculate the best route or solution, it might take 10 to 15 seconds to start writing. Don't worry—it hasn't frozen! The wait is worth the highly accurate result.
- **Keeping your prompts too simple:** DeepSeek thrives on detail, especially in "Deep Think" mode. If you just ask it to "make a schedule," it has to guess too much. Tell it your preferences, your time limits, and your goals to get the best outcome.
- **Forgetting to enable "Deep Think":** If you're asking a complex question and getting a very quick, but potentially less detailed, answer, double-check that you've remembered to toggle "Deep Think" on. It makes a big difference for intricate tasks.

What this means for you

- **In everyday life:** You can use DeepSeek to budget your weekly groceries, plan a holiday itinerary with driving times and stops, or get complex topics (like how home loans work) explained in simple, easy-to-understand terms.
- **For your business or work:** You can use it to draft professional emails for tricky situations, solve complex spreadsheet formulas, or brainstorm creative marketing ideas for your local business, knowing it will think through the nuances.
- **If you're just getting started:** Keep the "Deep Think" toggle on for any task that involves planning, problem-solving, or making decisions. Treat it like a friendly colleague who loves helping you double-check your plans and figure things out.

Try it now

Let's test DeepSeek's brainpower right now. Open your new DeepSeek chat, make sure **Deep Think** is toggled on, and send this prompt:

> "Explain the difference between a solar eclipse and a lunar eclipse as if I am a ten-year-old. Use a fun analogy using kitchen items like oranges and torches."

Using Claude to Understand Complex Topics for Study or Work

6 steps · Read online

Feeling overwhelmed by a new subject or a tricky report? Claude can be your personal tutor, helping you unravel complex information into easy-to-grasp insights. This guide will show you how to use Claude to simplify anything from scientific papers to business jargon, making learning less daunting and boosting your confidence.

Before you start

- An Anthropic account. The free tier of Claude is usually enough for these basic explanation tasks.
- Access to the Claude website (claude.ai) on a computer or mobile device.
- Roughly 10-15 minutes to follow along with the steps.

Step 1 — Open Claude and Start a New Chat

Your first step is to open Claude and begin a fresh conversation. This gives you a clean slate to explore your topic. On your device, open your web browser and type `claude.ai` into the address bar, then press Enter. Once the page loads, look for a "New Chat" button or a similar option, often found on the left side of the screen or as a prominent button in the main window. Clicking this will clear any previous conversation you might have had, getting you ready to ask about your complex topic. If it looks different: Claude's interface might vary slightly over time. If you don't see "New Chat", look for an icon like a plus sign (+) or a speech bubble to start a new conversation. You'll know it worked when you see a blank message box, ready for your input, usually with "Ask Claude anything" or similar text.

Step 2 — Introduce the Topic

Now it's time to tell Claude what you want to understand. You'll give it a "prompt", which is just the instruction you give to an AI. Think of it like telling a smart assistant exactly what you want it to do. In the message box at the bottom of the screen, type a clear request asking Claude to explain your complex topic. Be specific about the subject. Once you've typed your request, press the Enter key on your keyboard or click the 'Send' button (often a paper airplane icon) next to the message box. Claude will then begin to process your request and generate an initial explanation. Example: `Explain the concept of 'quantum entanglement' in physics.` You'll know it worked when Claude starts typing out its response directly in the chat window, appearing as a new message above your prompt.

Step 3 — Ask for Simpler Terms

Claude's initial explanation might still use some technical language. To make it truly accessible, you'll ask for "simpler terms" or an "analogy" – explaining something unfamiliar by comparing it to something familiar, like explaining how a computer processes information by comparing it to a kitchen preparing a meal. In the message box, type a follow-up request, building on the previous answer. For instance, you can ask it to rephrase the explanation. Claude will then take its previous response and try to simplify it, often using less technical jargon or more common words. Example: `Can you explain that in simpler terms, like you're talking to a high school student?` If it looks different: If Claude seems to misunderstand, clarify that you want the *previous* explanation rephrased. You'll know it worked when Claude provides a new, less complicated explanation, often using everyday language or a relatable comparison.

Step 4 — Request Examples or Scenarios

Abstract explanations are much easier to grasp when you have concrete examples. This step helps you anchor the concept in real-world scenarios. Type another request in the message box, specifically asking for examples related to your topic. You might want examples from everyday life, or specific scenarios related to your study or work. Claude will then generate a list of examples or describe practical scenarios that illustrate the concept you're learning, helping to solidify your understanding. Example: `Please give me two everyday examples where this concept might apply.` You'll know it worked when Claude provides several distinct examples, often using bullet points or numbered lists, that demonstrate the complex topic in action.

Step 5 — Deepen Understanding with Follow-up Questions

Even with simpler terms and examples, you might still have specific parts of the topic that are unclear. This is where you can dive deeper into specific points. Claude has a "context window," which means it remembers what you've discussed so far in the conversation. Identify any specific terms or ideas from Claude's previous answers that you still find confusing. Type a direct question about that specific point in the message box. Claude will use the existing conversation to provide more detailed clarification on that particular aspect, without you needing to re-explain the whole topic. Example: `You mentioned 'decoherence' in the first explanation. Can you elaborate on what that means specifically?` You'll know it worked when Claude focuses its response precisely on the term or idea you asked about, offering a more in-depth explanation than before.

Step 6 — Summarize and Verify

Finally, it's a good idea to consolidate your learning by asking for a summary and considering how to cross-reference important information. Remember that Claude, like all AIs, can sometimes experience "hallucinations" – where it confidently makes something up that isn't true or

doesn't exist. Type a request for a concise summary of the entire discussion. While Claude generates this, consider if the topic is critical (e.g., medical advice, financial figures, legal information). For critical topics, always verify the information with reliable sources like textbooks, academic journals, or expert opinions. Claude will provide a brief overview of the key points covered, allowing you to quickly review what you've learned. 💬 Example: `Can you summarise our conversation about quantum entanglement in three short paragraphs?` You'll know it worked when Claude presents a clear, condensed version of the main ideas from your chat.

Common mistakes

- **Being too vague in your initial prompt:** If you just say "Explain physics," Claude won't know where to start. **Fix:** Always be specific, like "Explain the concept of 'black holes' in astrophysics."
- **Overloading Claude with too many questions at once:** Asking for an explanation, examples, and a summary all in one prompt can sometimes lead to less focused answers. **Fix:** Break down your requests into smaller, sequential steps, as shown in this guide.
- **Trusting critical information without verification:** While Claude is a powerful learning tool, it should not be treated as the sole source of truth, especially for sensitive or crucial data. **Fix:** Always cross-reference facts, figures, and important definitions with reputable, human-verified sources to ensure accuracy.

What this means for you

- **In everyday life:** You can use Claude to understand complex news articles, decode technical terms in product manuals, or get a simpler breakdown of health information from your doctor. It turns confusing information into something you can actually use.
- **For your business or work:** Use Claude to quickly grasp industry reports, understand new software features, or simplify complex legal or financial documents. This can save you time and help you contribute more confidently in meetings.
- **If you're just getting started:** Pick one confusing email you received or an article you skimmed recently. Copy and paste a paragraph you didn't understand into Claude and ask, "Can you explain this paragraph to me in simple terms?"

Try it now

Find a recent news article about a scientific breakthrough, a new piece of technology, or a complex economic policy. Choose one paragraph that you find particularly hard to understand, then paste it into Claude and ask, "Explain this to me like I'm new to this topic." See how quickly it can clarify things for you!

How to Use ChatGPT Voice Mode to Practice Conversations in a New Language

6 steps · Read online

> **In short:** ChatGPT's voice mode turns the app into a patient speaking partner that replies out loud, waits for your answer and never gets tired. You'll need the mobile app, an updated version, and on the free plan you may get only short chats.

ChatGPT's voice mode is one of the most underrated ways to learn a language. Instead of staring at vocabulary lists, you get a real back-and-forth: you speak, the AI speaks back, and you can interrupt, ask "how do you say…?" and try again without anyone rolling their eyes. It works on your phone, in your earphones, while you're doing the dishes. It won't replace a real teacher, but for daily practice it quietly beats most apps.

Before you start

- A free ChatGPT account (sign up with an email at chat.openai.com or in the app). Voice mode may be more limited on the free plan — a Plus subscription gives longer chats and the more natural-sounding advanced voice.
- The **ChatGPT mobile app** on iOS or Android (voice mode lives in the app, though the website is catching up in some regions).
- A working microphone and either earphones or a quiet spot — the AI listens to your voice and replies out loud.
- A target language in mind (Spanish, Mandarin, French, anything ChatGPT can chat in).
- About ten minutes for your first practice session.

Step 1 — Open ChatGPT and switch on voice mode

Tap the **ChatGPT app icon** on your phone — a black-and-white flower-like shape, or just "ChatGPT" in your app drawer. The home screen appears with a chat box at the bottom. Look for a **microphone icon** or a **sound-wave icon** next to the text box, usually on the right side. Tap it once. A new full-screen view opens, the keyboard slides away, and a pulsing circle appears with the words "Listening…" or a sound-wave animation. If you don't see a microphone, look for a **headphones icon** instead, or check the **model picker** at the top of the chat — some versions tuck voice mode inside a dropdown labeled with the current model's name.

Example: nothing to type yet — just tap the icon.

You'll know it worked when the keyboard disappears and a "Listening…" circle or sound-wave animation appears on screen.

Step 2 — Pick a voice you want to chat with

Before you start talking, look for a small **settings icon, gear symbol or row of circles** — usually at the top or bottom of the voice screen. Tap it and a list of named voices appears (the names vary by app version, but they're usually short friendly words). Tap each one to hear a short sample. Choose the voice whose accent and pace feel easiest for you to listen to for a long stretch — there's no "correct" choice, only what your ear prefers. If you can't find a voice picker, don't worry — the app usually picks a sensible default and you can return to this later.

Example: pick the voice whose short sample feels the clearest in your earphones, then move on.

You'll know it worked when you tap a name and hear a short sentence spoken in that voice.

Step 3 — Set the scene for your practice

Speak one short sentence out loud to introduce yourself and your goal. The AI will reply out loud in your target language, and your words will appear as text on screen so you can read along. This is the moment to set the level, the situation, and how strict you want corrections to be. If you don't ask for a specific situation, ChatGPT usually just starts a friendly chat — which is fine, but a tiny bit of context makes practice much more useful straight away.

Example: "I'd like to practice ordering at a café in Italian. I'm a beginner, please speak slowly, and correct my biggest mistakes gently."

You'll know it worked when the AI replies in your chosen language and matches the level you asked for.

Step 4 — Have the actual conversation

Just talk. Say what you would say in the situation you set up, in your target language, using whatever words you have. Pause for a beat after each sentence and the AI will reply out loud. Don't worry about mistakes — keep the flow going. If you didn't catch a word, ask the AI to repeat, slow down or write it. You can also interrupt mid-sentence (just start talking again) and the AI will usually stop and listen. If no audio comes out, check your phone's volume, that your earphones are connected, and that the app hasn't been muted in your phone's settings.

Example: order your coffee — "Buongiorno. Vorrei un caffè e un cornetto, per favore."

You'll know it worked when the AI replies in spoken audio and you can keep the conversation going without touching the screen.

Step 5 — Ask for corrections and explanations

When you want feedback, just ask. You can do this in your native language or in your target language — both work. The AI will explain grammar, suggest a more natural phrase, or give you the word you were stuck on. Some versions of voice mode pop up a small notes panel or a list of corrections under the chat, but don't rely on that — asking out loud is the fastest route. If a correction is too fast or full of jargon, just say "explain that more simply" or "give me an example sentence" and try again.

Example: "Was that sentence correct? And how would a native speaker actually say 'I'd like a coffee'?"

You'll know it worked when you get a clear explanation or a more natural version of what you tried to say.

Step 6 — Wrap up and save what you learnt

When you've had enough, say something like "thanks, let's stop here" or tap the **end-call or close** button — usually a red phone icon or an X in the corner. The conversation turns into a normal text chat you can scroll back through. Ask for a short recap of what you practiced, then either screenshot it, tap the **share icon** (often a square with an arrow pointing up) to send it to yourself, or simply leave the chat open in your history for next time. This recap is your revision material — reading it later is where new words actually stick.

Example: "Thanks! Can you list three new words I used today with a quick English meaning for each?"

You'll know it worked when you get a tidy list (or a short summary) that you can save, screenshot or revisit in your chat history.

Common mistakes

- **Starting cold without a scenario.** Just saying "let's speak French" usually leads to small talk that goes nowhere. Pick a real situation (ordering food, asking for directions, a job interview) and the practice becomes useful almost immediately.
- **Trying to be perfect.** The whole point is making mistakes in a safe space. If you freeze up, drop into English for one sentence and switch back — fluency comes from messy attempts, not from being correct every time.
- **Forgetting the recap.** A spoken chat evaporates fast. Always end with "give me three words I learned today" or similar, and save the answer. Without this last step, much of the learning is gone by the next morning.

What this means for you

- **In everyday life:** Twenty minutes of voice practice while you're walking the dog or doing the dishes adds up faster than a weekly class — and it fits into time you'd otherwise just be listening to music.
- **For your travel or work:** You can rehearse the exact phrases you'll need before a trip or a meeting with a foreign client, then walk in feeling noticeably calmer and more natural.
- **If you're just getting started:** Pick one tiny situation (buying a train ticket, greeting a neighbor), do it for five minutes a day, and review the recap each evening — that's a routine most people can actually keep.

Try it now

Open the ChatGPT app, tap the microphone icon, and say: *"Let's practice introducing ourselves in Japanese. I'm a complete beginner, please speak slowly."* Have a two-minute chat, then ask for three words to remember. That's your first practice session done.

Organize your work with Claude Projects – a step-by-step beginner's guide

6 steps · Read online

By the end of this guide you'll have a brand-new Claude Project that holds all your relevant chats and documents, and you'll know how to ask Claude to draw on everything at once. It's ideal for anyone who wants a clean, goal-focused AI workspace without needing any technical background.

Before you start

- **Claude account:** Sign up for a free Anthropic Claude account (you'll need an email and a password).
- **Device:** A desktop or laptop with a modern web browser (Chrome, Edge, Firefox or Safari).
- **Plan:** Projects are available on Claude — check the current plan details for file size limits.
- **Time:** Expect about 5–7 minutes to complete all steps.

Step 1 — Sign in and open the Claude dashboard

Open your browser, go to **claude.anthropic.com**, and click the **Log in** button. You'll see a screen with fields for your email and password; enter them and press **Enter**. After logging in, the main chat view appears – a white space with a text box at the bottom. Look for the navigation area in the sidebar; you should see a section labeled "Projects". If you don't see it immediately, look for a menu option called "Dashboard" or similar. Example: `email@example.com` + `YourPassword` → press **Enter**.
You'll know it worked when you see a list of past conversations and a side panel showing your Projects.

Step 2 — Create a new Project

In the Projects area, find a button to create a new project – it may say "+ New Project" or show a similar create option. Click it; a small dialog pops up asking for a name and an optional description. A **Project** is just a digital folder that groups together chats and files for a single aim – think of it as a binder you keep on your desk. Type a clear name (e.g., `Q3 Marketing Campaign`) and a short description (e.g., `Copy, ad ideas and feedback for the next quarter`). Example:

```
Name: Q3 Marketing Campaign
Description: All copy, ad ideas and feedback for the upcoming quarter
```

You'll know it worked when the new project appears in your project list, and the main area switches to an empty workspace titled with your project name.

Step 3 — Add existing chats to the Project

With your project open, look for a way to add existing conversations – there may be an "Add conversation" button or similar option. Click it; a panel slides in showing recent chats with check-boxes beside each title. A **chat** (or conversation) is simply a previous dialogue you've had with Claude. Tick the boxes next to the chats you want linked (e.g., "Ad copy brainstorming" and "Email subject line ideas"). After selecting, press the **Add** button at the bottom of the panel. Some versions may let you drag a chat tile into the project area instead of using check-boxes – if you don't see check-boxes, try dragging. Example: Tick "Ad copy brainstorming" and "Email subject line ideas", then click **Add**.
You'll know it worked when those chat titles appear as cards inside the project workspace, indicating they're now part of the same folder.

Step 4 — Upload supporting files

Next, locate an upload option in the project pane – it may be labeled "Attach file", "Upload file", or show as an icon. Click it; a file-chooser window opens. Navigate to a document on your computer (PDF, text file, spreadsheet or image) and select it, or simply drag the file onto the designated area. Claude will process the file, extracting its text so it can be referenced later. If you only see a large rectangle that says "Drag files here or click to upload", that's the same upload area. Example: Drag `CampaignBrief.pdf` from your desktop onto the upload box.
You'll know it worked when the file shows up as a line item inside the project, often with a tiny icon indicating its type (PDF, sheet, etc.) and a loading bar that disappears once processing finishes.

Step 5 — Write a prompt that uses everything

Now you can ask Claude to work across the whole project. In the text box at the bottom of the project screen, type a clear instruction that refers to the combined material. A **prompt** is simply the question or task you give the AI. For example, ask it to summarize the ideas from the attached chats and the PDF, then suggest three headline options. Press **Enter**. Claude will read the linked chats and the uploaded file, then start typing a response that pulls information from all sources. If the answer seems vague, try re-phrasing with more specifics – Claude usually produces better results when the prompt is detailed. Example:

```
Summarize the key ideas from all attached chats and the brief, and suggest three headline options.
```

You'll know it worked when Claude's reply includes references to the chat content and the PDF, such as "Based on the ad copy brainstorming chat..." or "From the CampaignBrief.pdf...".

Step 6 — Share the Project (optional)

If you want teammates to see or edit the workspace, look for a **Share** button – it may be shown as an icon or found inside a "More actions" menu. Click it; a dialog appears offering a copy-able link or an email-invite field. Choose **Copy link**, then paste the URL into an email or chat message. Some versions may place the sharing options inside a dropdown menu; if you don't see the option directly, look for a menu with additional actions. 💬 Example: Click **Share**, then press **Copy link**.
You'll know it worked when a small banner appears saying "Link copied to clipboard" or when an invitation confirmation message shows up.

Wrap-up — What you've achieved and where to go next

You now have a fully functional Claude Project that brings together past chats, files, and a single prompt that pulls from everything. This workspace can be reused: simply add new chats or files as the task evolves, and ask Claude to update its output. In future sessions you can duplicate the project (look for a "Duplicate" or "Copy" option) to keep a record of each campaign or research topic. Remember that Claude can read the content you upload, but it cannot send emails or publish documents for you – you'll need to copy the generated text into your own tools.

Common mistakes

- **Skipping the description:** Without a clear name or description you may forget the project's purpose.
  Fix: Add a brief description when creating the project; you can edit it later by clicking the project title.
- **Linking unrelated chats:** Adding conversations that aren't relevant can confuse Claude's responses.
  Fix: Only select chats that directly support the current goal; remove any that feel off-topic.
- **Forgetting to upload files:** If a file isn't attached, Claude can't reference its content.
  Fix: After each upload, double-check that the file appears in the project list; re-upload if the name is missing.

What this means for you

- **In everyday life:** Planning a home renovation? Create a Project for the remodel, add design sketches, budget spreadsheets, and chat notes about contractors – then ask Claude to summarize costs and suggest next steps.
- **For your business or work:** Use a Project to collate marketing copy drafts, product spec PDFs, and internal feedback chats, allowing Claude to generate a unified campaign brief in minutes.
- **If you're just getting started:** The easiest first move is to set up a simple "Weekend Planning" project, add a few personal chats and a PDF itinerary, and watch Claude combine everything into a single plan.

Try it now

In the next two minutes, open Claude, create a new Project called **"My First Project"**, add one recent chat and upload a small text file, then type the prompt:
💬 Example: `Give me a quick summary of the chat and the uploaded file.`
If Claude returns a combined summary, you've successfully completed the core workflow – congratulations!

Work & email

Automate Email Drafts with AI for Faster Inbox Management

7 steps · Read online

By the end of this guide, you'll have a clever system that writes draft replies for you, helping you spend mere seconds instead of minutes on each routine inbox message. This is perfect for anyone who feels their email is a time-sink, whether you're handling personal messages, customer queries, or internal staff notes. It can draft the email, but remember, it can't send it for you – you'll always have the final say!

Before you start

- **A device with a web browser:** You'll be using web-based tools, so a computer or tablet is ideal.
- **A working email account:** You'll need an account like Gmail, Outlook, or any web-based service.
- **A free AI account with API access:** You'll need access to an AI service that offers an **API key**. Think of this like a special password that lets other programs connect to the AI securely. Examples include OpenAI's platform (for ChatGPT models), Claude, or Google's AI Studio. While signing up for these platforms is often free, using the API incurs small costs, typically a few cents per request. A "free chat account" alone usually isn't enough, as you need developer access.
- **A free account with a no-code automation service:** Tools like Zapier, Make, or IFTTT allow different apps to "talk" to each other without writing any computer code. We'll use Zapier as our example. Be aware that creating multi-step workflows, like the one in this guide, often requires a paid plan after a trial period or for higher usage, even if initial sign-up is free.
- **Time commitment:** This setup will take approximately 20-30 minutes to complete the initial steps and then a few minutes for testing and refining.

Step 1 — Sign up and find your AI API key

Your first step is to get an **API key** from your chosen AI tool. An API key is a unique string of characters that acts like a secret token, allowing external services (like our automation tool) to communicate with the AI securely without you manually logging in each time.

Open your web browser and go to the developer platform for your chosen AI (e.g., platform.openai.com for OpenAI, console.anthropic.com for Claude). After you log in or sign up, you'll typically see a dashboard. Look for a menu item like "API Keys," "Settings," or "Developer" in the left sidebar or under your profile picture. When you click it, you should see an option to "Create new secret key" or "Generate API key." Click this button. A new, long string of letters and numbers will appear on your screen. Immediately copy this key and paste it into a secure location, like a password manager or a temporary text document. You usually won't be able to see it again after you close the window, so don't lose it!

You'll know it worked when you have a long, unique string of characters (your API key) copied and safely stored, confirming your AI platform is ready for connection.

Example (for OpenAI): Go to `platform.openai.com/api-keys`. If you're signed in, you'll see your API Keys page. Click the "+ Create new secret key" button. A new key will pop up. Click the green 'copy' icon next to it and save it.

Step 2 — Create a new email filter

Next, you'll set up a **filter** in your email account. A filter is like a personal email sorter that automatically watches for incoming messages that match specific criteria (like words in the subject line or sender address) and then performs an action. For this guide, we'll create a filter that identifies emails that need an AI-generated draft reply.

Open your email account in a web browser and navigate to your email settings. Look for options labeled "Filters," "Rules," or "Manage rules." For Gmail users, click the gear icon ( ) in the top-right corner, then select "See all settings," and finally, click on "Filters and Blocked Addresses." Here, you'll see a button like "Create a new filter." Click it, and a form will appear with fields such as "From," "To," and "Subject." In the "Subject" field, type a distinct word or phrase that will signal to your system that this email needs an AI-drafted reply, such as "AI Draft Request" (without the quotes). Leave all other fields blank for now. Then, click "Create filter" or "Continue." A warning might appear about not having actions selected; that's fine for now, just confirm.

You'll know it worked when you see your newly created filter listed among your email rules, indicating it's ready to watch for matching emails.

Example (for Gmail): Click > "See all settings" > "Filters and Blocked Addresses" > "Create a new filter." In the "Subject" field, type `AI Draft Request`. Click "Create filter" and confirm any warnings.

Step 3 — Craft your AI prompt template

Now it's time to write the instructions you'll give to your AI, known as a **prompt**. This is like telling a human assistant exactly how to draft a reply for you. A **prompt template** includes special placeholders, which are like empty boxes where the AI will insert specific information from the email it needs to reply to.

Open a simple text editor (like Notepad, Google Docs, or Apple Notes) where you can easily copy and paste text. Think about the kind of replies you want: friendly, formal, concise, or detailed. A good prompt is clear and tells the AI its role and what to do. For instance, you'll instruct it to acknowledge the sender and suggest a few steps. The crucial placeholder you'll use is `{{email_body}}` – this tells the AI exactly where to insert the actual content of the email it's replying to.

You'll know it worked when you have a complete, clear instruction, including `{{email_body}}`, saved in your text editor, which you feel confident will produce the kind of reply you're after.

Example prompt:

```
You are a professional customer service assistant. Write a polite, concise draft email reply in US English. Acknowledge the sender's message
```

```
Email content to reply to:
{{email_body}}
```

Step 4 — Set up the Zapier trigger

This step brings your email and AI together using an **automation service** like Zapier. A "Zap" in Zapier is an automated workflow that starts with a **Trigger** (an event that kicks things off) and then performs one or more **Actions**. Here, our trigger will be a new email matching the filter you set up in Step 2.

Log in to your Zapier account and click the "Create Zap" button. You'll see a screen to choose your Trigger. Search for and select your email provider (e.g., "Gmail"). For the "Event" field, choose "New Email Matching Search." You'll then be asked to connect your email account if you haven't already. In the "Search String" field, enter the exact phrase you used in your email filter in Step 2 (e.g., `subject:"AI Draft Request"`). Follow the prompts to test this trigger. Zapier will try to find a recent email that matches your criteria to ensure it's working.

You'll know it worked when Zapier successfully tests your trigger, indicating it can find new emails that match your filter criteria.

Example (for Zapier):

1. On the "Trigger" screen, search for and select "Gmail."
2. For "Event," choose "New Email Matching Search."
3. For "Search String," type `subject:"AI Draft Request"` (including the quotes).
4. Click "Continue" and follow the steps to connect your Gmail and test the trigger.

Step 5 — Add the AI action in Zapier

Now that Zapier can detect new emails, you need to tell it to send that email's content to your AI to generate a reply. This is the first **Action** step in your Zap.

After successfully setting up your trigger, click the "Add a step" button (often a plus sign or "Action" button). Search for and select your chosen AI tool (e.g., "OpenAI" or "Claude"). For the "Event" field, you'll typically choose an action like "Send Prompt" or "Generate Content." You'll be prompted to connect your AI account; this is where you'll paste your **API key** from Step 1. In the "Prompt" field, paste your entire prompt template from Step 3. Crucially, replace the `{{email_body}}` placeholder by clicking into the prompt field and selecting the "Body" or "Plain Body" option from your previous email trigger step. This tells the AI to use the actual email content. You might also be asked to select a model (e.g., "gpt-3.5-turbo"). Test this action to see if the AI generates a sample response.

You'll know it worked when the test of this action shows a generated text response from the AI, containing a draft reply based on a sample email.

Example (for Zapier):

1. On the "Action" screen, search for and select "OpenAI."
2. For "Event," choose "Send Prompt."
3. Connect your OpenAI account by pasting your API Key from Step 1.
4. In the "Prompt" field, paste your full template. Click into the prompt field where `{{email_body}}` was, and from the dropdown that appears, select "Body Plain" (or similar) from the Gmail trigger step.
5. Test the action.

Step 6 — Create the email draft in Zapier

The final step in your Zap is to take the AI's generated response and save it as a new draft email in your inbox. This is the second **Action** step.

After successfully testing your AI action, click "Add a step" again. Search for and select your email provider (e.g., "Gmail"). For the "Event" field, choose "Create Draft." You'll connect your email account again if needed. In the fields that appear, you'll map the AI's response to the body of your new draft. For "To," you can select the "From Email" from your initial trigger. For "Subject," you might write "RE: " and then select the

"Subject" from your trigger. For the "Body" field, select the output from your AI action step (e.g., "Choice 1 Text" or "Content"). You can also choose to include the original email if you like. Finally, activate your Zap by turning it "On."

You'll know it worked when Zapier successfully tests this action by creating a new draft in your email account, containing the AI's response.

Example (for Zapier):

1. On the "Action" screen, search for and select "Gmail."
2. For "Event," choose "Create Draft."
3. Connect your Gmail account.
4. For "To," select "From Email" from your initial Gmail trigger step.
5. For "Subject," type `RE:` and then select "Subject" from the initial Gmail trigger step.
6. For "Body," select "Choice 1 Text" from your OpenAI action step.
7. Test the action and then turn your Zap "On."

Step 7 — Test your automation and refine

The hard work is done – now for the fun part: testing! You want to make sure your AI assistant is drafting exactly what you expect.

Open your email client and send yourself a brand new email that perfectly matches the filter criteria you set in Step 2 (e.g., using "AI Draft Request" in the subject line). Go to your "Drafts" folder in your email account. After a minute or two, you should see a newly generated draft reply there. Read it carefully. Is the tone right? Is it too long or too short? Does it make sense? If the draft isn't quite right, go back to Zapier, edit your AI action step, and tweak the wording of your prompt template. Change things like "short, professional" to "brief, friendly and informal," or adjust the word count. Then, send yourself another test email and check the new draft. You can repeat this process as many times as needed to get the perfect reply. Always remember to review each draft before hitting send – the AI is there to assist, not replace your judgement.

You'll know it worked when you consistently receive AI-generated drafts in your drafts folder that you would feel comfortable sending, maybe with just a minor edit or two, matching your desired tone and content.

Example: After a test run, you find the draft is a bit too formal. Go back to Zapier, edit your OpenAI action, and change the prompt line "Write a polite, concise draft email reply" to "Write a friendly, informal, and very brief draft email reply." Then, send another test email to yourself.

Common mistakes

- **Not protecting your API key:** This key grants access to your AI account, potentially incurring costs. Store it securely in a password manager and never share it publicly. If it's compromised, immediately generate a new one and delete the old one in your AI tool's settings.
- **Setting the email filter too broadly:** If your email filter (from Step 2) is too general (like "contains 'the'"), you'll get draft replies for newsletters, spam, and personal messages you don't need help with. Make your filter specific to recurring, predictable email types (e.g., `subject:"AI Draft Request"` or `from:"[specific sender email]"`).
- **Skipping the final review step:** While AI is powerful, it can occasionally misinterpret tone, insert invented facts, or miss crucial nuances. Always give the draft a quick read before hitting send to avoid embarrassing mistakes or misleading your recipient.

What this means for you

- **In everyday life:** You can breeze through routine messages like "Can you share the file?" or "Thanks for the update" with a ready-made reply, freeing up minutes for the things you enjoy, like hobbies or spending time with loved ones.
- **For your business or work:** Customer-service teams can handle frequent, simple queries faster, keeping response times low and clients happy without needing to hire extra staff for basic tasks. It can also help you quickly respond to internal requests, keeping projects moving smoothly.
- **If you're just getting started:** The easiest first move is to set a filter for a single, very common phrase you receive (e.g., "thank you for your email") and let the AI draft a very short, one-sentence acknowledgment. You'll see the time saved instantly.

Try it now

In the next two minutes, open your email settings, create a new filter, and set it to watch for the subject word "test" (without quotes). For now, set the action to "Mark as important" and "Apply the filter to matching conversations." This tiny change will already start grouping similar emails, making it easier to apply the AI draft process later. Happy automating!

How to Get Quick Summaries from Long Documents with Claude

7 steps · Read online

Ever feel swamped by lengthy reports, articles, or complex documents? This guide will show you how to use Claude, an **AI (Artificial Intelligence)**, which is like a very smart computer program designed to think and learn, to quickly get the main points from almost any text. By the end, you'll be able to upload a document and ask Claude to summarize it for you, giving you the essential information you need in a flash.

Before you start

- You'll need a free Claude account (or a paid plan if you prefer). This will give you access to Claude's website.
- Access to the Claude website through your computer's web browser.
- A digital document you want to summarize, such as a PDF or a plain text file. Note that Claude has limits on the file size and length it can process, but free plans are usually sufficient for most common documents.
- This whole process should take you about 5-7 minutes.

Step 1 — Access Claude's Website

Your first step is to open your web browser and navigate to the Claude website. Think of Claude as a powerful **Large Language Model (LLM)**, which is like the highly intelligent engine behind an AI, specifically designed to understand and generate human language. You'll usually see a clean, welcoming chat interface where you can begin interacting with the AI.

You'll know it worked when you see a main chat area, often with a text box at the bottom ready for you to type in.

Step 2 — Start a New Conversation

To ensure you're starting fresh, it's a good idea to begin a new conversation with Claude. Look for a "New Chat" button, which is commonly found on the left-hand side of the screen or near the top. Clicking this will give you a clean slate, meaning Claude won't carry over any previous topics or instructions from old chats, so your summary request will be clear.

If it looks different on your screen, don't worry! Apps change often. You might see a small plus (+) icon, or a menu option like "Start new conversation." Simply click whatever seems to open a fresh chat.

You'll know it worked when the main chat window is clear, with no previous messages, and the message box at the bottom is empty, waiting for your input.

Step 3 — Upload Your Document

Now it's time to get your document into Claude. Look for a paperclip icon, usually found near the message box where you type. This icon is universally used for attachments. Clicking it will open your computer's file explorer, allowing you to browse and select the PDF or text file you wish to summarize. Claude can handle various formats, but PDFs and plain text are generally the most reliable for summarisation. Remember, very large files might not upload on a free plan, as Claude, like all AIs, has a **context window** – this is like the AI's short-term memory, limiting how much text it can 'remember' at once.

If you don't see a paperclip, look for an "Attach file" text link or a similar icon that suggests adding a file.

You'll know it worked when the name of your uploaded document appears in the message box, usually next to your cursor, indicating it's ready to be sent.

Step 4 — Craft Your Summary Request

With your document attached, it's time to tell Claude what you want. This instruction is called a **prompt**, which is simply the command you give to the AI. In the message box, type out a clear request for a summary. Being specific helps Claude understand exactly what you're looking for, rather than just giving a generic overview. For instance, you could ask for specific details or a summary tailored to a particular audience.

If your message box is hard to find, it's usually a long, horizontal bar at the very bottom of the chat window.

You'll know it worked when your complete summary request appears in the message box, ready to be sent to Claude. 💬 Example: "Please summarize this document for me, highlighting the main points and key takeaways relevant to small businesses."

Step 5 — Send Your Request

Once your document is attached and your summary request is typed, the next step is to send it to Claude. Look for a "Send" button or a paper plane icon, usually located at the right end of the message box. Clicking this will submit both your document and your prompt to the AI. Claude will then begin processing the information, reading through your document and generating a summary based on your instructions. This process usually happens quite quickly.

If the button looks different, it might be an arrow or just the word "Send". Just click the button that looks like it will submit your message.

You'll know it worked when your message disappears from the box, and Claude's response starts appearing in the main chat window, beginning with your requested summary.

Step 6 — Review and Refine the Summary

After Claude generates the summary, take a moment to read through it. AI responses are usually very helpful, but sometimes they might not hit every single point you need, or you might want more detail on a specific section. This is where you can interact further. If you need more information or a different perspective, just type another prompt in the message box, referring back to the document or the summary Claude just gave you. For example, you could ask Claude to elaborate on a particular section or rephrase the summary for a different audience.

If the answer isn't quite what you expected, don't worry, that's normal with AIs. Just try rephrasing your question or adding more specific details in your next prompt.

You'll know it worked when you've either received a summary that fully meets your needs or have asked follow-up questions that yield the specific additional information you were looking for. 💬 Example: "Can you list the top three challenges mentioned in the document and suggest potential solutions?"

Step 7 — Save or Copy Your Summary

Once you're happy with Claude's summary, you'll likely want to save or copy it for your own use. You can typically highlight the text of Claude's response using your mouse, then right-click and select "Copy." You can then paste this summary into a document, email, or wherever you need it. Some AI tools also offer specific "Copy" buttons next to their responses or even export options, though these vary.

If you don't see a direct "copy" button, highlighting the text and using your computer's standard copy (Ctrl+C on Windows, Command+C on Mac) and paste (Ctrl+V or Command+V) functions will always work.

You'll know it worked when you've successfully pasted the summary into another application or saved it to a file, making it ready for your tasks.

Common mistakes

- **Uploading unsupported files or files that are too large:** Claude works best with common document types like PDFs, Word documents, and text files. If you try to upload an image or a video, or a document that's many hundreds of pages long, it might not work.
  - **The fix:** Convert your file to a supported format like PDF or plain text. If it's very long, you might need to split it into smaller sections or use a paid plan with a larger context window.
- **Vague summary requests:** Just saying "Summarize this" might give you a general overview, but it might not be what you actually need. Claude is good, but it can't read your mind!
  - **The fix:** Be specific! Ask for "three key takeaways," "a summary for someone who knows nothing about this topic," or "the pros and cons of this proposal." The more detail you give in your **prompt**, the better the result.
- **Not verifying critical information:** While Claude is powerful, it's an AI and can sometimes **hallucinate**, which is when an AI confidently presents false or inaccurate information as if it were true. It generates text based on patterns, not absolute truth.
  - **The fix:** Always cross-check any critical facts, figures, or sensitive information from the summary with the original document or another reliable source, especially if you plan to use the summary for important decisions. Claude can draft an excellent summary, but it can't verify the truthfulness of *every single detail* for you, especially in complex or niche topics.

What this means for you

- **In everyday life:** Imagine quickly understanding the long terms and conditions for a new app, getting the gist of a lengthy news article, or grasping the main points of a new recipe or instruction manual without reading every single word. Claude can save you precious time.
- **For your business or work:** This skill is invaluable for professionals. You can rapidly review lengthy reports, condense meeting minutes, extract key findings from research papers, or quickly understand client briefs, allowing you to focus on decision-making rather than sifting through pages of text.
- **If you're just getting started:** Pick a news article you're interested in online. Copy the main text, paste it into a new Claude conversation, and ask, "Summarize this for me in three sentences." It's a quick, easy way to get a feel for how it works!

Try it now

Find a short PDF document (perhaps a product manual or a short article) or copy-paste the text from a recent online news story into Claude, and then ask it to "Give me the top three main points from this text." See how quickly you can get to the core information!

Turn Messy Notes into a Polite Meeting Agenda in Five Minutes

7 steps · Read online

> **In short:** Paste your messy email thread or bullet list into any AI assistant with a clear instruction about timing, owners, and tone, then iterate. You'll get a polite, structured agenda in under five minutes — though you'll still need to send the invites yourself.

A 30-minute meeting with seven bullet points, three owners, and zero structure usually ends with 20 minutes wasted on "wait, who's covering that?" — and another 10 on a follow-up email nobody reads. This guide shows how to feed that mess into any AI assistant and walk out with a timed, owner-clear agenda you can actually send. It's written for people who already use AI for everyday tasks and want one more workflow that saves real time at work.

Before you start

- An account on any AI assistant you already use — ChatGPT, Claude, Gemini, Copilot, Le Chat, or DeepSeek all handle this pattern. No new signup needed.
- The messy source material within reach: an email thread to copy from, a notes app with bullet points, or a transcribed voice memo.
- About five minutes once you've done it once; a little longer the first time as you build your reusable template.

Step 1 — Gather the messy material in one place

Open your AI assistant in a browser or app and start a fresh chat. Paste your raw material — the email thread, the bullet list, the voice-note transcript — straight into the message box, and add two short sentences of context before you hit send: what the meeting is for, how long it is, and who's coming. You're not asking for the agenda yet; you're staging the inputs.

Example: "Below is the email thread about our Q3 launch. Goal: 30-minute sync to unblock the design freeze. Attendees: Maria, John, Sarah."

You'll see your pasted text sitting in the input area, with the send button below it. After you send, the assistant confirms receipt and begins working. If your assistant has a "Projects" or "custom instructions" feature, you can save the context sentence there for next time — but a plain chat works fine without it.

You'll know it worked when the assistant's first reply acknowledges the inputs you provided and starts producing structured output rather than asking what you meant.

Step 2 — Ask for a first draft with explicit structure

Send a follow-up instruction in the same chat telling the assistant how to shape the draft: numbered topics, a duration in minutes next to each, a named owner per topic, a polite opener, and a small buffer between items. The clearer your structure request, the less you'll need to fix later.

Example: "Now produce the draft. Use a polite opener, list each topic with a duration in minutes and a single owner, leave a 2-minute buffer between topics, and end with a one-line confirmation request."

You'll see a numbered agenda appear, with minutes and names attached to each line. If your assistant formats the response as a table by default — some do — reply with "please rewrite as a numbered list with bold topic names" and the next version arrives in the shape you wanted.

You'll know it worked when your chaotic items are reordered into clear blocks with names and minutes, and you can read the whole meeting in one glance.

Step 3 — Calibrate the politeness

Reply in the same chat asking for tone adjustments. The first draft is almost never the right tone — it usually sounds either too stiff ("Please find below the agenda…") or too casual for a workplace you share with senior people. Tell the assistant which direction to move.

Example: "The opener feels stiff. Loosen it. Our team uses first names and avoids phrases like 'circling back'. Keep the same structure and timings."

You'll see a revised version where the wording changes but the structure stays. If the assistant doesn't pick up your tone signal, paste one sentence of how you'd actually open the meeting and ask it to match that register — that's the fastest shortcut to a natural voice.

You'll know it worked when the opener reads like something you'd say out loud without wincing.

Step 4 — Lock down the owners and timing

Send another short instruction asking the assistant to verify two things: each topic has exactly one named owner, and the minutes add up to the meeting length. This is where most agendas quietly fall apart — two topics assigned to the same person back-to-back, or a 45-minute agenda for a

30-minute slot.

Example: "Confirm each topic has exactly one owner. Add up the minutes and show the total at the end. If anything overflows, suggest what to cut."

You'll see a revised agenda where totals match and names sit beside each item. If the assistant doesn't add a "Total: X minutes" footer, ask for one — it's the simplest way to verify the math yourself before you copy anything out.

You'll know it worked when the total at the bottom equals your available time, and no item is missing a name.

Step 5 — Add a closing line with next steps

Ask the assistant to add a short "Next steps" block at the bottom: who sends the calendar invite, who brings what (slides, doc, numbers), and when the follow-up happens. This is the part most agendas skip, and it's the part that decides whether the meeting actually moves anything.

Example: "Add a 'Next steps' block at the end: Maria sends the calendar invite by tomorrow morning; John brings the latest design mockup; follow-up sync the following Monday."

You'll see a new block appear below the agenda. If your team uses a specific tool — Notion, Google Docs, Slack — name it in your prompt and the assistant will adjust the closing line to fit that workflow.

You'll know it worked when each person at the meeting knows exactly what to do before it starts and after it ends.

Step 6 — Copy the result into your real-world tool

Select the final agenda in the chat and paste it into the place your team will actually see it — an email draft, a shared doc, a calendar description, or your meeting software. This step is small but it matters: the moment the agenda lives outside the chat, you can edit it, share it, and update it without re-asking the AI.

If your assistant offers a "copy" or "export" button, use it; otherwise highlight the text and use your system copy shortcut. The text lands in the new tool with light formatting shifts depending on the destination — that's normal.

You'll know it worked when the agenda sits in the same place your team already works, and you didn't need to retype a single bullet.

Step 7 — Save your prompt as a reusable template

Highlight the full chat — your original messy input, the structure request, the tone instructions, the owners-and-timing check, the closing block — and save it as a reusable template. A plain text file in your notes app works. If your assistant supports saved prompts or custom instructions, use that instead so the template loads automatically.

Example: a doc titled "Agenda from messy notes" with three labeled blocks: [PASTE INPUT], [STRUCTURE REQUEST], [TONE NOTES].

Next time you have a chaotic thread, you paste the template, swap in the new content, and skip the prompt-writing step. The first time costs you ten minutes; every meeting after that costs you one.

You'll know it worked when the next agenda takes you under a minute to draft.

Common mistakes

- **Treating the first draft as final.** The first reply is a sketch. Always run one tone-and-owners pass — usually two — before you copy anything out.
- **Owners stay vague ("the team", "marketing").** Vague owners mean nobody follows up. One named person per topic, every time, no exceptions.
- **Forgetting the AI can't send the invite.** Once you copy the agenda out, the chat ends and the tool stops helping. You still need to send the calendar invite, share the doc, and chase the owners yourself.

What this means for you

- **In everyday life:** The same pattern tidies up a 10-minute prep before a school meeting, a volunteer group call, or a family planning session — paste the chaos, ask for structure, lock in who does what.
- **For your business or work:** Every recurring meeting — weekly stand-up, monthly review, project sync — gets a polished, owner-clear agenda in minutes instead of 30 minutes of wrestling with bullet points.
- **If you want to go further:** Try the same input-to-structure pattern on other meeting artefacts: a recap email after the meeting, a project kickoff brief from a Slack thread, or a status report from scattered updates. The skill is the same — stage the messy input, request the shape you want, iterate on tone and owners.

Try it now

Open any AI assistant, paste your messiest recent email thread, and send this single line: "Turn this into a 30-minute agenda with one named owner per topic and a polite opener." Watch what comes back — that's your starting point for the iteration loop above.

Build a simple review form for your small business with GitHub Copilot

7 steps · Read online

> **In short:** GitHub Copilot can build a "Leave a review" form (HTML, CSS, and JavaScript) in about 15 minutes inside a free browser editor. You'll still need a free service like Formspree to actually deliver reviews to your inbox — plain HTML alone can't save data on its own.

If you run a small shop, café, or local service, you already know how much a few honest reviews help. You probably also know that hiring a developer to build a "Leave a review" form can cost more than you'd like to spend. This guide walks you through having GitHub Copilot write that form for you — and explains, in plain words, what it can do and what it can't.

Before you start

- A free GitHub account (sign up at github.com — all you need is an email).
- A modern browser like Chrome, Edge, or Firefox.
- About 15–20 minutes of quiet time.
- A short list of 3 fields you want on your form (most review forms just need: customer name, star rating, and a comment box).
- Heads-up: GitHub Copilot offers both free and paid plans — check GitHub's pricing page for the current options before you start.

Step 1 — Open a free Codespace (your code editor in the cloud)

Sign in to github.com, then head to github.com/codespaces and click the green **"New codespace"** button. Pick the **"Blank"** template and choose a small machine (2-core is plenty for a tiny form). A new tab opens with a full code editor — it looks a lot like VS Code, but it runs in your browser. Wait about 30 seconds for it to load. If the Codespaces button is hard to find, create a new empty repository first (click the **"+"** in the top-right, then **"New repository"**), then click the green **"Code"** button and pick the **"Codespaces"** tab. You'll know it worked when you see a dark editor screen with a file panel on the left and a small terminal panel at the bottom. 💬 Example: name your repository `review-form` so it's easy to find later.

Step 2 — Make a new file for your form

In the left-hand file panel, look for an icon that looks like a piece of paper with a folded corner — that's the **"New File"** button. Click it, type `index.html` as the file name, and press Enter. A blank white editor opens in the middle of the screen — that's where your form will live. If that button isn't where you expect, right-click anywhere in the empty file panel and choose **"New File"** from the menu that pops up. You'll know it worked when `index.html` shows up in the left panel and the editor in the middle is empty, with a line number 1 on the left edge.

Step 3 — Ask Copilot to build the form for you

Look for the **Copilot Chat icon** in the left sidebar — it usually looks like a small speech bubble or a sparkle ✨. Click it. A chat panel slides open on the right. In the box at the bottom, type your request in plain English. Be specific about the three fields you want. If you don't see a chat icon, try the keyboard shortcut **Ctrl + Shift + I** on Windows/Linux, or **Cmd + Shift + I** on Mac — that opens the same chat panel. 💬 Example: *"Write a clean HTML form for a small bakery called 'Sunrise Sourdough'. Ask for the customer's name, a star rating from 1 to 5, and a comment box. Include a nice 'Submit review' button."* You'll know it worked when Copilot replies with a block of HTML code in the chat, and a small **"Apply"** or **"Copy"** button sits next to it.

Step 4 — Paste the code into your file and save it

Click into your empty `index.html` file so the cursor is inside it. Then press **Ctrl + V** (or **Cmd + V** on Mac) to paste the code from the chat. If a small popup asks you to confirm, click **"Accept"**. Save the file with **Ctrl + S** (or **Cmd + S**). If your keyboard shortcuts don't match, look at the top menu bar for **File → Save** instead — every code editor has that. You'll know it worked when the file no longer has a dot next to its name in the left panel (dots mean unsaved changes). The tab at the top should now read `index.html` with no little asterisk. 💬 Example: if Copilot's reply was a chat message instead of a code block, copy it from the chat, then paste it into the file the same way.

Step 5 — Ask Copilot to style it

Now make it look nice. Go back to the Copilot Chat panel and type a follow-up like the example below. Copilot replies with CSS — the part that controls colors, spacing, and fonts. Click the **"Apply"** button if it appears; otherwise click **"Copy"** and paste the CSS at the top of your file just after the opening `<head>` tag. If "Apply" isn't visible, look for a small **"Insert at cursor"** link near the code block — most editors offer one. You'll know it worked when your file has new lines starting with `<style>` and ending with `</style>`, and the code inside has words like `color`, `padding`, or `border-radius`. 💬 Example: *"Now add CSS that makes the form look warm and friendly — soft cream background, rounded corners, and a brown 'Submit' button. Keep it simple."*

Step 6 — Make it actually deliver the review to you

A form sitting on a page by itself can collect information in the browser, but it can't email you the results without a small extra step. The honest thing here: Copilot can write the front-end (what visitors see), but to receive reviews in your inbox you'll need a free form service. **Formspree** has a free tier that works well — sign up, get a form endpoint (a long web address), then ask Copilot to point your form at it. In the chat, type something like the example. Copilot will rewrite the `<form>` tag so the data goes to Formspree, which then forwards it to your email. If Formspree's setup looks confusing, ask Copilot instead: *"Change the submit button so it opens the user's email app with the review already filled in."* That quick shortcut uses a `mailto:` link and works without any third-party service. You'll know it worked when you click **Submit** in the preview and either (a) Formspree shows a thank-you page, or (b) your email app opens with the review text ready to send. 💬 Example: *"Update the form so it submits to Formspree. Use this endpoint: https://formspree.io/f/your-code-here, and change the button text to 'Send my review'."*

Step 7 — Preview it and tweak what you don't like

Right-click your `index.html` file in the left panel and pick **"Open in browser"** (or look near the top of the file's tab for a small browser icon). A new tab opens showing your form. Try filling it out and submitting to see how it behaves. If "Open in browser" isn't there, look at the bottom-right corner of the editor for a small **"Go Live"** or **"Live Preview"** icon — Codespaces sometimes hides it there. You'll know it worked when your form looks decent on screen, the fields are easy to click, and submitting it doesn't show an error. If something looks off, go back to Copilot Chat and describe the problem in plain words — for example: *"The submit button is too small on phones — make it bigger."*

Common mistakes

- **Pasting the code back into the chat instead of the file.** Many beginners paste Copilot's reply into the chat again, creating an endless loop. The fix: click into the empty editor area before pressing paste, then save with Ctrl+S.
- **Forgetting to save.** Codespaces often auto-saves, but not always. The fix: glance at the file tab — if there's a dot or asterisk next to the filename, it isn't saved yet. Press Ctrl+S to be sure.
- **Trusting every line Copilot writes.** Copilot is helpful but it makes mistakes — it might invent a field name, miss a closing tag, or use outdated HTML. The fix: read the suggested code, run the preview, and ask Copilot to fix anything weird instead of just pasting and publishing.

What this means for you

- **In everyday life:** you can swap the example prompt for any small thing on your site — a contact form, a "Book a table" button, an FAQ section — and Copilot does the same job. Once you've done this once, the second form takes half the time.
- **For your business or work:** a simple review form you control (no monthly fee to a third-party tool) is a small but real upgrade. You own the code, you can change it any time, and you can hand it to a developer later if you outgrow it.
- **If you're just getting started:** the easiest first move is exactly what Step 1 says — open a Codespace, name it `review-form`, and type your prompt into the chat. Twenty minutes from now you'll have a working page.

Try it now

Open github.com/codespaces, start a blank Codespace, and paste this exact prompt into Copilot Chat: *"Make me a clean, friendly HTML page with a 'Leave a review' form. Fields: name, star rating 1–5, and a comment box. Add a 'Submit review' button and simple warm styling."* Save the file, hit preview, and see what Copilot hands you in under a minute.

How to Clean Up a Flooded Email Inbox Using AI in One Afternoon

7 steps · Read online

> **In short:** You can use any AI assistant (ChatGPT, Claude, Gemini, or Copilot) to group your unread email into folders and write a one-paragraph summary of what needs attention today. Most people finish in under three hours.

By the end of this guide, you'll have a tidy inbox with folders that make sense, plus a single short paragraph at the top telling you what actually matters right now. The AI does the thinking; you do the clicking.

Before you start

- An account with any AI assistant you can already use — ChatGPT, Claude, Gemini, Copilot, or Le Chat all work. The free version is enough.
- Your email account (Gmail, Outlook, Yahoo, Apple Mail, anything) open in a web browser.
- About 2–3 hours of quiet time — most of it is waiting, not typing.
- A notes app or blank document where you can paste text (Notepad, Google Docs, Notes — anything that saves text).

Step 1 — Pick which AI you'll use, and open it in a new tab

Open the AI assistant of your choice in a browser tab and sign in. You'll see a chat box at the bottom of the screen with a blank text field waiting for your first message. The conversation is private to you — nothing you paste is sent back to your email provider.

If you've never used one before, ChatGPT, Claude, or Gemini are all safe picks. The free version is plenty for this task; paid plans just let you paste longer messages. If you can't find your bookmark, search "[tool name] login" in Google — the first result is almost always the real site.

Example: open chatgpt.com, claude.ai, gemini.google.com, or copilot.microsoft.com.

You'll know it worked when you see a typing box with a "Send" arrow or paper-plane icon, and the AI greets you when you say hello.

Step 2 — Make a copy of your unread emails the AI can read

In your email account, search for `is:unread` (Gmail), `unread:true` (Outlook web), or use the "Unread" filter in the left sidebar. The screen fills with every message you've ignored — that's your starting pile.

Pick the most recent 50–100 unread emails. Copy the subject lines and senders into a plain text document with `Ctrl+C` (Windows) or `Cmd+C` (Mac). If your email client offers "Export to .csv" for the unread folder, that's even faster — but only export unread mail, never your whole account.

Example: in Gmail, tick the checkbox at the top of the inbox. If a "Select all conversations" link appears, click it. Open the ⋯ menu and choose "Export."

You'll know it worked when you have a file with rows like "Sender — Subject — Date — first 2 lines."

Step 3 — Ask the AI to group the emails into 5–7 folders

Switch back to your AI tab and paste the list in. Then ask it to suggest a folder structure — copy the prompt below as-is.

The AI will respond with a small list of folder names (5–7 is the sweet spot — too many and you stop using them) and which emails belong where. It will not move anything itself; it only suggests. That part is still on you.

Example prompt: *"Here is a list of my unread emails. Group them into 5–7 folder names that make sense, and tell me which emails go in which folder. Use plain folder names a normal person would understand. Don't add explanations yet — just the list."*

You'll know it worked when you see something like: "Personal & Family — 12 emails. Bills & Receipts — 8 emails. Work — 27 emails. Newsletters — 41 emails. Calendar & RSVPs — 6 emails."

Step 4 — Create those folders in your email client

In your email account, create the folders the AI suggested. Most email apps have a "+" or "Create new label/folder" button in the left sidebar — if you don't see it, right-click an existing folder name and look for "New label," "New folder," or "Create."

Name them exactly as the AI wrote them, with the same spelling and capitalization. Lowercase, plural, no emoji — keeping it boring makes it stick. If the AI suggested something like "WorkStuff" and you'd prefer "Work," rename it now. The point is a structure you'll actually use.

Example: in Gmail, look in the left sidebar, scroll to "Labels," click the "+" next to it, and type "Bills & Receipts." In Outlook web, right-click "Folders" in the sidebar and choose "New folder."

You'll know it worked when every new folder name shows up in the sidebar, ready to receive mail.

Step 5 — Move messages into folders, starting with the small groups

Go back to your inbox and tackle the smallest group first — usually "Calendar & RSVPs" or "Personal." Open each email, click "Move to," and pick the matching folder. Do not read anything. The goal is an empty inbox, not total comprehension.

For the biggest group (often "Newsletters"), speed this up: search `category:promotions` (Gmail), `category:updates`, or use the bulk-select checkbox at the top of the list, then choose "Move to" → "Newsletters." This is the only step that feels slow — the rest is fast.

Example: type `label:unread from:newsletter` in the search box, select all results, move them to your "Newsletters" folder in one click.

You'll know it worked when your main inbox counter shows a much smaller number than before, and the moved messages no longer appear in the inbox list.

Step 6 — Ask the AI to write your "what matters now" summary

Now paste the remaining urgent emails (the ones still in your inbox, or the Work and Personal folders) back into the AI and ask it to write a short paragraph telling you what actually needs attention today.

The AI is good at this — it sorts through noise, spots deadlines and asks, and skips pleasantries. Your job is to give it enough context: who is waiting on you, what is due this week, what can wait a month.

Example prompt: *"Here are the remaining unread emails. Write me one short paragraph (under 100 words) telling me what I actually need to act on today, what's due this week, and what I can ignore. Start with the most urgent."*

You'll know it worked when the AI returns a single paragraph beginning with the deadline or action that's most urgent, like: "Reply to the school about Thursday pickup by tomorrow morning — they need a yes or no. Confirm the dentist appointment on Friday. Ignore the three shipping notifications..."

Step 7 — Pin that summary somewhere you'll see it

Take the paragraph the AI wrote and put it where your eyes land first in the morning. The easiest spot: the very top of your inbox (most email apps let you "Pin" or "Star" a message), or a sticky note on your desktop, or the first line of today's to-do list.

This is the whole point of the exercise. A clean inbox feels good for ten minutes — a clear "what matters now" keeps you focused all week. The folders are just a parking lot; the paragraph is the compass.

Example: copy the paragraph into a new email titled "Today" addressed to yourself, and pin it at the top of your inbox.

You'll know it worked when tomorrow morning, before you open any other email, you can read that one paragraph and know what to do first.

Common mistakes

- **Exporting your whole email account instead of just unread.** The AI does not need ten years of receipts to group your unread mail — and pasting your entire inbox into a chatbot is risky. Stick to unread, and only the last few hundred.
- **Letting the AI invent clever folder names like "CriticalPath_Q3_Items."** If a folder name is not something you'd write on a sticky note, change it. Boring folder names get used; clever ones get ignored.
- **Reading every email as you move it.** That's not cleanup, that's procrastination wearing a costume. Move first, read later — the summary tells you what actually matters.

What this means for you

- **In everyday life:** you'll stop dreading the inbox icon. The folder structure is a parking lot for things you don't need to act on; the summary is your morning compass.
- **For your work or business:** if you get a lot of client email, do this once a month. The folder names might shift slightly (a new project, a new client), but the pattern stays the same.
- **If you're just getting started:** pick one AI you already trust, do Steps 1–4 today, and save 5–7 for tomorrow. Two shorter sessions beat one exhausting one.

Try it now

Open your email account, search for unread messages, and copy the first 30 subject lines into your AI chat. Ask it for 5 folder names. That's it — ten minutes of work, and you'll already see the shape of the cleanup.

Check your resume for ATS-friendly keywords using Hugging Face's Inference API

7 steps · Read online

You'll learn how to scan your resume for keywords that applicant tracking systems (ATS) love, using free AI tools. This helps your resume get past the first filter and into a human's hands.

Before you start

- A **free Hugging Face account** (sign up here)
- A **resume in PDF or Word format** (saved on your device)
- A **computer or laptop** (phone works, but a larger screen is easier)
- **No coding skills needed** — this guide is for everyone
- **Total time:** about 10 minutes

💡 *What is ATS?* Think of it as a robot that scans your resume before a human does. It looks for keywords from the job ad. If your resume doesn't have enough of those words, it might get rejected before you even know.

Step 1 — Create a free Hugging Face account

Go to huggingface.co and click **Sign Up** in the top-right corner. Fill in your email, choose a username and password, then click **Create account**. Check your email and confirm your account.

You'll see your Hugging Face profile page. This is your dashboard — your home base for using AI models.

💬 Example:

```
Email: yourname@example.com
Username: resumehelper
Password: MyPassword123!
```

You'll know it worked when you see your profile picture and username in the top-right corner.

Step 2 — Open the Inference API page

On the Hugging Face homepage, look for **Products** in the top menu. Hover over it and click **Inference API**. This is a free tool that lets you run AI models without installing anything.

You'll see a page with a text box and a list of AI models. This is where you'll paste your job ad and resume.

💡 *What is an API?* Think of it like a waiter in a restaurant. You tell the waiter what you want (your job ad and resume), and the waiter brings it to the kitchen (the AI model) and brings back the answer — without you needing to cook.

You'll know it worked when you see a text box that says **Enter your prompt here** and a button that says **Submit**.

Step 3 — Choose an ATS-friendly keyword finder model

In the text box, type or paste this model name exactly:

```
microsoft/Phi-3-mini-4k-instruct
```

Then press **Enter** or click **Submit**.

This model is a smart AI assistant that can read your job ad and resume and tell you which keywords are missing.

💡 *What is a model?* Think of it like a super-smart student who has read thousands of job ads and resumes. It knows what words to look for.

You'll know it worked when you see a chat window open with a welcome message from the AI.

Step 4 — Write your job ad and resume into the chat

In the chat box, type or paste this exact prompt (copy the whole block):

```
I need help finding ATS-friendly keywords for a job application.
Here is the job ad:
```

```
[PASTE JOB AD HERE]

Here is my resume:
[PASTE RESUME TEXT HERE]

List the top 10 keywords from the job ad that are missing from my resume.
```

Replace `[PASTE JOB AD HERE]` with the full text of the job ad you're applying for. Replace `[PASTE RESUME TEXT HERE]` with the text of your resume. To copy text from a PDF, open the PDF, select the text, and press **Ctrl+C** (Windows) or **Cmd+C** (Mac). Paste it into the chat with **Ctrl+V** or **Cmd+V**.

💬 Example:

```
I need help finding ATS-friendly keywords for a job application.
Here is the job ad:
We are looking for a digital marketing specialist with experience in SEO, Google Analytics, and content creation.

Here is my resume:
I have 3 years of experience in digital marketing, including social media management and email campaigns.
```

You'll know it worked when the AI replies with a list of keywords, like:

> The top 10 keywords missing from your resume are: SEO, Google Analytics, content creation, PPC, keyword research, on-page optimization, conversion rate optimization, Google Ads, Ahrefs, backlink building.

Step 5 — Ask the AI to rewrite a sentence with the missing keywords

Copy one of the missing keywords from the AI's reply. Then, in the same chat, type:

```
Rewrite this sentence from my resume to include the keyword [KEYWORD]:
"I have 3 years of experience in digital marketing, including social media management and email campaigns."
```

Replace `[KEYWORD]` with the keyword you copied, like **SEO** or **Google Analytics**.

💬 Example:

```
Rewrite this sentence from my resume to include the keyword SEO:
"I have 3 years of experience in digital marketing, including social media management and email campaigns."
```

You'll know it worked when the AI replies with a rewritten sentence that includes the keyword naturally, like:

> I have 3 years of experience in digital marketing, including SEO strategy, social media management, and email campaigns.

Step 6 — Save your updated resume

Copy the rewritten sentence from the AI's reply. Open your resume file (Word or Google Docs), find the sentence you want to update, and paste the new version over it. Save the file with a new name like `Resume_ATS_Updated.pdf`.

💡 *Why save a new file?* This way, you always have your original resume to go back to if you need to.

You'll know it worked when you have a new file saved with your updated resume text.

Step 7 — Run the process again to double-check

Repeat Step 4 with your updated resume. Paste the same job ad and your new resume text into the chat. Ask the AI:

```
List the top 10 keywords from the job ad that are still missing from my resume.
```

You'll know it worked when the AI replies with a shorter list of missing keywords, or says:

> All the top keywords from the job ad are now included in your resume.

Common mistakes

- **Mistake:** Pasting the job ad or resume as an image or screenshot. **Fix:** Copy the text from the PDF or Word file and paste it as plain text. Images can't be read by the AI.
- **Mistake:** Using a phone to copy and paste large blocks of text. **Fix:** On a phone, it's easy to lose your place. Use a computer or laptop if you can. If you must use a phone, paste smaller chunks of text at a time.
- **Mistake:** Expecting the AI to rewrite your whole resume perfectly in one go. **Fix:** The AI is great at spotting missing keywords, but it can't rewrite your whole resume in one sentence. Update one sentence at a time, then check again.

What this means for you

- **In everyday life:** You'll spend less time guessing which keywords to add to your resume. The AI does the hard work for you, so you can focus on the rest of your job application.
- **For your business or work:** If you're hiring, you can use this same method to check if your job ads include the right keywords. It helps attract candidates whose resumes will pass the ATS filter.
- **If you're just getting started:** This is the easiest first move to make your resume ATS-friendly. No coding, no expensive tools — just free AI and a few minutes of your time.

Try it now

Open your resume file and copy the first sentence. Go to huggingface.co, open the Inference API, and paste this prompt:

```
I need help finding ATS-friendly keywords for a job application.
Here is the job ad:
[PASTE A JOB AD YOU'RE APPLYING FOR]

Here is my resume:
[PASTE THE FIRST SENTENCE OF YOUR RESUME]

List the top 5 keywords from the job ad that are missing from my resume.
```

Press **Submit** and see what the AI suggests. You'll have your first ATS keyword list in under 2 minutes.

Write a Polite Complaint Email to a Company with DeepSeek

7 steps · Read online

> **In short:** DeepSeek can draft a polite, well-structured complaint email in under a minute — but you'll need to add the real details (order numbers, dates, names) and send it from your own account. It can't contact the company for you.

When something goes wrong with a product or service, the email you write in frustration usually gets ignored. A short, calm message that says what happened, what you want, and by when is far more likely to get a real reply. DeepSeek is great at shaping that kind of message — you describe the problem in your own words, and it turns your rambling into a clean, firm but polite email.

Before you start

- A phone, tablet, or computer with internet access.
- A web browser (Chrome, Safari, Edge, Firefox — any modern one). Or the DeepSeek app if you prefer.
- A rough idea of what went wrong: when it happened, what you bought, and what you'd like the company to do about it.
- Any order numbers, receipts, or screenshots handy (helpful, but not required).
- About 10 to 15 minutes total.

Step 1 — Open DeepSeek in your browser

Go to **chat.deepseek.com** in your web browser. The page loads a chat screen with a text box near the bottom labeled something like "Send a message" or "Message DeepSeek." No account is needed for a basic chat, so you can start typing right away. If you have the DeepSeek app installed, open it instead and you'll see the same chat layout. Example: type nothing yet — just confirm the box is there. You'll know it worked when the chat screen is empty and ready, with a blinking cursor in the message box.

Step 2 — Describe what happened in plain words

In the chat box, write a short, honest summary of the problem. Don't worry about grammar or politeness — just dump the facts the way you'd tell a friend. Include the company name, the product or service, when it happened, and what went wrong. DeepSeek reads everything in this message and uses it to draft the email later. Example: "I bought a coffee maker from BrewCo on March 12, order #44291. It arrived broken — the carafe was cracked and water leaks everywhere. I emailed them a week ago and nobody replied. I want a replacement or a full refund." You'll know it worked when DeepSeek replies by asking a follow-up question or jumping straight into a draft. If it asks "What outcome do you want?", move to Step 3.

Step 3 — Ask for a polite complaint email

Type a clear follow-up: "Please turn this into a polite complaint email I can send to the company. Keep it firm but respectful, around 150 words, with a subject line." DeepSeek now produces a full email — greeting, body paragraphs, and a sign-off — usually in a couple of seconds. The subject line usually appears at the top. Example: "Please turn this into a polite complaint email I can send to BrewCo customer service. Keep it firm but respectful, around 150 words, and give me a subject line too." You'll know it worked when you see a complete email in the chat, with a "Subject:" line, a "Dear..." greeting, and a closing like "Sincerely."

Step 4 — Tell DeepSeek what outcome you want

The first draft is usually generic. Tell DeepSeek specifically what you want the company to do — a refund, a replacement, an apology, a callback, a coupon. The clearer you are, the more useful the next draft becomes. You can also mention a reasonable deadline, like "within 7 business days." Example: "Please rewrite it so I clearly ask for a full refund of $89 within 7 business days, and mention that if I don't hear back I'll file a chargeback with my bank. Keep the same polite tone." You'll know it worked when the new version includes your specific ask, in plain language, not a vague "I hope you can help."

Step 5 — Adjust the tone if needed

Read the email out loud. Does it sound like you? If it's too stiff, say "make it sound warmer and more conversational." If it's too soft and you want to sound firmer, say "make the tone more direct without being rude." DeepSeek rewrites the whole email in seconds, so don't be shy about asking two or three times. Example: "Make it a bit warmer — I'm not angry, just disappointed. Keep the firmness but add one sentence saying I have liked their products in the past." You'll know it worked when the new version feels like something you'd actually be comfortable sending under your own name.

Step 6 — Add your real details and read it carefully

Now do the human part. Open the email in your own mail app (Gmail, Outlook, Apple Mail — anything), and copy the draft over. Fill in the missing pieces: the actual recipient email (usually support@example.com or a contact form address), the real order number, dates, your name, and any return or tracking numbers. Read it once more to make sure it sounds like you and every fact is correct. Example: replace "[Your

Name]" with your real name, "[Order #]" with the real order number, and double-check the dollar amounts. You'll know it worked when the email reads as a complete, accurate message with no placeholder text like "[insert date here]."

Step 7 — Send the email from your own account

Hit send from your own email account. DeepSeek can't send emails for you — it's a chat tool, not a mail program. After sending, keep the message handy; if you don't hear back in the time you mentioned (often 5 to 7 business days), you can simply reply to your own thread with "Just following up on my email below." Example: write a short follow-up — "Hi, I'm following up on my email from March 20 about order #44291. I haven't received a reply yet. Please let me know how you'd like to resolve this." You'll know it worked when the email lands in your "Sent" folder and you have a copy you can quote later if needed.

Common mistakes

- **Being too angry in the prompt.** DeepSeek often mirrors the tone of what you give it. If you rant, you may get a rant back. Stay matter-of-fact and let the tool do the polishing.
- **Forgetting the specific ask.** A complaint that doesn't say "refund, replacement, or reply by Friday" often gets a generic "we're sorry" reply. Always name what you want.
- **Sending without reading.** AI drafts sometimes invent details — a wrong date, a fake policy, a tone that isn't yours. Always read it once, top to bottom, before sending.
- **No receipt of what you sent.** Screenshot or forward the final email to yourself so you have proof of what was said and when.

What this means for you

- **In everyday life:** A broken toaster, a missed delivery, a noisy hotel room, a wrong phone bill — any of these can become a 10-minute fix instead of a week of ignored messages.
- **For your business or work:** If you invoice or buy from other companies, the same trick works for late payments, missing deliveries, or poor service. A polite, specific email usually beats a phone call.
- **If you're just getting started:** Pick the smallest, most recent annoyance and try one email. You'll see how much easier it is when you don't have to start from a blank page.

Try it now

Think of one small frustration from the last month — a delayed package, a wrong charge, a rude reply from a store. Open DeepSeek, type two sentences about what happened, and ask it to turn that into a polite email. Read the draft, fix one detail, and save it in your drafts folder. You'll have a real, ready-to-send message in under five minutes.

Drafting Professional Presentations from Scratch Using Microsoft Copilot in PowerPoint

7 steps · Read online

Imagine turning a messy idea or a blank slide deck into a polished, ready-to-use presentation in under 10 minutes. With Microsoft Copilot in PowerPoint, you can draft outlines, generate slide titles, and fill in speaker notes—all from a simple prompt. This guide walks you through every step, from opening PowerPoint to exporting your finished slides.

Before you start

- A **Microsoft 365 account** (free or paid)
- The **PowerPoint desktop app** (Windows or Mac) or **PowerPoint on the web**
- A **device with internet access**
- Roughly **10–15 minutes** of your time
- Optional: a **Microsoft 365 Copilot license** (paid) for full features

> ⚠️ Note: Some Copilot features may require a paid Microsoft 365 Copilot license. If you don't have it, you'll still be able to use basic Copilot features like generating outlines and slide titles.

Step 1 — Open PowerPoint and start a new presentation

Open the **PowerPoint app** on your computer or go to PowerPoint on the web. Click **New blank presentation** to start with a blank slide deck.

You'll see a single slide titled "Click to add title." The ribbon at the top shows tabs like **Home**, **Insert**, and **Design**.

If you don't see the **Home** tab, look for a menu at the top with words like **File**, **Home**, or **View**.

You'll know it worked when you see a blank slide with a title and subtitle placeholder.

Step 2 — Open Copilot in the PowerPoint ribbon

On the **Home** tab in the ribbon, look for a button labeled **Copilot** (it looks like a sparkle or a chat bubble with stars). Click it to open the Copilot pane on the right side of your screen.

You'll see a chat box that says "Ask me anything about this document." This is where you'll type your instructions.

If you don't see the Copilot button, make sure you're using the latest version of PowerPoint. Look for **File > Account > Update Options > Update Now**.

You'll know it worked when the Copilot pane opens on the right with a blinking cursor ready for your prompt.

💬 Example: Type or paste:

```
Help me create a presentation about the benefits of remote work for a team of 20 software developers.
```

Step 3 — Give Copilot a clear prompt to generate an outline

In the Copilot chat box, type a clear, specific prompt. Include:

- Your **topic**
- Your **audience**
- Any **key points** you want covered

After you press Enter, Copilot will process your request and show a message like "Generating outline..." in the pane.

You'll see a new section titled **Outline** appear in the Copilot pane, with 4–6 suggested slide titles and bullet points.

If the outline isn't quite right, you can type:

```
Add a slide about work-life balance and another about team collaboration tools.
```

You'll know it worked when you see a structured outline with slide titles and bullet points in the Copilot pane.

💬 Example:

```
Create an outline for a 10-minute presentation on 'How Remote Work Boosts Developer Productivity'. Include slides on flexibility, tools, and
```

Step 4 — Insert the outline into your slides

Click the **Add to presentation** button under the outline in the Copilot pane. Copilot will automatically create new slides based on the outline.

You'll see your slide deck grow from one blank slide to 5–7 slides with titles and bullet points.

If the slides look too plain, don't worry—you can edit them later. If the Copilot pane disappears, click the **Copilot** button again to reopen it.

You'll know it worked when you see multiple slides with titles and bullet points in your deck.

Step 5 — Ask Copilot to generate speaker notes for each slide

Select the first slide. In the Copilot pane, type:

```
Add speaker notes to this slide explaining the benefits of remote work for developers.
```

Copilot will write 2–3 sentences of speaker notes under the slide. Repeat this for each slide by selecting it and asking Copilot to add notes.

You'll see a small note icon appear on each slide. Click the icon to read the notes.

If speaker notes don't appear, make sure you've selected the slide first. If Copilot doesn't respond, refresh the pane or check your internet connection.

You'll know it worked when you see speaker notes appear under each slide.

Example: For a slide titled "Flexibility," Copilot might add:

```
Speaker notes: Remote work lets developers choose their most productive hours. Studies show 60% of developers report higher focus at home. M
```

Step 6 — Refine content with follow-up prompts

With the outline and notes in place, ask Copilot to improve or expand any slide. For example:

```
Rewrite the slide titled 'Team Collaboration Tools' to include specific examples like Microsoft Teams and GitHub.
```

Copilot will update the slide content in real time. You can also ask it to:

- Shorten the text
- Add a quote
- Include a data point

You'll see the slide content change immediately after you press Enter.

If Copilot doesn't update the slide, try clicking **Refresh** in the Copilot pane or re-selecting the slide.

You'll know it worked when the slide content updates to match your new prompt.

Example:

```
Add a quote from Satya Nadella about remote work culture to the last slide.
```

Step 7 — Export or present your presentation

When you're happy with the slides, save your presentation (**File > Save As**). To present, click **Slide Show > From Current Slide** or press **F5**.

You'll see your slides full-screen, ready to present.

If you want to share it, go to **File > Share > Send a Copy** to export as PDF or PowerPoint file.

You'll know it worked when you can present or share a complete, professional-looking deck.

Common mistakes

- **Mistake:** Writing a vague prompt like "Make a presentation." **Fix:** Be specific: include topic, audience, and key points. Example: "Create a 5-slide presentation on AI tools for small business owners."

- **Mistake:** Expecting Copilot to write perfect content in one try. **Fix:** Use follow-up prompts to refine slides. Ask Copilot to shorten, add data, or include a quote.
- **Mistake:** Not checking speaker notes or slide content for accuracy. **Fix:** Always review Copilot's output. Edit or remove anything that doesn't fit your message.

What this means for you

- **In everyday life:** Need to present a project update at work? Use Copilot to draft slides in minutes instead of hours. Great for team meetings, client pitches, or internal reviews.
- **For your business or work:** Create consistent, professional presentations without hiring a designer. Use it for quarterly reviews, training decks, or onboarding materials.
- **If you're just getting started:** Start with a simple prompt and let Copilot do the heavy lifting. You'll save time and look polished—even if you're not a designer.

Try it now

Open PowerPoint, start a new presentation, and open Copilot. Type:

```
Help me create a 5-slide presentation about the benefits of remote work for a team of 20 software developers.
```

Press Enter. Review the outline, add it to your slides, and start presenting in under 10 minutes.

Digitize Handwritten Notes into Editable Text with AI

6 steps · Read online

Do you have notebooks full of brilliant ideas, meeting notes, or shopping lists that you can never quite find when you need them? This guide will show you how to quickly turn those physical scribbles into digital text using AI, so they're always searchable and editable, right on your phone or computer. It's for anyone who wants to bring order to their paper notes without retyping a single word. While AI is great at "reading" your handwriting, it can't understand *meaning* or summarize for you yet – it just converts the words.

Before you start

- **A smartphone or tablet with a working camera.** Ensure your device has enough battery life and storage for a few photos.
- **Access to an app that offers AI-powered OCR.** This means **Optical Character Recognition**, a technology that "reads" text from images. Think of it like a very clever digital assistant looking at your picture and identifying each letter and word, then turning it into text you can actually edit. The "AI-powered" part means it uses advanced artificial intelligence to understand different handwritings and layouts, making it much more accurate than older systems. Many free note-taking or scanning apps now include this feature, such as Google Keep (available on Android and iOS) or Microsoft OneNote (available on Android and iOS).
- **A free Google or Microsoft account.** You'll likely need one of these to use the apps suggested, though you probably already have one set up on your device. Ensure you're logged in before you start.
- **Roughly 5-10 minutes** to follow these steps for your first digitized note.

Step 1 — Prepare your notes for capture

Before you take a photo, you need to set up your handwritten notes in a way that helps the AI "read" them clearly. You'll see your physical notes lying on a surface. Your goal is to give the AI the best possible image to work with, as good, even lighting and legible handwriting make a big difference for the AI to understand your scribbles correctly. Think of it like a librarian trying to read tiny, smudged print – the clearer, the better!

Once prepared, your notes will be ready for you to take a photo. If your usual handwriting is quite messy, try to write a little slower and clearer for important notes you plan to digitize, as this significantly improves the AI's accuracy.

Example: Place your handwritten notes on a plain, uncluttered surface, like a clear desk. Position them under natural light from a window, if possible, or use a well-lit lamp to avoid casting shadows over your writing. Make sure the paper is flat and smooth, with no creases. You'll know it worked when your notes are lying flat, look easy to read in the available light, and are clear of any distracting objects, ready for a photo.

Step 2 — Open your scanning app

Now it's time to open your chosen app. You'll see your device's home screen with various app icons. You need to locate and tap on the icon for your OCR app. For this guide, we'll use Google Keep as the primary example, which is free and widely available. Microsoft OneNote is another option that includes document scanning. Once opened, the app will load, often showing your existing notes or presenting a default view.

If you don't already have one of these apps, you'll need to download it from your device's app store (Google Play Store for Android or Apple App Store for iOS) and install it first. Look for the app icon, which usually looks like a colorful note for Google Keep or a blue notebook icon for Microsoft OneNote.

Example: Tap the "Google Keep" app icon on your home screen or app drawer (it looks like a yellow and white notepad). Alternatively, tap the "Microsoft OneNote" app icon (a blue notebook icon). You'll know it worked when the app opens and you see its main interface, such as a list of notes or a camera view, indicating it's ready for you to take action.

Step 3 — Take a clear photo

With your chosen app open, you'll typically see options to create a new note or use the camera. You need to find and tap the camera or scan icon to activate your device's camera within the app. The screen will then switch to a camera interface, often with specific modes like "Document" or "Scan" in apps like OneNote to help you capture text accurately.

Position your device directly above your notes, aiming for a straight-on shot without shadows, ensuring all your writing fits within the frame. Then, tap the shutter button. The app will capture the image, and some might even automatically detect the edges of your document and crop the image for you, making it look neater.

Example: In Google Keep, tap the plus (+) icon at the bottom of the screen, then choose the camera icon (it might be labeled 'Take photo' or 'Image'). In Microsoft OneNote, select 'Camera' from the options at the bottom, then hold your phone steady over your notes and tap the shutter button. You'll know it worked when you see a sharp, unblurred image of your notes presented within the app, often with the document edges automatically detected and cropped, ready for the next step.

Step 4 — Initiate text recognition with AI

With your clear photo captured in the app, it's time for the AI-powered OCR to do its magic. You'll typically see your captured image on the screen, possibly with a menu or a button related to text. You need to find and tap an option that tells the app to "read" the text from the image. This button might say something like "Scan text," "Recognize text," "Grab image text," or simply show a 'T' icon.

Once you activate this function, the AI will process the image you've taken, which usually takes just a few seconds. A new screen or panel will then appear, displaying your handwritten words now typed out as digital text. If you don't see an obvious "Grab image text" button, try tapping on the image itself, or look for a three-dot menu (⋯) or a gear icon (⚙) that might contain the option you need.

Example: In Google Keep, after your photo is captured, tap on the image to view it full screen. Then, look for a three-dot menu (⋯) in the top-right corner of the screen and tap the option labeled 'Grab image text'. In Microsoft OneNote, after taking a photo, tap the 'Text' icon that appears on the captured image to extract the text. You'll know it worked when a new text box or a separate screen appears, displaying your handwritten words now typed out as digital text, ready for you to review.

Step 5 — Review and make corrections

The AI is incredibly smart, but it's not perfect! Especially with unique handwriting, faint lines, or unusual symbols, it might make a few mistakes. You'll see the extracted text, usually in an editable field, displayed on your screen. Take a moment to read through this digitized text carefully and correct any errors.

As you make corrections, the text on your screen will update instantly, just like in any word processor. If the text isn't directly editable, look for an "Edit" button or a cursor you can tap to activate the editing mode. Remember, the AI is a great starting point, but it usually cannot perfectly capture every nuance of human handwriting, so your review is crucial for accuracy and ensures your digital notes are truly useful.

Example: Carefully read through the digitized text. If you spot a word like "flow" incorrectly appearing as "flew," tap on "flew" and type "flow" to correct it. Also, remove any extra spaces, line breaks, or unwanted characters the AI might have picked up from smudges or background textures on the paper. You'll know it worked when you've corrected all the obvious mistakes, and the digital text accurately reflects your original handwritten notes, making them much clearer than your original scrawl.

Step 6 — Save or export your editable text

After reviewing and correcting your digitized notes, you can now save them or share them. You'll see options on your screen to save the note within the app, copy the text to your device's clipboard, or share it to other applications. This final step means your old scribbles are now **editable text** (you can change it whenever you like) and **searchable text** (your device can find it instantly when you search).

Once you take action, your app will usually confirm that the note has been saved, or the text will be transferred to your chosen destination. If you don't see an explicit "Save" button, look for an arrow pointing left (←) or a "Done" button, which will usually save your changes within the app you're using. Free versions of apps might have limits on how many notes you can save or what export options are available, but basic saving and copying are almost always free.

Example: Tap the 'Save' or 'Done' button within your chosen app to store the note. Alternatively, look for a 'Share' icon (often a three-dot connection or an upward-pointing arrow) to copy the text to your phone's clipboard or send it directly to an email, a document app like Google Docs or Microsoft Word, or cloud storage like Google Drive or OneDrive. You'll know it worked when your app confirms the note is saved (e.g., you see it in your main notes list), or when you successfully paste the text into another application, ready for editing, searching, or sharing.

Common mistakes

- **Poor lighting or blurry photos:** Shadows or dim light make it incredibly hard for the AI to "read" your notes accurately, and blurry images are almost impossible for it to decipher. Always aim for bright, even light, and hold your phone steady when taking the picture. If the first attempt is blurry, simply delete it and retake the photo.
- **Not reviewing the extracted text:** It's easy to assume the AI got it all right, but even the best AI can make mistakes, especially with unique handwriting or complex layouts. A quick check and correction can save you from future confusion or embarrassment when you try to use the notes later. Always perform a quick read-through.
- **Using an app without a strong OCR feature:** Not all camera or note-taking apps have robust AI-powered OCR. If your results are consistently bad (lots of errors, missing text), try downloading a dedicated scanning app like Google Keep or Microsoft OneNote, which are known for good text recognition capabilities.

What this means for you

- **In everyday life:** Imagine quickly finding that recipe idea you scribbled down weeks ago, or having a digital copy of your kids' school notes that you can share instantly with a co-parent. No more rummaging through stacks of paper! This tool helps you keep important information organized and accessible, whether it's a shopping list, a quick thought, or important reminders.
- **For your business or work:** Effortlessly convert meeting notes, brainstorming sessions, or client feedback from paper to a digital format. This makes them searchable, easy to share with colleagues, and simple to integrate into digital documents, saving you valuable time on

transcription.

- **If you're just getting started:** Pick one short, clearly written note – like a shopping list or a phone number – and try to digitize it. It's a low-pressure way to see how well the AI works with your handwriting.

Try it now

Grab a piece of paper and quickly jot down three items on your grocery list. Now, open Google Keep (or Microsoft OneNote), take a clear photo of your list, and use the app's text recognition feature. In less than two minutes, you'll have your shopping list in an editable format, ready to paste into a messaging app or email!

How to turn a long PDF contract into a plain-English summary in 5 minutes

6 steps · Read online

> **In short:** Open any AI assistant that accepts PDF uploads, attach the file, and ask a few focused questions about deadlines, fees, and cancellation. You'll get a plain-English summary in minutes — but always read the original key clauses yourself, because AI summaries can miss important details.

Contracts hide the important parts deep in the small print: the auto-renewal clause, the 30-day cancellation window, the fee that "may apply" later. An AI assistant can pull all of that into a short, readable summary, but only if you ask the right questions. This guide shows the technique that works in any major AI assistant that accepts PDF uploads — ChatGPT, Claude, Gemini, Copilot, Le Chat, and others.

Before you start

- An AI assistant that can read PDF files you upload. Most major assistants do this on their free tier, but very long contracts (over roughly 100 pages) may need a paid plan. Check the upload limit on the assistant's pricing page if your document is unusually long.
- The contract itself, saved as a PDF on your computer or phone.
- About 10 minutes. The summary itself is fast; the careful read-through at the end takes a moment.
- A quiet moment to actually read what the assistant tells you.

Step 1 — Start a fresh chat and upload the PDF

Open your AI assistant and begin a brand-new conversation. A clean slate gives you the cleanest answer — older chats can confuse what the assistant reads in the new document. Once the chat is open, look for a **paperclip**, **attachment**, or **plus icon** near the message box. Tap that, then choose the PDF from your files. Some assistants also let you drag and drop the file straight onto the chat window.

If you don't see a paperclip, look for a **+ icon** or a button labeled something like **Add files** or **Upload**. On phone apps the icon often hides inside a small menu at the side of the text box.

Example: "I'm uploading a contract. Please don't analyze it yet — I'll ask my questions one at a time."

You'll know it worked when the file name appears above the message box (or as a small chip beside it) and the assistant confirms it can see the document. If nothing happens after 30 seconds, the file may be too large — check the assistant's file-size limit before assuming the tool is broken.

Step 2 — Ask for a plain-English overview first

Before you drill into the details, ask the assistant to give you a one-paragraph summary of what the document is, who it's between, and what it covers. This grounds the rest of the conversation: once the assistant has the big picture in mind, its later answers tend to be more accurate. Type your question in the message box where the uploaded file is already attached, and press Enter or the send arrow.

If the answer is longer than a paragraph or feels like a sales pitch, ask it to "shorten to one paragraph and skip the legal jargon." If the assistant gives you only a vague intro, ask a follow-up: "What kind of contract is this — a service agreement, a lease, a sales contract, or something else?"

Example: "In plain English, what is this document, who are the parties, and what is the main thing being agreed to?"

You'll know it worked when the assistant gives you back a short paragraph (3 to 6 sentences) that names the type of contract, the people or companies involved, and the core purpose. If any of those three is missing, ask one more focused question before moving on.

Step 3 — Ask the three questions that actually matter

Now the real work. In the same chat, type three focused questions, one after another: **deadlines** (when does this start, when can I leave, when does it auto-renew), **fees** (what does it cost, what changes after the first period, what extra charges might apply), and **cancellation** (how do I end this, what's the notice period, what happens if I miss it). Ask all three in one message so the assistant can answer them together. Press Enter and wait — usually under a minute.

If the assistant answers only one of the three, send the missing questions in a fresh message. If the answer feels too short, ask "Please list every date, dollar amount, and cancellation step you can find in the document."

Example: "From this contract, list: 1) every deadline or time limit I need to know, 2) every fee, charge, or price that might apply, and 3) exactly how I cancel, including any notice period. Use bullet points."

You'll know it worked when the assistant gives you a structured, easy-to-scan answer with three clear sections. If any answer is "I can't find that in the document," ask it to point to the page or section it looked at — sometimes the answer is there and the assistant missed it.

Step 4 — Push for the things people forget to ask

Once you have the big three, push one more layer. Contracts hide risk in the small print: what happens if you miss a payment, who pays legal fees in a dispute, whether the company can change terms mid-contract, and which country's laws apply. Ask these as a follow-up message in the same chat. The assistant now has the document **and** your earlier answers in its context window (the conversation history it can refer to), so it can build on what it already told you.

If the assistant's answer feels generic, paste in the actual clause: "I'm looking at section 8.2. What does this mean in plain English, and what's the worst-case scenario for me?"

💬 Example: "What happens if I miss a payment, who covers legal fees in a dispute, can the company change the terms mid-contract, and which state's laws apply? Point me to the section for each."

You'll know it worked when the assistant points to specific sections or page numbers and explains each risk in a sentence a non-lawyer can follow. If it just says "this is a standard clause" without explaining the actual risk, ask "what's the worst that could happen to me, in concrete terms?"

Step 5 — Ask for a one-page checklist you can save

Now ask the assistant to compress everything into a short, printable checklist you can keep next to the contract. A checklist is easier to skim than a long summary, and it's the format most people actually re-read later. Ask in the same chat — the assistant still has the document and your earlier conversation. The result should fit on one screen: action items, key dates, and money amounts, plain and short.

If the checklist comes back too long, ask it to "cut this to the 10 most important items I'd forget in six months."

💬 Example: "Summarize everything we just found into a one-page checklist I can save: dates I must remember, amounts I'll be charged, and the steps to cancel. Keep it under 15 lines."

You'll know it worked when you can copy the answer into a notes app and read it in under a minute. If the checklist still feels like a wall of text, ask the assistant to put each item on its own line with the date or amount in bold.

Step 6 — Spot-check three sections in the original PDF

AI summaries are good, but they're not perfect. The assistant can misread a clause, skip a footnote, or paraphrase something in a way that changes the meaning. Before you rely on the summary, open the original PDF and read three things yourself: the **cancellation clause** (usually labeled "Termination"), the **fees section** (often "Fees," "Charges," or "Payment Terms"), and any **date or deadline** the assistant flagged. If those three match what the assistant told you, the rest of the summary is very likely accurate.

If you can't find the section the assistant mentioned, use the PDF's search function (Ctrl+F on a computer, or the magnifying glass icon in most phone PDF apps) and type a keyword from the clause. If something doesn't match, send a follow-up question: "Section X says something different — can you re-read and reconcile?"

💬 Example: "Cancellation clause says 30 days, fees say a charge after year one, and auto-renewal is 60 days before expiration — can you re-read sections 7, 9, and 12 and tell me if your summary matches?"

You'll know it worked when the three sections you read match the summary, and you feel comfortable either signing the contract or deciding not to. If anything is off, send the assistant the exact text and ask it to explain the difference — that's how you catch the small print that changes everything.

Common mistakes

- **Asking one giant question.** A single "summarize this contract" prompt usually gives a vague, hard-to-scan answer. Splitting it into overview, the three big questions, the hidden risks, and a checklist gets you a result you can actually use.
- **Trusting the summary without checking.** AI assistants can misread clauses, especially around fees and dates. Always read the cancellation, fees, and deadline sections yourself before you sign. The summary is a map, not a replacement for the territory.
- **Letting a long chat get messy.** After 20 or 30 messages, the assistant can lose track of what it told you earlier. If the conversation gets long, copy the checklist into a new chat, re-upload the PDF, and start fresh. Cleaner chats give cleaner answers.
- **Skipping the free-tier limits.** Some assistants let you upload small files on the free plan but cap the size on paid plans. If your upload fails, check the assistant's file-size limit before assuming the tool is broken.

What this means for you

- **In everyday life:** Leasing a new apartment, renewing a phone plan, signing up for a gym membership — the same five-step question set works for any contract. Once you've done it twice, the second one takes about three minutes.
- **For your business or work:** Vendor agreements, client contracts, partnership letters, freelancer terms — all of these can be summarized the same way. The checklist step is especially useful when you need to brief a partner or a manager who hasn't read the document.
- **If you want to go further:** Try a second AI assistant on the same PDF and compare the two summaries. Two independent reads catch more than one — and they don't always agree on which clauses matter most. Just don't paste a confidential contract into a tool your

company hasn't approved for outside use.

Try it now

Pick the last contract you signed — apartment lease, phone plan, gym membership, anything with more than five pages. Upload it to your AI assistant and ask the three big questions from Step 3. Save the answer somewhere you can find it in six months, when the renewal notice arrives. You'll never read a long contract the same way again.

Unlock Document Insights with Google NotebookLM

6 steps · Read online

Imagine having a smart assistant that can read all your scattered notes, long reports, or important documents and then answer your questions using *only* the information you've provided. That's the power of Google NotebookLM. By the end of this guide, you'll be able to upload your own documents, ask intelligent questions, and quickly pull out key information or summaries without sifting through pages yourself. It's perfect for students, researchers, or anyone dealing with lots of text.

Before you start

- **A Google Account:** You'll need a standard Google account (like the one you use for Gmail or Google Drive) to access NotebookLM.
- **A Web Browser:** Access NotebookLM through a web browser on your computer. While it works on tablets, a computer provides the best experience for this guide.
- **An Internet Connection:** A stable connection is needed to upload documents and interact with the AI.
- **A few documents:** Have 1-3 documents (e.g., a PDF report, a text file of notes, or a Google Doc) ready to upload. These will be your "sources".
- **Cost:** NotebookLM is currently available for free.
- **Time:** This guide will take about 10-15 minutes to follow.

Step 1 — Access Google NotebookLM

First, you'll need to open Google NotebookLM in your web browser. This is where you'll create a dedicated space for your documents and questions. You'll see a welcome screen or an overview of your existing Notebooks, which are like individual project folders. If you're logged into your Google account, it should open directly. If it asks you to sign in, use your Google account details.

If it looks different: If you see a welcome tour or introductory pop-ups, simply click through them to get to the main interface. Look for a "Create new Notebook" or a large "+" button to start. **You'll know it worked when:** You see the NotebookLM interface, likely with a main content area and a sidebar, and perhaps a prompt to create your first Notebook.

Step 2 — Create a New Notebook

A "Notebook" in NotebookLM is like a digital binder where you keep all the related documents for a specific topic or project. It helps keep your information organized and ensures the AI only uses the sources relevant to that project. To create one, look for a prominent button or option to start a new project.

What you see: On the main screen, you'll typically find a large "New Notebook" button or a similar option in a sidebar. Clicking it usually prompts you to give your new Notebook a name. **What happens next:** A small pop-up window will appear, asking you to name your new workspace. Give it a descriptive title that reflects the documents you'll be adding, so you can easily find it later. **If it looks different:** If you don't see a "New Notebook" button, look for an icon like a "+", a folder, or text that says "Create". The exact placement might vary, but it's usually at the top or in a left-hand navigation pane. 💬 Example: `My Research Notes on Renewable Energy` **You'll know it worked when:** A new, empty Notebook with your chosen name appears on your screen, ready for you to add "sources".

Step 3 — Add Your Source Documents

"Sources" are the documents you want NotebookLM to read and analyze. These can be PDFs, Google Docs, text files, or even web pages saved as PDFs. The AI will then use *only* these sources to answer your questions, acting like a super-smart indexer for your own content. You can usually add multiple files to a single Notebook.

What you see: Inside your new Notebook, you'll see an area labeled "Sources" or "Add sources", often with an "Upload" button or an option to connect to Google Drive. Click the "Upload" button. **What happens next:** A file explorer window (or Google Drive selector) will open, allowing you to choose documents from your computer. Select one or more relevant files. NotebookLM will then process these files, which might take a few moments depending on their size and number. **If it looks different:** Look for an icon representing a document, a paperclip, or a cloud upload symbol if "Upload" isn't immediately visible. If you're using a Google Doc, there might be a direct integration option for Google Drive. **You'll know it worked when:** Your uploaded documents appear listed under the "Sources" section in your Notebook. You might see a "Processing" status initially, which will change to "Ready" when the AI has finished reading them.

Step 4 — Ask a Question Based on Your Sources

Now that NotebookLM has your documents, you can start asking it questions. Remember, the key here is that the AI (which is powered by a large language model, or LLM – a sophisticated computer program that understands and generates human-like text) will *only* refer to the information within the sources you provided. This prevents "hallucinations" – when an AI confidently makes up facts – and ensures its answers are grounded in your specific content.

What you see: In the main area of your Notebook, there's usually a chat box or a "Ask your sources" field, similar to what you'd find in a regular chat AI. Type your question here. **What happens next:** The AI will process your question, consult *only* the uploaded sources, and then generate an answer directly below your prompt. It will often include small citations (like "Source 1, Page 3") pointing back to where it found the information in your documents. **If it looks different:** The chat box might be labeled differently, perhaps "Query sources" or "Enter prompt". Just ensure you're typing your question into the area designated for interacting with your documents. 💬 Example: `What are the key challenges mentioned in the "Renewable Energy Report.pdf"?` **You'll know it worked when:** You receive a detailed answer from NotebookLM, specifically referencing your uploaded documents, with citations next to the relevant parts of its response.

Step 5 — Explore and Refine with Follow-up Questions

The conversation doesn't have to end with the first answer. NotebookLM is designed for back-and-forth interaction. If an answer is too broad, too short, or if you want to dive deeper into a specific point, you can ask follow-up questions. This helps you refine your understanding and extract exactly what you need.

What you see: After receiving an answer, the chat box typically remains open, or a "New prompt" button appears. You can type another question directly into this field, building on the previous interaction. **What happens next:** NotebookLM will take your new question and the context of your previous conversation into account, again pulling information solely from your sources. The new response will appear below the previous one, continuing the thread. **If it looks different:** If the chat box clears, simply re-type your follow-up question. The AI remembers the context of the current conversation within the Notebook. 💬 Example: `Can you summarise the solutions proposed for reducing these challenges?` **You'll know it worked when:** You get a new answer that directly addresses your follow-up question, further refining the information based on your sources, and continuing the conversational flow.

Step 6 — Generate an Overview or Summary

One of NotebookLM's most powerful features for a beginner is its ability to quickly generate an "Overview" or summary of your entire Notebook or specific documents. This is incredibly useful for getting the gist of large amounts of text without reading every word, or for quickly preparing for a meeting or study session.

What you see: Look for a button or option in your Notebook, often labeled "Overview", "Summarize Notebook", or "Generate Summary". It might be in the sidebar or at the top of the main conversation area. Click it. **What happens next:** NotebookLM will quickly process all the documents within that specific Notebook and generate a concise summary or a list of key topics and insights. This often appears as a new section or card in your Notebook interface. **If it looks different:** The feature might be represented by an icon, such as a bulleted list or a document with an arrow, or be found under a "Tools" or "More actions" menu. **You'll know it worked when:** A new, structured summary or overview appears in your Notebook, highlighting the main points and themes from all your uploaded sources.

Common mistakes

- **Expecting general knowledge answers:** NotebookLM *only* answers from your uploaded sources. If you ask a question not covered in your documents, it won't be able to help, and might even tell you it can't find the information. **Fix:** Always ensure your question can be answered by the content you've provided.
- **Uploading unsupported file types:** While NotebookLM supports common formats like PDFs and Google Docs, it can't read every type of file (e.g., image files, spreadsheets, or proprietary formats). **Fix:** Convert your documents to a supported format like PDF or plain text before uploading.
- **Asking overly broad or vague questions:** A very general question might lead to a generic or incomplete answer, even with relevant sources. **Fix:** Be specific with your prompts. If you want details on a particular section, mention it. If you want a specific type of information, ask for it directly (e.g., "List the pros and cons of X").

What this means for you

- **In everyday life:** If you're planning a trip, you could upload travel guides and articles into a Notebook and ask it to highlight the best restaurants or activities in your destination, saving hours of reading.
- **For your business or work:** You can drop in client reports, internal policies, or research papers and quickly get summaries, compare different sections, or pull out specific data points for presentations, making your research much faster and more accurate.
- **If you're just getting started:** Try uploading a favorite recipe or a hobby-related article. Then, ask NotebookLM simple questions like "What are the main ingredients?" or "What's the first step?" to get a feel for how it extracts information.

Try it now

Open Google NotebookLM, create a new Notebook, and upload two or three of your own PDF documents, like a household bill, a product manual, or an article you've saved. Then, ask it: "What are the key points in these documents?"

Summarize Any Long Document or Article with AI

6 steps · Read online

Tired of scrolling through a 20-page policy, a lengthy news article, or a long email thread? This guide will show you how to use a free AI tool to quickly get a short, reliable summary. In just a few minutes, you'll be able to cut through the noise and focus only on what truly matters.

Before you start

- **A device with internet access:** A computer, tablet, or smartphone is perfect.
- **A free AI chat tool:** We'll use general instructions that work for popular options like ChatGPT (find it at `chat.openai.com`) or Gemini (at `gemini.google.com`). You'll need to create a free account and log in.
- **A document or article to shorten:** This could be a webpage, a PDF where you can select text, an email thread, or any plain text.
- **Rough total time:** About 6 minutes to follow all the steps.

Step 1 — Copy your document text

First, open the document or article that you want the AI to summarize. Once it's open, you need to select all the relevant text. On a computer, you can often do this by clicking your mouse at the start of the text and dragging to the end, or by using a keyboard shortcut like `Ctrl+A` (Windows) or `Cmd+A` (Mac) to select everything. With the text highlighted, press `Ctrl+C` (Windows) or `Cmd+C` (Mac) to copy it to your **clipboard**. The clipboard is a temporary storage area on your device, like a hidden notepad, where copied items wait to be pasted.

You'll know it worked when the text is highlighted (often in blue or gray), and you've used the copy command. Although nothing visibly changes on your screen after copying, the text is now safely on your clipboard, ready for the next step. If you can't select text, try opening the document in a different program (e.g., a web browser for PDFs) or look for a "Select All" option within the application's menu.

Example: Open a long email thread. Press `Ctrl+A` (Windows) or `Cmd+A` (Mac) to select the entire email body, then press `Ctrl+C` (Windows) or `Cmd+C` (Mac).

Step 2 — Open your AI chat tool

Next, open your preferred free AI chat tool in your web browser or dedicated app. For instance, you might type `chat.openai.com` into your browser's address bar for ChatGPT, or `gemini.google.com` for Gemini. If you're using a phone or tablet, simply launch the app. If you're prompted, make sure you're logged into your free account.

You'll typically see the AI tool's main chat interface, which usually features a large, empty text box at the bottom of the screen. This is where you'll type your messages to the AI. The tool is now waiting for your instructions. If your screen looks different, or you see old conversations, look for a "New Chat" button, a plus sign icon (`+`), or a similar option to start a fresh, blank conversation.

Example: Open your web browser and navigate to `gemini.google.com`. Log in to your account if prompted, then ensure you have a blank conversation ready.

You'll know it worked when you see an empty chat box, ready for you to type your first message, usually labeled "Message [AI name]..." or "Type your prompt here...".

Step 3 — Craft your summary request (Prompt)

Now, in the AI's chat box, you'll type your instruction for the summary. This instruction is called a **prompt** – it's like giving a clear job description to the AI. Be specific about the desired length, format (e.g., bullet points), and tone for your summary. A good prompt helps the AI understand exactly what you need.

You'll see the words you type appearing in the chat box. At this stage, the AI is just receiving your input, and you haven't sent anything yet. If the chat box shows pre-written suggestions, simply ignore them and type your own request. Being clear upfront saves time later.

Example: Type this into the chat box (do not press Enter or Send yet):

> "Please summarize the following document in three concise bullet points. Keep the language neutral and factual, avoiding any personal opinions or interpretations."

You'll know it worked when your complete summary request is typed into the box, clearly visible, but you haven't sent it to the AI yet.

Step 4 — Paste your text and send

After typing your summary request, move your cursor to the end of your prompt, then immediately paste the long document text that you copied in Step 1. You can do this by pressing `Ctrl+V` (Windows) or `Cmd+V` (Mac). The pasted text will appear right after your instruction in the same chat

box. Once both your prompt and the entire document text are in the box, press Enter or click the "Send" button (often a paper plane icon) to send everything to the AI.

You'll see your prompt and the long document text appear in the chat history. After a moment, the AI will start generating its response. This processing is called **inference**, which is when the AI "thinks" about your request and creates an output. If your document is extremely long (e.g., hundreds of pages), the AI might tell you it's too much text for its **context window**. The context window is like the AI's short-term memory – the amount of information it can "remember" and process at once. If this happens, you might need to break your document into smaller chunks and summarize them separately, or consider a paid plan, as some free versions may have usage restrictions.

Example:

> "Please summarize the following document in three concise bullet points. Keep the language neutral and factual, avoiding any personal opinions or interpretations.
>
> [Now, paste your copied document text here, directly after the prompt.]"

You'll know it worked when you see your message disappear into the chat history, and the AI begins to type out its summary, usually appearing as new text below your input.

Step 5 — Check and refine the summary

Once the AI has generated its summary, carefully read through the bullet points it provided. Compare them against the original document to ensure they are accurate, don't miss any crucial details, and match the tone you requested. Sometimes, AI can "hallucinate" – meaning it confidently makes up information that isn't present in the original text, or misinterprets something. If anything feels off, needs clarification, or isn't quite what you wanted, simply type a follow-up question or instruction in the chat box.

You'll see the AI's initial summary in the chat. When you ask a follow-up question, your new prompt will appear, and then the AI will try to adjust its response based on your feedback. If the AI completely misunderstands or gives a poor response, you might need to try rephrasing your initial prompt more clearly or start a new chat conversation altogether.

Example:

> "The third bullet point mentions 'Section 5'. Please confirm the exact wording of Section 5 from the original text and update that bullet point to be more precise. Also, can you ensure the overall tone is strictly professional?"

You'll know it worked when the AI responds with an updated summary that addresses your feedback, making it more accurate and aligned with your specific needs.

Step 6 — Finalise and use your summary

Once you are completely happy with the AI's summary, it's time to make it your own. You can copy the final bullet points from the chat interface to your clipboard. Most AI tools allow you to highlight the text with your mouse (click and drag) and then use the copy command (`Ctrl+C` or `Cmd+C`). Some tools may also have a specific "Copy" button or icon next to the AI's response to make this easier.

You'll see the polished, accurate summary on your screen. After copying, the text is on your clipboard, ready to be pasted into another application, like an email, a report, or your personal notes. If there isn't an obvious "copy" button, simply highlight the AI's text and use your device's standard copy function.

Example: Paste the finalised bullet points into a new email, a team chat application like Microsoft Teams, or a note-taking app on your phone to share or save.

You'll know it worked when the summary text is successfully copied (you can test this by pasting it into a simple text editor like Notepad or Notes) and you can then paste it into your desired application.

Common mistakes

- **Too vague a request:** Asking for "a summary" without specifying length or format can result in a lengthy, unstructured response. **The fix:** Always be specific. Say "in three bullet points," "one short paragraph," or "a summary suitable for an executive" in your prompt.
- **AI "hallucinations" or opinions:** AI tools can sometimes confidently present incorrect information or add their own interpretations that aren't in the original text. **The fix:** Always include phrases like "keep the language neutral and factual; do not add your own opinions" in your initial prompt, and critically review the summary against the original document. Don't be afraid to ask for corrections.
- **Text too long for the AI:** Trying to summarize extremely long documents (e.g., an entire book or a very dense legal document) might hit the AI's **context window** limits. **The fix:** Modern AI tools have much larger context windows than before, but extremely long documents may still require breaking them into smaller sections, or using a paid plan if available.

What this means for you

- **In everyday life:** The next time you receive a long work email, a detailed news article, a community newsletter, or even a lengthy recipe blog post, you can quickly get the main points without reading every single word. This saves you valuable time and helps you stay informed efficiently.
- **For your business or work:** Condense lengthy reports, meeting minutes, customer feedback surveys, or policy documents into concise, actionable bullet points. This can help you prepare for presentations faster, communicate key information efficiently to colleagues, or quickly grasp the core of complex documents. Remember, while AI is incredibly helpful, it cannot provide legal or medical advice, and human review is crucial for critical documents.
- **If you're just getting started:** Pick one long email thread you've received today, or a lengthy article you saved to read later. It's an easy first step that will quickly show you the power of AI summaries and build your confidence in using these tools.

Try it now

Open any long email thread you've received in the last week

How to Prepare for a Job Interview Using a Pasted Job Post

6 steps · Read online

> **In short:** Paste a job listing and a short note about your background into your AI assistant, and you'll come away with ten likely interview questions, two talking points for each, and five thoughtful questions to ask the interviewer — typically in under ten minutes of back-and-forth.

A good interview answer sounds rehearsed without sounding rehearsed. The trick is having specific stories and numbers ready before you sit down — and that's exactly what an AI assistant is good at when you give it the right inputs. In fifteen minutes, you can walk into a room with three lists ready: questions they'll probably ask, talking points for each, and questions you'll ask them. It can draft the prep, but it can't read the room or know which story from your past will actually land — that's still on you.

Before you start

- **An AI assistant you already use** — ChatGPT, Claude, Gemini, Copilot, or any other chat-style assistant. The steps are the same.
- **The full job description** copied to your clipboard. Don't paraphrase yet; the more detail the assistant sees, the sharper the output.
- **A short paragraph about you** — your current role, years of experience, and one or two wins you'd actually mention in an interview.
- **About 15 minutes** — enough time to do the back-and-forth and skim the result.

Step 1 — Open a fresh chat and name it after the role

Click **New chat** (or **New conversation** — look for a pencil-and-paper icon or a + sign) in your assistant. Before you type anything, give the chat a name you'll recognize later, like "Acme Senior Designer interview" or just the company plus the role.

Why a new chat? Older messages from unrelated work can leak in and confuse the assistant. A fresh, named thread keeps everything you build here tied to this one interview — easy to find the night before, easy to revisit if a second round is scheduled.

Example: name it "Acme – Marketing Lead interview prep."

You'll know it worked when a new empty chat opens with the name you chose (or with an editable title field at the top of the screen).

Step 2 — Paste the job description and your background in one message

In your new chat, send a single message that contains two clearly labeled blocks: the job description, and a paragraph about you. Use clear labels like **"JOB POSTING:"** and **"ABOUT ME:"** so the assistant knows where one ends and the other begins. Very long descriptions (over a few thousand words) may need to be split into two messages if your app refuses the paste.

What happens next: the assistant usually replies with a short acknowledgement ("Got it — ready when you are") or jumps straight into a summary. Either is fine.

If it looks different: if pasting a long posting fails, try uploading the posting as a file — most assistants accept PDFs and text files in the same chat where you usually send text.

Example:

```
JOB POSTING:
[Paste the full job description here — responsibilities, requirements, and any "nice to haves"]

ABOUT ME:
[Paste 4–6 sentences: current role, years of experience, two recent achievements with numbers if you have them]
```

You'll know it worked when the assistant confirms it has the context — usually with a brief recap, a clarifying question, or simply by being ready to answer your next prompt.

Step 3 — Ask for ten likely interview questions

Send a single, specific prompt asking for ten questions the hiring manager is likely to ask, based on what you just pasted. Ask for a mix of categories — behavioral, technical, and situational — so the list isn't one-note.

Why ten? Fewer than that and you may miss the ones that come up. More than ten and you'll feel overloaded. Ten is the sweet spot: enough to practice, small enough to actually read.

Example: "Based on the job posting and my background above, give me ten interview questions a hiring manager is most likely to ask me. Mix behavioral, technical, and situational questions. For each, add one sentence on what they're really trying to find out."

If it looks different: if the list comes back generic, follow up with "make these specific to [the company or team from the posting]" and ask for a fresh list. A vague list is usually a sign the original prompt needed more context, not a sign the assistant is broken.

You'll know it worked when you have a numbered list of ten specific, role-tied questions — not stock ones like "tell me about yourself."

Step 4 — Get two concrete talking points for each question

Without starting a new chat, follow up with a second prompt asking for two concrete talking points under each of the ten questions. Ask for stories, examples, or numbers drawn from your own background — and tell the assistant to flag any question where your pasted background looks thin and you'll need to dig deeper yourself.

What happens next: the assistant writes a structured answer, often with a short headline per question and a couple of ideas underneath. This is the moment the prep actually becomes useful.

If it looks different: if the response is too long to scan, add "keep each talking point to two sentences max" in your next prompt, or ask it to rewrite the same list more concisely.

Example: "Now for each of those ten questions, give me two concrete talking points I could use, drawing from my background. Keep each point to one or two sentences. Flag any question where my pasted background seems thin."

You'll know it worked when each of the ten questions now has two short, specific points you could picture yourself saying out loud.

Step 5 — Ask for five smart questions to ask back

In the same chat, ask for five questions you, the candidate, should ask the interviewer. Ask for ones that show you did your homework — about the team, the goals, or how the role is measured. The point is to leave the conversation with real signals, not rehearsed fluff.

Example: "Give me five strong questions I, the candidate, can ask back at the end of the interview. Make them specific to this role and company. Skip anything generic like 'what's the culture like'."

If it looks different: if any of the questions come back feeling templated, follow up with "replace any you can't tailor to this exact role" — most assistants will rewrite them on the spot.

You'll know it worked when you have five questions you wouldn't have thought of on your own, and each one is something you'd genuinely want answered.

Step 6 — Pin or save the chat and skim the highlights

Take thirty seconds to skim the chat and star, pin, or copy the talking points that felt most true to you. Different apps handle this differently — look for a **star icon** next to an answer, a **pin** option in a menu, or just copy the best points into a notes app where you'll actually see them again.

Why bother? Two reasons. First, you'll want to reread the chat the morning of the interview. Second, picking the strongest two or three talking points now means you practice those — not all twenty.

If it looks different: if your assistant has no pin or star option, the chat's title alone usually makes it easy to find later. Or paste the talking points into your phone's notes app under the company name.

Example: pick the talking points that felt truest to your real work — maybe one about leadership, one about a measurable result, one about a hard problem you solved — and re-read them once before you walk into the interview.

You'll know it worked when you can open the chat or your notes in under ten seconds and find the three talking points you actually plan to use.

Common mistakes

- **Pasting only the company blurb.** If the description is six paragraphs long and you paste just the first one, the assistant's questions will be vague. Paste the whole thing — responsibilities, requirements, and any "nice to haves" — so the assistant has the full picture.
- **Asking for "interview questions" without context.** A generic request gets generic answers. Always include both the job description and your background in the same chat; that's what makes the list specific to *this* role and *your* level.
- **Stopping after the first list.** The ten questions are the start, not the end. The follow-up that asks for talking points is where the real prep happens. Skip it and you'll have a list of questions with nothing ready to say.
- **Treating the output as a script.** The assistant may suggest a story you don't actually have. Treat the talking points as a draft, not a script — swap in your own real examples whenever you can.
- **Reusing the same chat for unrelated work.** Once you start mixing topics, the assistant's answers drift. Keep this chat dedicated to one interview until you're done with it.

What this means for you

- **In everyday life:** the same recipe works for any "I have a document and I want to be ready" moment — a school presentation, a doctor's appointment where you want a list of questions, or a tough conversation with a landlord. Paste the source, share your context, ask for what you need.
- **For your business or work:** if you ever sit on the other side of the table, paste the job description and a candidate's resume into the same flow and you'll get a quick read on likely strengths, gaps, and good first-round questions — useful when screening applications.
- **If you want to go further:** the next level is saving this whole structure as a **Custom Instruction** or a saved template in your assistant, so the next time you prep, you only paste the job posting and your background — the rest of the structure is already loaded. Many assistants also let you create a dedicated "Project" or named workspace for one company, so multiple chats (prep, follow-up, salary research) live together.

Try it now

Open a new chat in the assistant you already use, name it after a real interview you have coming up (or one you're imagining), and paste one real job description plus four sentences about your background. Then send the prompt from Step 3. Ten minutes from now, you'll have a list you didn't have ten minutes ago.

Write a Resume From Scratch (Even With No Work Experience)

6 steps · Read online

> **In short:** A resume is a one-page story of what you've done, what you're good at, and what you want next — and you can build one from school, volunteering, caregiving, side jobs or hobbies. You don't need formal work experience; you need to describe your real life in the language employers recognize.

Writing a resume from zero feels harder than it is. Most "first resumes" stay blank because people think they have nothing to put on them. In reality, you've already handled deadlines, solved problems, dealt with people, and learned new things — you just haven't framed that work as work yet. This guide walks you through turning everyday experience into a clean, honest resume a hiring manager can read in under a minute.

Before you start

- **A computer or phone** with a notes app, Google Docs, Microsoft Word, or a free resume builder (look for one by searching "free resume template").
- **About 60–90 minutes** for your first draft. After that, updates take 10 minutes.
- **A list of what you've actually done** — school, classes, projects, babysitting, family jobs, club roles, volunteering, online courses, certifications. Don't filter yet; just dump everything.
- **A job (or two) you'd like to apply to.** Even a saved job posting helps you match words later.

Step 1 — Pick a clean, one-page layout

Open Google Docs and click **Template Gallery** in the top-right, or search for "resume" in Microsoft Word's template search. Pick something simple: name at the top, contact line below it, then clear section headings (think "section headers" — the bold labels that divide your resume into parts like Experience and Education). Avoid heavy designs, photos, or color blocks; hiring managers scan, they don't admire.

If templates feel overwhelming, just open a blank document and type your name big at the top. You'll add structure next. If you don't have Word or Docs, free builders like the ones from job search sites work too — search "free resume builder" and pick one with no signup wall.

Example: section headers to copy — `Summary`, `Experience`, `Education`, `Skills`, `Projects`.

You'll know it worked when you have a blank one-page document with your name at the top and five empty sections waiting for content.

Step 2 — Write a 2-sentence summary at the top

Under your name and contact info, add a **summary** — two short sentences that say who you are and what you're aiming for. Think of it as the cover of a book: it tells the reader what to expect inside. Use plain words, not job titles you haven't earned.

If you've never had a job, your summary can lean on what you're studying, what kind of work you want, and one strength. Keep it specific — "hardworking" doesn't say much, but "organized, comfortable with spreadsheets, and quick to learn new tools" paints a picture. If you don't know exactly what you want yet, write a **target** instead: "Looking for an entry-level role in retail, admin, or customer service."

Example: "High school senior finishing in May 2026. Strong with customers, cash handling, and staying calm under pressure. Looking for part-time retail or front-desk work."

You'll know it worked when someone reading just those two sentences could guess what kind of job you'd be good at.

Step 3 — Turn life experience into an "Experience" section

Create an **Experience** section heading and list anything where you showed up, did something, and someone (even a family member or teacher) noticed. Jobs count, but so do babysitting siblings, helping at a family store, school projects, volunteering at a food bank, mowing lawns, Discord servers you moderate, or a class group you led. For each entry, use this format: **Title — Where — Dates** (even if dates are approximate), then 2–4 short **bullet points** underneath describing what you did.

Each bullet should follow a simple recipe: *did what, using what skill, with what result*. You don't need fancy words; you need honest ones. If you watched three younger cousins every afternoon while your parent worked, that's "supervised and cared for three children daily, planned meals, handled minor conflicts, and maintained a consistent routine." That's a real bullet.

Example (babysitting): `Cared for three children (ages 4–9) after school each day. Prepared snacks, helped with homework, and organized weekly craft activities.`

Example (school project): `Led a 4-person team on a science fair project. Researched topics, delegated tasks, and presented findings to a panel of 5 judges.`

You'll know it worked when your experience list has 3–6 entries and every bullet describes something you actually did.

Step 4 — Build an "Education" section and a "Skills" list

Add an **Education** section right below Experience. Include your school (or "Self-taught" / "Online coursework"), graduation year (or expected year), and one line of relevant classes or achievements — anything that hints at what you know. No degree yet? No problem: list what you ARE studying, plus any courses, workshops, or certificates, even free ones from sites like Coursera, Khan Academy, or YouTube tutorials.

Then add a **Skills** section. This is just a comma-separated list of tools and abilities: typing speed ("60+ WPM" if true — WPM means words per minute), languages, software (Google Sheets, Canva, TikTok editing, whatever you actually use), and people-skills ("customer service," "calm under pressure"). Skip anything you can't honestly do.

💬 Example education entry: `Lincoln High School — Expected graduation May 2027. Coursework: Business Math, Intro to Computer Science, Spanish II.`

💬 Example skills line: `Skills: Google Docs, Sheets, Canva, social media scheduling (Meta Business Suite), conversational Spanish, POS cash handling.`

You'll know it worked when a stranger could scan these two sections and quickly see what you know and where you learned it.

Step 5 — Add optional sections that strengthen your story

Two more sections often help first-time resumes: **Projects** and **Volunteering**. A project can be anything you built or organized — a class group, a club, a YouTube channel, a small Etsy shop, a Discord community, a coding exercise. Volunteering covers anything you did for free that helped people — church groups, school events, neighborhood cleanups, food drives, coaching a kid's team.

You can also add **Languages**, **Certifications** (CPR, food handler's permit, even a free online certificate you actually finished), or **Awards** (perfect attendance, honor roll, employee-of-the-month at a family business — yes, that counts). If you have nothing extra, skip these — a tight two-page-equivalent on one page beats a sparse two pages.

💬 Example project entry: `Project: Neighborhood plant-swap app — designed a simple scheduling tool in Google Sheets, promoted via local Facebook group, coordinated 3 successful swaps with 40+ attendees.`

You'll know it worked when your resume fills the page naturally and every line answers "what did this person actually do?"

Step 6 — Proofread, shorten, and save as a PDF

Read your resume out loud once. Fix anything that sounds awkward, too long, or that you can't defend in an interview. Most first drafts are too long; aim for one page. Cut filler like "hardworking team player" unless you back it up with an example. Replace vague verbs ("helped," "worked on") with specific ones ("organized," "built," "answered," "tracked," "taught").

Save the final file as a **PDF** (in Google Docs: **File → Download → PDF**; in Word: **File → Save As → PDF**). PDFs keep your formatting intact no matter who opens them. Name the file clearly: `FirstName-LastName-Resume.pdf`.

💬 Example: cut this — "Responsible for various tasks including helping customers." Replace with — "Greeted 80+ customers daily, restocked shelves, and handled returns at a family-run grocery store."

You'll know it worked when you can email the file, the formatting holds, and you feel ready to explain any line on it out loud.

Common mistakes

- **Mistake:** Leaving the page half-empty because you think you "don't have experience." **Fix:** List school projects, caregiving, volunteering, hobbies done seriously, or small paid gigs — anything with responsibility counts. If the page is still sparse, add a one-line "Interests" section (photography, gaming communities, language learning) and tie one to a skill.
- **Mistake:** Using only generic words like "hardworking" and "team player." **Fix:** Replace each with a concrete example — "trained two new volunteers on check-in procedures" beats "good trainer" every time.
- **Mistake:** Lying or stretching the truth. **Fix:** Stick to what really happened, then frame it precisely. A hiring manager will ask follow-ups; honest, well-described experience always beats padded claims.

What this means for you

- **In everyday life:** A resume is just a written summary of your life so far — once you've made one, updates for the next 10 years take minutes, not hours.
- **For your business or work:** If you hire people, a resume built this way tells you whether someone can actually describe what they did — a skill that predicts on-the-job clarity.
- **If you're just getting started:** Your first resume won't be perfect, and that's fine. Send it to one job you actually want, see what feedback comes back, and revise. Resume-writing is a skill, and it improves with each round.

Try it now

Open a blank document right now and type your name at the top. Underneath, write one sentence describing what you're good at, and one sentence describing the kind of work you'd like. That's your summary. You just started a real resume — the rest is filling it in, one honest line at a time.

How to Organize a Messy Email Inbox Into Simple Folders in One Sitting

6 steps · Read online

> **In short:** Sort by sender, archive or delete anything that doesn't need a reply, then create just three broad folders (To act on, To read later, To keep) so the system still works a month from now. The whole thing takes about 30 minutes.

A messy inbox isn't a personal failure - it's what happens when there's no system. Here's a 30-minute plan that creates one simple system, designed to keep working long after you've finished. It works in any major email app because the idea is the same everywhere: stop the noise at the door, then give the messages that matter a home.

Before you start

- A computer or phone with your email app or website open. The steps apply generally to Gmail, Outlook, Apple Mail, Yahoo, and similar services - exact button names vary, but the ideas don't.
- About 30-45 minutes of quiet time, with no other tabs open competing for attention.
- Your password nearby in case you get logged out, and a notepad (or a note on your phone) so you can jot down anything that needs a real follow-up.

Step 1 — Pick your window and open the inbox the right way

Choose a 30-45 minute block when you won't be interrupted - a Sunday morning, a lunch break, after the kids are asleep. Open your email app or website and sign in. The key move: switch your inbox view from "Unread" or "Priority" to something that shows everything - often labeled "All Mail" or "All Messages." This matters because you need to see the full mess before you can fix it. Think of it like clearing the table before you start sorting the dishes.

If your app doesn't have a clear "All Mail" option, look for a folder on the left sidebar labeled something like "All," "Everything," or "Archive." The goal is to see every message in one long list, not just the unread ones. On phone apps, the default view usually already shows everything in your inbox, which works fine.

Example: tell yourself out loud, "I'm spending the next 35 minutes on this. No answering anything that needs more than 60 seconds."

You'll know it worked when your full inbox is visible in one list - usually hundreds or thousands of messages, not just the unread ones.

Step 2 — Sort by sender so the pattern jumps out

Find the option to sort or group your email by sender. In many webmail apps this lives under a sort menu (look for a downward arrow, a "Sort by" option, or three vertical dots near the top of the message list). On a phone app, look in the top-right corner for a filter or sort icon. When you group by sender, the same address suddenly appears 20, 30, or 50 times - that's your noise made visible.

If you can't find a "sort by sender" option, search for one big sender's name instead - the one whose name you recognize the most in the list. The point is the same: figure out who's filling your inbox before you start moving anything.

Example: search "Amazon" and count how many emails you find. If the number is in the double digits, that's a sender you'll handle in Step 3.

You'll know it worked when you can see, at a glance, which senders dominate your inbox - usually a handful of marketing lists, a few retailers, and some auto-notifications.

Step 3 — Clear the noise: archive or delete

Start with your biggest offenders - the senders who appear most often. For each one, you have three honest choices: read and reply (rare), archive the whole batch (the default for anything that doesn't need a response), or delete it. Look for an "unsubscribe" link near the bottom of marketing emails and click it - that legally stops future mail from the same sender. "Archive" just means moving a message out of your inbox while keeping it findable, like putting a paper into a filing cabinet instead of leaving it on your desk.

If you're not sure which messages to keep, ask one question: "Will I need this in six months?" A receipt for a big purchase, yes. A flash sale from a store you visited once, no. Anything you can't answer becomes "archive" - still searchable, just not in your way.

Example: for the Amazon thread - open one message, click "unsubscribe" near the bottom, wait for the confirmation page, then select all Amazon emails (look for a checkbox at the top of the list) and click "Archive" or "Delete."

You'll know it worked when unsubscribe confirmations start arriving in your inbox saying "you've been removed" - that's a sign you've cut the supply line.

Step 4 — Create your three broad folders

Now make just three folders. Pick simple, broad names that won't go stale next month: "To act on" (anything requiring a real reply or decision), "To read later" (interesting but not urgent - articles, newsletters you'll actually open), and "To keep" (receipts, confirmations, contracts, anything you might need to find again). A "folder" here is just a named container for related emails - like a labeled drawer in that same filing cabinet.

Most email apps have a "+" or "Create new folder/label" button - look for it in the left sidebar, often just above the existing folder list. Right-clicking the sidebar often brings up the same option. Keep the names short - one or two words each.

If your app uses "labels" instead of folders (some do), the idea is the same: three broad categories that you tag messages with. You can apply more than one label to the same message later if you need to.

Example: name them "To act on," "To read later," "To keep" - exactly those three, no more. Resist the urge to add a fourth on the same day.

You'll know it worked when all three folders appear in your left sidebar and you can click into an empty one.

Step 5 — Sweep the remaining emails into the right folder

Go back to your inbox. Now that the noise is gone, what remains are messages that probably need your attention or your future self's attention. Work through them quickly, one at a time: open, decide, file. Most belong in one of your three folders. A few might need a quick reply - if it takes less than 60 seconds, reply now; if it takes longer, file it under "To act on" and keep moving.

If a message doesn't fit any of the three folders, archive it instead of creating a new one. The point of three folders is that they stay broad - resist the urge to add a fourth until you've lived with these for at least a month.

Example: a newsletter you want to read on the weekend → drag to "To read later." A package tracking update → drag to "To keep." A friend's question that needs a real answer → drag to "To act on."

You'll know it worked when your actual inbox (the view that doesn't include your new folders) is mostly empty - ideally under 20 messages.

Step 6 — Add one simple rule so it keeps working

Find your email app's "filter" or "rule" feature - usually under Settings, then a tab labeled "Filters," "Rules," or "Inbox rules." A filter is just an automatic instruction: "when an email arrives from this address, move it to that folder." Create one rule for your biggest remaining sender (often a newsletter you actually like, or a frequent work contact).

This is the step that keeps the system working months from now. Without a rule, you'll have to re-sort every week. With even one rule, the inbox stays close to empty on its own.

If your email app doesn't support automatic rules, do the next best thing: make a tiny daily habit of glancing at your inbox once a day and dragging anything obvious into the right folder - two minutes is usually enough.

Example: in Gmail settings, choose "Create new filter" → From: newsletter@example.com → "Apply the label: To read later." Every future issue lands in the right folder automatically, no work from you.

You'll know it worked when the next email from that sender shows up already filed - and you didn't have to do anything.

Common mistakes

- **Creating too many folders.** The temptation is to make one for "Bills," one for "Work," one for "Family" - and suddenly you have fifteen folders, half empty, and you can't find anything. Three broad folders beats ten narrow ones. Add a fourth only after a real, repeated need shows up.
- **Filing things you'll never look at again.** "To read later" can quietly become a graveyard of saved articles. If a newsletter has been sitting unread for two weeks, archive it unopened. The point of saving is to read, not to feel productive.
- **Skipping the unsubscribe step.** Just archiving the noise doesn't stop it from arriving tomorrow. The unsubscribe link is the actual fix - archive is just the bandage. Spend the extra 10 seconds per sender; it pays back forever.

What this means for you

- **In everyday life:** A clean inbox means you stop dreading the little red badge on your phone. You see the few messages that actually matter when you open the app, instead of scanning past 400 to find them.
- **For your business or work:** When you search for an old receipt or a past conversation, you find it in seconds because everything has a home. Tax season, warranty claims, and "what did they say last March?" all get easier.
- **If you're just getting started:** Tomorrow morning, spend 10 minutes only on Step 3 - clearing one big sender. Don't try to do all six steps at once yet. Get a feel for the rhythm before committing the full 30 minutes.

Try it now

Open your email app right now and find the single sender with the most emails in your inbox. Click "unsubscribe" at the bottom of one of their messages. That's it - one minute, one sender, one less source of noise tomorrow.

Practice job interview questions with Qwen before the real interview

6 steps · Read online

> **In short:** Qwen can play the role of an interviewer, ask you common questions, and give you honest feedback on how your answers come across — clarity, completeness, and confidence. You'll need about 20–30 minutes for one full practice run, and a free account is enough to start. It cannot read your body language or tone, so for video interviews, do at least one round on camera separately.

By the end of this guide, you'll have run a real mock interview with Qwen, gotten concrete feedback on your answers, and saved your best versions for the night before the real thing. This is for anyone with an interview coming up — first job, career change, or just wanting to sound more polished.

Before you start

- A Qwen account (free sign-up with email or a supported login).
- A device with a browser — a computer is easier for long answers, a phone works too.
- A job title or two in mind that you actually want to practice for.
- About 20–30 minutes for a full session, or 5 minutes for a quick drill.

Step 1 — Open Qwen and start a fresh chat

Go to Qwen's chat page in your browser, or open the Qwen mobile app if you have it installed. Look for a **"New Chat"** button or a small **"+"** or pencil icon — usually at the top of the sidebar on a computer, or right above the message box on a phone. Click or tap it. A fresh, empty conversation opens with a blank message box at the bottom. Starting fresh matters because leftover context from an earlier chat could push Qwen's questions in the wrong direction. If you don't see a New Chat button, try refreshing the page — sometimes the sidebar hides it on smaller screens.

Example: just type "Hi" first to confirm the chat is responsive, then move to Step 2 to set up the real practice.

You'll know it worked when you see an empty message box with no previous messages above it.

Step 2 — Tell Qwen what role you're practicing for

In the message box, type a short paragraph that gives Qwen the context it needs to ask realistic questions. The more specific you are about the job, the better the questions will be. Mention the job title, the industry, the seniority level, and one or two skills you want to highlight. Press Enter or tap send. Qwen usually replies with a short confirmation like, "Got it — let's practice for a marketing coordinator role." If the reply feels off or skips your details, scroll up, click the **regenerate** button (usually a circular arrow), or retype your setup more simply.

Example: "I have an interview next week for a marketing coordinator role at a mid-size SaaS company. I have 2 years of experience in social media and email campaigns. Please act as the hiring manager and interview me one question at a time. After each answer, give me feedback on clarity, completeness, and confidence, then move to the next question."

You'll know it worked when Qwen replies acknowledging the role and confirms it will go one question at a time.

Step 3 — Answer the first question like it's real

Read Qwen's question carefully, then answer as if you were sitting in the actual interview. Type your answer in the same box — don't overthink it, don't use perfect essay language, just answer the way you naturally would. The point is to practice, not to write something polished. If you blank out, write "I'm not sure" the way you might say it out loud, then keep going. The more honestly you answer, the more useful Qwen's feedback will be. If the question feels too easy, ask for a harder version before answering — this is normal interview prep, not a sign you're doing it wrong.

Example: if Qwen asks "Tell me about a time you handled a difficult customer," you might type something like: "At my old retail job, a customer wanted to return a laptop past the 30-day window. I listened to her frustration, explained the policy, then asked my manager if we could do a store credit as a one-time exception. She agreed, and the customer left happy. I learned that sometimes bending the rules a little keeps a customer for life."

You'll know it worked when you've typed a full, honest answer — anything between 3 and 12 sentences is a typical interview length.

Step 4 — Ask for specific feedback on your answer

After you finish typing, don't wait for Qwen to move on. Type a follow-up asking it to evaluate what you just said, and be specific about what kind of feedback you want — clarity, completeness, confidence, filler words, structure. Qwen usually gives you a paragraph or two of concrete suggestions, like "your answer is clear but you ramble in the second sentence" or "you named the result but didn't explain your specific role." If you want a number, ask for a score out of 10. If the feedback feels too soft ("great answer, well done!"), push back: ask Qwen to be tougher, since you'd rather hear it now than in a real interview room.

💬 Example: "Before the next question, can you give me feedback on my last answer? Specifically: (1) Was it clear and easy to follow? (2) Did I actually answer what was asked? (3) Did I sound confident or unsure? (4) What is the one thing I should improve? Please be honest — I'd rather hear it from you than from a real interviewer."

You'll know it worked when Qwen responds with specific, actionable feedback — not just a generic compliment, but actual comments about what was strong and what to improve.

Step 5 — Push yourself with follow-up and pressure questions

Real interviewers rarely stop at one question on a topic — they ask follow-ups like "and what would you do differently?" or "how did your manager react?" So once you've answered and gotten feedback, ask Qwen to dig deeper. Type a message asking for a follow-up, or simply say "ask me a harder version of that one." If you really want to stress-test yourself, ask Qwen to challenge you: "Now push back on my answer as if you don't believe me." This is where most of the real interview prep happens — the first answer is usually fine, but it's the second and third round that separates nervous candidates from confident ones.

💬 Example: "That feedback helps. Now please ask me a follow-up question about the same story — something a real interviewer would press me on. I'll answer, you give me feedback, and then we move on."

You'll know it worked when Qwen asks a sharper, more specific question that builds on what you said — not a brand-new topic.

Step 6 — Save your best answers for the day before

Once you've practiced for 20–30 minutes and feel you've landed some strong answers, ask Qwen to summarize your session. Then copy your best answers into a notes app, a Google Doc, or even a simple text file on your phone. The night before the interview, re-read your saved answers out loud — this refreshes the words in your mouth so they sound natural, not memorized. If Qwen gave you specific suggestions ("trim the first sentence, end with a number"), apply those edits now and save the cleaner version. Treat the saved file as your personal answer bank — you can reuse and refine it for future interviews.

💬 Example: "Can you give me a summary of the three strongest answers from this session, plus one thing I should remember to improve next time?"

You'll know it worked when you have a short document with your polished answers that you can re-read on the morning of the real interview.

Common mistakes

- **Asking Qwen to "just give me good answers."** That turns it into a cheat sheet, not practice. You learn by answering yourself and getting feedback, not by copying what Qwen writes.
- **Pasting ten questions into one prompt and answering them all at once.** That's a quiz, not an interview. Real interviews are conversational. Ask Qwen to go one question at a time so you can answer, get feedback, then continue.
- **Stopping after one round of feedback.** The first answer is rarely the best. Ask Qwen to ask the same question again with a follow-up, or to push back, or to make the question harder. Two or three rounds of pressure is where confidence actually builds.
- **Forgetting that Qwen can't read body language or tone.** Qwen evaluates what you *type*, not how you *sound*. If you're prepping for a video interview, do at least one practice round on camera or out loud, separately, to catch the things text can't show.

What this means for you

- **In everyday life:** Even outside job hunting, this same trick works for any high-stakes conversation — a parent-teacher meeting, a price negotiation, a tough email you're about to send. Practice the words before you have to say them.
- **For your business or work:** If you hire people or run interviews yourself, you can flip the script — ask Qwen to generate the five most common weak answers to a question you always ask, so you know what to listen for when you're the one in the chair.
- **If you're just getting started:** This is one of the most useful first things to try with any AI assistant. No technical setup, no cost to begin, and you'll immediately see why "AI as a practice partner" beats "AI as a cheat sheet."

Try it now

Open Qwen right now and run a single 60-second drill: ask it, "Give me one common interview question for a [your job title] role." Read the question, close the laptop, and answer it out loud in your room. That's it — your first mock interview is done.

How to write a polite but firm complaint email that actually gets a reply

6 steps · Read online

> **In short:** You can ask any AI assistant — ChatGPT, Gemini, Claude, Copilot, Le Chat, or another — to draft a complaint email that stays calm, lists the facts, and ends with one clear ask. The trick is giving it the right facts and tone instructions; it can write the draft in seconds, but you still need to read it aloud, cut it to one ask, and send it yourself.

You've had a bad experience with a company. Maybe a product arrived broken, a bill is wrong, or support stopped replying. You want to complain, but you don't want to sound rude — and you really want a reply this time. An AI assistant can help you write a message that's hard to ignore: short, factual, polite, and impossible to misread. Here's how to do it in about 15 minutes.

Before you start

- **An AI assistant you can type into.** Any major one works (ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek). A free account is enough for this task.
- **A real complaint in mind.** Pick one you actually need to send — vague practice complaints don't teach you much.
- **About 15 minutes.** Most of that time is reading the draft out loud and editing, not waiting for the AI.

One honest note before we start: the assistant can draft the email, but it can't send it for you, and it can't guarantee a reply. What it *can* do is remove the hardest part — finding the right words when you're frustrated.

Step 1 — Write down the five facts a rep will need

Before you open any AI tool, get the facts out of your head and onto paper (or a notes app). The biggest reason complaint emails get ignored is that they're vague. A sentence like "Your service is terrible" gives a customer-service rep nothing to act on.

Open a blank note and answer five questions in one or two short sentences each: (1) What did you buy or sign up for, and when? (2) What went wrong? (3) What have you already tried to fix it? (4) What do you want them to do now? (5) What's your deadline for a response?

Example: "Bought a blue wool coat from Example Outfitters on March 4 (order #1234). It arrived March 10 with a torn lining. Emailed support on March 11, no reply. I'd like a replacement shipped by March 25, or a full refund."

You'll know it worked when you can read the note out loud in under 30 seconds and someone else would understand what happened without asking follow-up questions.

Step 2 — Open your AI assistant and start a fresh chat

Open the AI assistant you normally use, on whatever device is handy — phone, tablet, or computer. When the chat window opens, start a new conversation rather than continuing an old one. A fresh chat gives the assistant a clean slate, so it won't accidentally mix in tone or facts from a previous topic.

Look for a button like a **"New chat"**, **"+"**, or **pencil icon** near the top or side of the screen. On some apps it lives inside a menu (often three dots or three lines). If you can't find it, just refresh the page or close and reopen the app — most assistants start you on a blank conversation by default.

Example: type "I want to draft a complaint email" as your first line, just to confirm the chat is responding.

You'll know it worked when you see a blank message box at the bottom of the screen and the assistant is ready for your prompt.

Step 3 — Paste a prompt that names the facts, the tone, and the ask

Now copy the five facts from Step 1 and paste them into the chat, followed by a short set of instructions. The clearer your instructions, the closer the first draft will be to something you'd actually send.

Type something like the prompt below — adjust the wording to your situation, but keep the same four pieces: facts, tone, length, and ask.

Example: "Please draft a polite but firm complaint email. Facts: [paste your five facts here]. Tone: calm, professional, no sarcasm, no insults. Length: under 150 words. End with one specific ask and a reasonable deadline of 5 business days. Do not invent any order numbers, dates, or names I didn't give you."

You'll know it worked when the assistant produces a complete email — greeting, facts, what you want, deadline — in one block of text that reads like a real human wrote it. If the first draft comes out too long, too stiff, or too emotional, that's normal; you'll fix that in the next step.

Step 4 — Read the draft out loud and tame the tone

Don't send the first draft yet. Read it out loud, slowly. Your ear catches things your eyes miss — awkward phrases, sentences that sound angrier than you feel, or phrases that sound so polite they sound sarcastic.

If a sentence makes you wince, copy that one sentence back into the chat and ask the assistant to rewrite it in a calmer voice. Don't throw the whole draft away; most of it is probably fine. Common rewrites you'll want: "I'm extremely disappointed" becomes "I'd like this resolved"; "This is unacceptable" becomes "I'd like to understand what happened"; "You need to fix this immediately" becomes "I'd appreciate a response by [date]."

If the draft came out too apologetic — sounding like you're asking permission to complain — say so directly. 💬 Example: "Make this sound more confident. I'm being reasonable, but I expect a real answer."

You'll know it worked when you can read the email aloud without rolling your eyes, and a friend would describe the tone as "firm but polite" — not angry, not groveling.

Step 5 — Cut the email down to one clear ask

A complaint email that asks for three different things often gets three different "we'll look into it" replies and nothing actually happens. Read the draft and circle the single most important thing you want.

If the email mentions a refund, a replacement, an apology, and a discount, pick the one that matters most and cut the others. You can always bring up the others in a later email if the first ask is refused. Then paste the draft back into the chat and ask for the trim.

💬 Example: "Shorten this to one ask: I want a replacement coat shipped by March 25. Remove the part about the discount."

You'll know it worked when the last paragraph of the email contains exactly one sentence that begins with "I'd like" or "Please" and ends with a specific date. Everything else in the email is supporting context, not a wish list.

Step 6 — Send it, save the thread, and set a follow-up reminder

Copy the final draft into your real email app, fill in the real order number, your name, and any attachments (photos of the damage, receipts, screenshots of the original chat). Read it one last time to make sure no detail was lost in the copy-paste — small errors like a wrong date can undo an otherwise strong email.

Before you hit send, take a screenshot of the email and save the AI chat somewhere you can find it. If you have to follow up later, you'll have a record of exactly what you asked for. Then set a calendar reminder for two business days after your stated deadline.

💬 Example follow-up: "Following up on my email below — I haven't received a response. I'd still like a replacement coat shipped by March 25. Please let me know either way."

You'll know it worked when the email is sitting in your sent folder, you have a date on your calendar for the follow-up, and you feel — for the first time in this whole mess — that you've said exactly what you needed to say.

Common mistakes

- **Letting the AI invent facts.** If you say "I ordered something last month," the assistant may invent an order number or date that sounds real but isn't. Always paste your *own* facts, and add the line "Do not invent any details I didn't give you."
- **Sending the first draft without reading it.** AI drafts are usually 80% right and 20% wrong in tone. The 20% is what makes an email sound robotic or slightly off. Always read it out loud.
- **Asking for too many things at once.** "I want a refund, a replacement, an apology, and a discount" usually gets you nothing. Pick one ask, and bring up the others later if needed.
- **Copying a tone that doesn't sound like you.** If the draft sounds too stiff or too formal, paste in a sentence you've actually written before — a text to a friend, an older email — and ask the assistant to match that voice.

What this means for you

- **In everyday life:** The next time a delivery is wrong, a bill is off, or a service lets you down, you'll have a way to complain that takes 15 minutes instead of an evening of staring at a blank screen.
- **For your business or work:** The same method works for vendor disputes, late invoices, or client complaints. A calm, one-ask email is also the right tone when you're escalating an issue to a manager — you sound professional, not emotional.
- **If you're just getting started:** The very first email to draft this way is the hardest. After that, you'll have a prompt template you can reuse with new facts pasted in.

Try it now

Open your AI assistant, paste this prompt, and fill in the five facts from Step 1: *"Draft a polite but firm complaint email. Facts: [your facts]. Tone: calm and professional. Length: under 150 words. End with one specific ask and a 5-business-day deadline. Do not invent any details."* Read the result out loud, trim it to one ask, and you'll have a sendable email in under 10 minutes.

Translate and Explain a Confusing Email or Letter with Perplexity

6 steps · Read online

> **In short:** Open Perplexity, paste the confusing message, and ask it to translate, summarize, and draft a reply. Free to use on the web; Pro adds longer messages and file uploads. It can draft a reply, but it cannot send the email for you.

That long email from your landlord, the tax letter in German, the polite-but-passive-aggressive message from a client — they all sit in your inbox looking scary. Perplexity can turn any of them into two things you actually need: a plain-English summary of what they want, and a calm suggested reply you can edit before sending.

Before you start

- A free Perplexity account (sign up with email, Google, or Apple) — works in any modern browser on phone or computer.
- The confusing message itself: either copied as text, or as a clear photo/screenshot (uploading may require Pro).
- About **5–8 minutes** for your first message, faster once you get the hang of it.
- Honest note: Perplexity is great at language and summarizing, but it's not a lawyer, accountant, or doctor. For official notices (taxes, contracts, legal threats), treat its summary as a starting point and confirm with a professional before you act.

Step 1 — Open Perplexity and start a new thread

Go to **perplexity.ai** in your browser (or open the Perplexity app if you have it installed). On the home page you'll see a big text box in the middle that says something like **"Ask anything..."** — that's where the conversation happens. Click or tap inside the box once. A blinking cursor appears, and the page is ready for your question. If you're already signed in, you'll also see a small **"New Thread"** or + button at the top-left or top-right — tap it so your message doesn't mix with an older chat. If the layout looks different (Perplexity updates its design often), look for a clearly empty text area with a **Send** button (paper-plane icon ✈ or arrow) at the end. 💬 Example: don't type yet — just leave the cursor blinking in the box. You'll know it worked when the text box is empty, your cursor is blinking inside it, and no earlier conversation is showing above.

Step 2 — Paste the confusing message into the chat

Open your email or letter in another tab, window, or app. Select **all the text** (Ctrl+A on Windows, Cmd+A on Mac, or long-press → Select All on a phone). Copy it (Ctrl+C / Cmd+C / Copy in the menu). Then tap back into Perplexity's text box and **paste** (Ctrl+V / Cmd+V / Paste). The full message lands inside the box as one block of text — you should see the words from the letter sitting there. If the message is very long, Perplexity may show a small character count or a warning; trim the greeting and signature if needed. If your screen looks different, you may also see a **paperclip icon** 📎 or an **upload button** — that's for attaching files or photos of letters instead of pasting text. 💬 Example: paste a paragraph like "Sehr geehrte Mieter, wir möchten Sie daran erinnern, dass die Miete für November bis spätestens zum dritten Werktag des Monats zu überweisen ist..." — the exact wording matters, so paste the real thing. You'll know it worked when the entire message shows up in the box, no strange symbols, and you can still see the send button (✈ or arrow) at the end.

Step 3 — Add a clear instruction below the pasted text

Don't hit send yet. Press **Enter** to make a new line inside the same box, then type what you want Perplexity to do with the message. Be specific — a vague "what does this mean?" gets a vague answer. The clearer the instruction, the more useful the reply. If you only see a single-line text box (some mobile views hide multi-line), look for a small **expand** icon or a + that opens a fuller composer. On the free plan you can usually ask anything; Pro users get longer pasted messages and file uploads — check the current plan details on Perplexity's pricing page if you're unsure which you have. 💬 Example: "Please translate this email into plain English, list the three things the sender wants me to do, and tell me how urgent it sounds." You'll know it worked when the box contains both the pasted message **and** your instruction underneath, and you can read both clearly before sending.

Step 4 — Send and read the summary

Tap the **Send** button (✈ paper-plane icon, or an arrow ↵ at the end of the box). A short "thinking" animation appears — Perplexity is reading your message and writing its answer. After a few seconds, a reply appears under your message, with the translation at the top, then the bullet list of what the sender wants, then the urgency note. Skim it slowly: any word you don't understand, highlight and copy, then ask a follow-up like "What does [word] mean in everyday language?" If the answer doesn't quite hit the mark, click the small **Regenerate** button (often a circular arrow ↻ near the reply) to get a fresh attempt — sometimes the second version is clearer. If the screen looks different on your device, look for a **Rewrite**, **Try again**, or **New answer** button near the bottom of the reply. 💬 Example follow-up: "Can you explain 'bis spätestens zum dritten Werktag' like I'm 10 years old?" You'll know it worked when you see a clear English version of the email, a short list of what the sender wants, and a plain-language sense of how urgent it really is.

Step 5 — Ask for a suggested reply in the same chat

Stay in the **same conversation** (don't start a new thread) so Perplexity remembers the original message and its summary. In the text box at the bottom, type a request for a draft reply. Be polite about tone — Perplexity matches the energy you ask for. The reply will appear below your message, ready for you to copy. If you want a different vibe (firmer, more apologetic, more formal), just ask in a follow-up — it's faster than starting over. Free users can do this as many times as they want; Pro only matters if you need to paste very long messages or upload PDFs. 💬 Example: "Now write a short, polite reply in English confirming I'll pay by Friday and asking for a receipt." You'll know it worked when a reply draft appears that you could realistically send after light editing — and you can copy it with a single tap or click.

Step 6 — Copy, edit, and send the reply yourself

Highlight Perplexity's suggested reply and **copy** it (Ctrl+C / Cmd+C / long-press → Copy). Open your email app, start a new message, **paste** the draft into the body, and then **edit it in your own voice** — add your name, your real deadline, any details only you know. Perplexity can draft, but it cannot send email or access your inbox, so this last step is always yours. Read the whole thing once out loud before clicking send; if any sentence sounds weird in your voice, change it. If your interface looks different, look for a **Copy** icon (two overlapping squares 📋) at the corner of the reply, or simply long-press the reply on a phone to bring up the copy menu. 💬 Example edit: change "I will remit payment by Friday" to "I'll send the payment by Friday — thanks for the reminder, Maria." You'll know it worked when you've sent a clear, calm reply that you actually wrote (with Perplexity's help), and the original confusing email no longer feels scary.

Common mistakes

- **Pasting without an instruction.** If you only paste the email and hit send, Perplexity may summarize it but won't translate or draft a reply. Always add a short instruction line under the paste — it changes the result completely.
- **Trusting the summary on legal or medical letters.** Perplexity can miss legal nuance. For tax notices, court letters, or anything health-related, treat its plain-English version as a heads-up, not a final answer — confirm with the right professional before paying or signing anything.
- **Sending the draft without editing.** AI-written text sometimes has odd phrasing, wrong names, or invented details. Always read the whole reply once, fix anything that doesn't sound like you, and remove anything that isn't true.

What this means for you

- **In everyday life:** the next time a long email from your bank, landlord, school, or insurance company lands in your inbox, you'll have a two-minute way to understand it and reply without stress — even if it arrives in another language.
- **For your business or work:** if clients or colleagues send you formal letters, contract updates, or foreign-language requests, you can draft a professional reply in the right tone in seconds, freeing you to focus on the actual decision rather than the wording.
- **If you're just getting started:** the easiest first move is to pick one email you've been ignoring because it felt too long or too formal, paste it into Perplexity, and ask for a three-bullet summary in plain English — that's it.

Try it now

Open perplexity.ai, find one email you've been avoiding, paste it into the box, and type underneath: *"Translate this into plain English and give me three bullet points of what the sender actually wants from me."* Send it, read the result, and notice how much lighter the email feels five minutes later.

How to Use Gemini Deep Research for Thorough Reports

6 steps · Read online

Imagine having a tireless personal research assistant who can scour dozens of websites, cross-reference facts, and write a polished report in minutes. This guide will show you how to use Gemini's powerful capabilities to gather deep insights for work, study, or personal projects, just like a professional researcher.

Before you start

- A Google account: You'll need a standard Google account to access Gemini.
- Access to Gemini: Ensure you can open Gemini in your web browser (gemini.google.com) or use the Gemini app on your mobile device. Some advanced features might perform better or offer more depth with a Gemini Advanced subscription, but the core research principles apply to all versions.
- A device: A computer, laptop, tablet, or smartphone with internet access.
- A clear topic or question: Have a specific subject in mind that you want Gemini to research.
- Rough total time: This process usually takes 5-10 minutes, depending on the complexity of your research topic and how many follow-up questions you ask.

Step 1 — Start a new chat and enable web search

First, let's open Gemini and make sure it has permission to search the internet. An AI (Artificial Intelligence) tool like Gemini is powered by an LLM (Large Language Model) – think of it as the super-smart brain that helps the AI understand and generate human-like text. To perform deep research, this brain needs access to the wider internet.

When you open Gemini, you'll see a clear text box at the bottom of the screen where you type your messages. Look towards the top-right or top-center of the screen for a toggle switch or button that usually says "Extensions" or shows a puzzle piece icon. Clicking this often reveals a list of tools, including "Google Workspace" and "Google Search." Ensure the "Google Search" extension is switched 'on'.

If it looks different: On some devices or versions, the "Extensions" might already be active or there might not be a visible toggle. If so, don't worry, Gemini often suggests using Google Search automatically when you ask a research question. Just ensure you're starting a fresh conversation, usually by clicking a "New Chat" button (often a '+' icon) in the side menu.

You'll know it worked when: The toggle for "Google Search" appears switched on, usually indicated by a blue or green color, or Gemini doesn't prompt you to enable search when you type your first question.

Step 2 — Craft a detailed research prompt

A prompt is simply the instruction you give to an AI. To get a high-quality, in-depth report, you need to be very clear with your prompt. Don't just ask a one-sentence question. Instead, tell Gemini who the report is for, what specific areas to cover, and how you want the final document organized. Think of it like giving instructions to a new research assistant.

Type your detailed instructions into the text box. The more detail you provide, the better Gemini can understand your needs and focus its search.

Example:

> "Act as an expert researcher for a homeowner. Please conduct a deep, multi-step search of the web to investigate current trends in home solar battery storage for residential properties in the US. Focus on average cost ranges, typical lifespan, available government rebates or incentives, and environmentally friendly recycling options. Organize the report with clear headings, bullet points, and provide specific examples where possible. Include links to your sources."

You'll know it worked when: Gemini acknowledges your request (often by saying "Okay, I can help with that...") and begins to indicate it's searching the web or preparing to ask clarifying questions.

Step 3 — Answer Gemini's clarifying questions

One of the cleverest parts of Gemini's deep research process is that it doesn't just guess what you want. It will often reply with two or three clarifying questions to narrow down the scope of its search. This is Gemini using its "brain" (the LLM) to ensure it understands your request perfectly, saving you time later.

Read Gemini's questions carefully. These might be about geographical focus, specific types of technology, or the target audience for the report. Type your concise answers in the text box below.

If it looks different: Gemini might sometimes jump straight into generating the report if your initial prompt was extremely clear. If it doesn't ask questions, that's fine; simply move to the next step.

Example: If Gemini asks, "Are you interested in global statistics or a specific country for the solar battery research?" you might reply:

> "Focus primarily on American residential properties, especially those in suburban and regional areas, not just cities."

You'll know it worked when: Gemini confirms it has understood your answers, perhaps by saying "Got it," and then transitions to the next stage of gathering information.

Step 4 — Observe Gemini gathering information

Once you've clarified the details, Gemini will begin its thorough search. You'll often see a visual progress bar or a series of messages indicating what the AI is doing in real-time. These messages might pop up within the chat window, showing stages like "Searching for battery lifespans" or "Analyzing recycling policies." This is Gemini actively using its Google Search extension to find and process information from various websites, much like a human researcher would.

This step takes longer than a standard AI reply—often up to a minute or two—because the assistant is reading actual web pages rather than relying purely on its pre-existing training memory.

If it looks different: Sometimes, the progress indicators might be subtle, or Gemini might just show a spinning wheel. As long as the chat isn't frozen and you see some activity, it's working.

You'll know it worked when: You see clear messages indicating Gemini is searching the web ("Using Google Search...", "Searching for..."), followed by text beginning to appear that resembles your requested report structure.

Step 5 — Review and verify the generated report

When the report is ready, Gemini will present a structured document, usually complete with clear headings, bullet points, and, crucially, citations or links to the websites it used. This is your chance to put on your detective hat.

Before you copy the text for use, perform a quick spot-check. Click on a few of the provided links to make sure they are active and support the claims made in the report. This protects you against "hallucination," which is the term for when an AI confidently invents facts or links that do not actually exist. While rare with Gemini's search capabilities, it's always wise to verify important details.

If it looks different: Ensure you scroll to the end of the report. The citations are often small numbers or footnotes. If you don't see any, you can politely ask Gemini: "Can you provide the source links for this report?"

You'll know it worked when: You've clicked on a few links and confirmed that the information in the report aligns with the content on the source websites.

Step 6 — Refine or ask follow-up questions

The first report Gemini provides is a great starting point, but you might want to dive deeper or change the focus slightly. This is where the interactive nature of AI truly shines. You can continue the conversation, asking for more details on a specific section, requesting a different format, or even asking Gemini to summarize complex points.

Type your follow-up questions or requests into the text box. This allows you to iteratively improve the research until it perfectly meets your needs.

Example:

> "Can you expand on the environmental impact of lithium-ion battery production? Also, please summarize the key points of this report in three bullet points, suitable for a quick presentation slide."

You'll know it worked when: Gemini provides new information, rephrases existing content, or summarizes the report based on your latest instructions, showing its ability to build on previous responses.

Common mistakes

- **Being too broad:** Asking for a report on "the history of technology" is too vast for a single deep dive. *The fix:* Narrow your focus to a specific decade, tool, or industry, or break it into smaller prompts.
- **Skipping the fact-check:** Never copy and paste a report straight into an important document or project without clicking the source links to verify the numbers and claims. *The fix:* Always dedicate a few minutes to spot-checking sources and ensuring accuracy.
- **Ignoring clarifying questions:** If Gemini asks for clarification, don't skip it or give a vague answer. Your response is crucial for guiding the AI to find the exact data you need, rather than general information. *The fix:* Take a moment to provide a clear, concise answer to each of Gemini's questions.

What this means for you

- **In everyday life:** You can quickly compare complex options

How to Draft a Polite Follow-up Email After a Job Interview

6 steps · Read online

> **In short:** Send a short, specific thank-you email within 24 hours of the interview — three or four sentences, one fresh detail from the chat, and a clear next step. Keep it warm, not needy; this guide shows you exactly what to write and what to avoid.

A good follow-up email is one of the few small things still fully in your control after an interview. It takes about ten minutes to write, and it quietly reminds the hiring manager who you are. Below is a simple, friendly process you can follow the same day, even if you've just walked out of the meeting room still buzzing.

Before you start

- **Time needed:** roughly 10–15 minutes, ideally within 24 hours of the interview.
- **What you need:** the interviewer's name, their email address, and the job title you applied for. A notes app (your phone is fine) where you scribbled two or three things that came up in the conversation.
- **Tool:** any email app or webmail — Gmail, Outlook, Apple Mail, whatever you already use. No special software required.
- **Optional:** an AI assistant (such as ChatGPT, Gemini or Claude) can help you polish your draft, but it cannot send the email for you, and it doesn't know what was actually said in the interview — that's your job to feed in.

Step 1 — Jot down three things you actually talked about

Open your notes app the moment you can after the interview — even two minutes in the parking lot is better than waiting. Write down three specific things that came up: a question they asked, a story you told, a detail about the role or the team. These are the bricks you'll build the email from. **What you see:** your phone's notes app, the date at the top, and probably a few messy half-sentences. **What happens next:** you'll have raw material to pull from, so the email can sound like *you* and not like a template you copied off the internet. **If it looks different:** if you didn't take notes, try to recall one question you were asked and one thing you said that landed well. Even one believable detail is enough. **You'll know it worked when** you have at least two concrete points you can reference in the email — something the interviewer would recognize as "ours, not just any interview." 💬 Example: *Asked about handling an angry customer → I told the café story. Mentioned the new Nashville office opens in March.*

Step 2 — Open a new email and set the subject line

Open your email app and click **Compose** (or **New Message**). For the subject line, keep it short and easy to scan — the interviewer gets a lot of mail. A clear subject line is like a polite wave from across a crowded room: it says "this is the person you met today, not a stranger." **What you see:** a blank email with a To, Subject and message area. **What happens next:** the moment you type their name correctly, the email already feels more personal than half the inbox they'll skim today. **If it looks different:** on some apps the button says **New email**, **Write**, or shows a pencil icon — any of these will do. **You'll know it worked when** the subject clearly says thank-you, names the role, and is short enough to read fully on a phone screen. 💬 Example: `Subject: Thank you – Marketing Coordinator interview`

Step 3 — Open with a warm, specific thank-you

In the first one or two sentences, thank the interviewer by name and mention something specific from your conversation. This is the part that turns a generic "thanks for the interview" into a real, human note. **What you see:** the cursor blinking at the top of the message body. **What happens next:** the reader instantly knows this email is about them and the meeting they actually had — not a copy-paste sent to ten companies. **If it looks different:** if you interviewed with a panel, address the lead person by name and use "and the team" for the rest. **You'll know it worked when** the first paragraph would only make sense to someone who was in that room. 💬 Example: *Hi Sarah, thank you so much for taking the time to chat with me this morning about the Marketing Coordinator role. I really enjoyed hearing how the team is planning the March launch.*

Step 4 — Briefly reinforce why you're a fit

Now add two or three sentences that gently remind them of your strengths, tied to something you actually discussed. This is where your notes from Step 1 come in. Think of it like adding one fresh photo to a friend's birthday card — small, but it shows you remembered. **What you see:** the email is taking shape — it's getting longer than a text but shorter than a letter. **What happens next:** the interviewer gets a quick, evidence-based recap of why you belong in the role, without you having to repeat your whole CV. **If it looks different:** if the conversation was mostly about the company or the role itself, you can skip this part and go straight to Step 5 — a thoughtful question can carry the same weight. **You'll know it worked when** the reader can see, in plain words, that you understood the role and can do it. 💬 Example: *Our chat about managing the café's social media during the weekend rush reminded me how much I enjoy fast-paced, customer-facing work — that's exactly the energy I'd love to bring to your team.*

Step 5 — Close with a clear, low-pressure next step

End with one short sentence that says you're looking forward to hearing from them, and sign off politely. Avoid phrases like "I really need this job" or "Please let me know ASAP" — they add pressure, and pressure is the opposite of what you want. **What you see:** the email might be around 80–120 words in total. That's the right length. **What happens next:** the email lands in their inbox as a calm, professional note — easy to reply to, easy to forward, easy to remember. **If it looks different:** if you don't know the next steps (some processes are slow), just say "I'm happy to provide anything else you need" — it's honest and useful. **You'll know it worked when** the email closes without begging, without bragging, and without leaving the interviewer unsure about what you want. 💬 Example: *I'm excited about the role and would love to hear how the next steps look. Thanks again, Sarah — have a lovely rest of your week.*

Step 6 — Do a quick proofread and hit send

Read the whole email out loud once, slowly. Check the interviewer's name is spelled right, the role title matches the job ad, and there are no obvious typos. Then send. **What you see:** the final email, maybe 100 words on screen. **What happens next:** the email is on its way. Most of the work is done. **If it looks different:** if you spot a mistake after sending, don't panic — a small typo is not worth a follow-up apology. Save that energy for the next step in the process. **You'll know it worked when** you can read it back and think, "yes, that sounds like me on a good day." 💬 Example (final email): *Hi Sarah, thank you so much for taking the time to chat with me this morning about the Marketing Coordinator role. I especially enjoyed hearing how the team is planning the March launch. Our conversation about handling busy weekend rushes reminded me how much I love fast-paced, customer-facing work, and I'd be thrilled to bring that energy to Wynnum Co. I'm excited about the role and would love to hear how the next steps look. Warm regards, Jordan.*

Common mistakes

- **Sending one identical email to everyone.** If you can't tell which interviewer it was for, they can tell you didn't write it. Always swap at least one specific detail.
- **Writing five paragraphs.** A follow-up is a thank-you card, not a second cover letter. If it scrolls on a phone, it's too long.
- **Sounding desperate or pushy.** Lines like "I really need this opportunity" or "Any update would be appreciated" add pressure. Keep the tone calm and confident.
- **Forgetting to proofread the name.** Getting the interviewer's name wrong undoes the whole "personal" effect. Double-check before you send.

What this means for you

- **In everyday life:** The same habit — a short, specific, timely message — works after a coffee chat, a school meeting, a meeting with a tradie, or a volunteering shift. It makes you the person people remember kindly.
- **For your work or job search:** A clean follow-up email slightly tilts the odds in your favor because most candidates don't bother, and many who do bother send something generic. Yours doesn't have to be brilliant — it just has to be real.
- **If you're just getting started:** Don't try to write the perfect email. Write a rough one, read it out loud, and send it within 24 hours. Done beats perfect every time.

Try it now

Within the next five minutes, open your notes app and write down three things you remember from your last interview (or your last important meeting). That's the hard part done — the email writes itself from there.

How to practice for a job interview with an AI mock interviewer tonight

6 steps · Read online

> **In short:** You can run a realistic mock interview with any major AI assistant tonight by giving it the job description, asking it to play interviewer, and stress-testing one tough answer at a time. It won't replace a real human, but it will sharpen what you actually say.

Job interviews don't reward memorization — they reward clarity under pressure. Tonight you can build that clarity with a practice partner that already lives in your phone or browser: an AI assistant. The trick isn't to ask random questions. It's to run a short, repeatable session that stress-tests your real answers before a human ever hears them. Below is a method you can use tonight, and again the night before any interview.

Before you start

- An account with any major AI assistant — **ChatGPT, Claude, Gemini, Microsoft Copilot, and Le Chat** all work for this. Use whichever you already have; don't sign up for a new one tonight.
- The job description for the role. Copy-paste the whole posting, even if it's long — the AI uses every line.
- A quiet 30–45 minutes, ideally with headphones and a microphone so you can hear yourself.
- Optional: a notes app open on your phone so you can save weak spots at the end.

Step 1 — Open a fresh chat with your AI assistant

Open the AI assistant of your choice in a browser or app and start a brand-new conversation. What you'll see is a blank chat box — no prior context, no leftover questions from yesterday. This matters because the assistant draws on the whole chat thread when it replies, so a clean slate means a clean practice session. In ChatGPT, look for a top-left button that says **New chat**; in Claude, Gemini, and Copilot, look for **New chat** or a + icon near the conversation list. On mobile, you may need to swipe to your conversation history first before the button appears. 💬 Example: just open the chat — no typing yet. You'll know it worked when you see an empty message box with a placeholder like "Message…" or "Ask anything."

Step 2 — Tell the AI exactly what role you're practicing for

Paste the job description into the chat, then add two or three lines about your background. What you'll see is a long message appear, followed by the assistant's reply — usually a short acknowledgment like "Got it, here's how I'll approach this." The assistant now has the context it needs to ask realistic questions instead of generic ones. If your tool has a **Custom instructions** or **Memory** setting, you can paste the role there so it persists across chats; otherwise just paste it into the first message of this session. On Claude and ChatGPT those settings live under your profile picture; on Gemini they're under **Settings**. 💬 Example: "I'm preparing for a Senior Customer Success interview at a B2B SaaS company. Here is the job description: [paste]. My background: 4 years in support, promoted once, led a small team. Please ask me one question at a time, like a real interviewer would." You'll know it worked when the assistant references your actual role or company in its next reply — not a generic "tell me about yourself."

Step 3 — Start with the one question you find hardest

Ask the assistant to ask you the single question you're most worried about — usually a behavioral or scenario question, not the opener. What you'll see is the assistant posing it the way a real interviewer would: *"Tell me about a time you dealt with a difficult customer."* Read the question, then **answer out loud** — even if you're typing, speak the answer aloud so you catch filler words, awkward pauses, and rambling. This is the most underused part of practice, and the part that actually changes how you sound in the room. If your assistant has a voice mode (ChatGPT and Gemini both do on desktop and mobile), switch to it now; otherwise read your typed answer out loud to yourself once you're done. 💬 Example: "Ask me one question at a time. Start with: 'Tell me about a time you had to win back an unhappy customer.'" You'll know it worked when you've actually spoken a full answer of roughly 60–90 seconds — not just thought one.

Step 4 — Ask for specific, structured feedback

Once you've answered, prompt the assistant to score your reply and point to one improvement. What you'll see is a short rubric — usually a 1-to-5 rating on clarity, specificity, or impact, plus one thing to tighten. This is where generic praise is useless, so ask for criticism directly; the assistant now becomes a coach instead of a yes-machine. If it forgets and just says "great answer!", tell it again to be tougher — most assistants default to polite, especially in voice mode. On the free plan of all the assistants listed above, text feedback is unlimited; voice mode may have a small daily cap. 💬 Example: "Rate my answer from 1 to 5 on clarity, specificity, and impact. Give me one strength and one specific thing to improve. Be honest — I'd rather hear it from you than from an interviewer." You'll know it worked when you can name one concrete sentence or word you'd change before tomorrow.

Step 5 — Run a five-question dress rehearsal

Once your warm-up answer feels solid, ask the assistant to run a full short interview. What you'll see is it moving through five questions in sequence — one at a time, waiting for your reply before continuing. Treat this like the real thing: no peeking at notes, no redoing an answer, and

try to keep each answer between 60 and 90 seconds. If your assistant has voice mode, use it for this round; speaking back-to-back is closer to the actual interview than typing. Set a phone timer for 10 minutes so you don't run over. If the screen shows a **Stop** or **Regenerate** button, ignore both and just answer naturally. 💬 Example: "Now run a 5-question mock interview, one question at a time. Wait for my answer before moving on. Mix behavioral and role-specific questions." You'll know it worked when you've answered five questions back-to-back out loud and feel a bit tired — that's the real interview muscle kicking in.

Step 6 — Capture your three weakest spots before you close the chat

End the session by asking the assistant to summarize the patterns it noticed — vague answers, missing numbers, weak openings. What you'll see is a short list of two to four patterns, plus the specific questions that exposed them. Screenshot it or copy it into your notes app before you close the chat, because tomorrow you'll want to practice those exact spots. This single summary is what turns a random practice session into a repeatable method: each night you run the same drill, you knock down one more weak pattern. Be aware — the assistant may flatter you here too, so push for the harshest pattern if its list feels too soft. 💬 Example: "Based on this whole session, list the three patterns that hurt my answers most, and the one question I should drill tomorrow." You'll know it worked when you have a 3-item checklist saved somewhere outside the chat.

Common mistakes

- **Reading your answer instead of saying it.** Typing an answer and reading it back to yourself feels safe but doesn't catch the rambling, filler words, or long pauses a real interviewer hears. Speak every answer out loud, even if it feels awkward.
- **Letting the assistant be too polite.** Most AI assistants default to "great answer!" — which is useless for practice. If your feedback feels like a hug, tell the assistant directly to score you from 1 to 5 and give one specific improvement.
- **Practicing only "tell me about yourself."** That opener is rarely the question that decides the interview. Spend at least half your session on the two or three harder behavioral and scenario questions for the actual role — that's where the real anxiety hides.

What this means for you

- **In everyday life:** If you're between jobs or thinking about a change, this same loop works for promotion conversations, salary reviews, or even asking a landlord to fix something — anywhere a clear, structured answer matters more than a clever one.
- **For your business or work:** Use this method to rehearse client pitches, performance reviews, or tough conversations with a colleague. The "AI asks, you answer out loud, get one specific critique" pattern generalizes to almost any high-stakes talk.
- **If you're just getting started:** Don't try to do all six steps tonight. Do Step 1, paste the job description, and answer one tough question out loud. That's already more practice than most candidates do.

Try it now

Open your AI assistant right now, paste the job description into a fresh chat, and type: *"Ask me one question at a time. Start with the hardest behavioral question for a [role] interview."* Then answer it out loud. That's the whole loop in two minutes — and the easiest way to find out where you actually stand tonight.

Polish Your Resume Bullet Points with Le Chat in 10 Minutes

6 steps · Read online

> **In short:** Paste a plain description of what you did at work into Le Chat, ask it to rewrite that as a resume bullet, and you'll get back a sharper, more impressive line in about ten minutes — even if writing isn't your strongest skill.

Most people describe their work the way they talk in real life: short, plain, missing the impressive parts. Recruiters scan your resume in seconds, though, and plain sentences get skipped. Le Chat — Mistral's AI assistant — can take your rough notes and turn them into the kind of bullet point a hiring manager actually pauses on. You stay in charge of the truth; Le Chat just helps with the wording.

Before you start

- A free Le Chat account (sign up with an email at Mistral's chat page).
- A short, rough description of one thing you did at a job — even a single messy sentence is fine.
- A web browser on a phone, tablet, or computer.
- About ten minutes, and one promise to yourself: only put **real** numbers on your resume. Le Chat can suggest placeholders, but the final figures must come from your actual experience.

Step 1 — Open Le Chat and start a fresh chat

Go to Le Chat in your browser and sign in if you aren't already. Once the page loads, look for a button labeled something like **New chat** or **Start new conversation** — usually at the top of the chat list on the left side of the screen, or shown as a plus icon. Tap or click it. A blank conversation opens, with an empty message box near the bottom often labeled something like **Message Le Chat** or **Ask anything**, though yours may use different words.

If you can't find the new chat button, look for a **pencil icon**, a +, or a small **chat bubble with a plus sign** — those are the most common shapes. On a phone, the option is often tucked inside a **menu icon** (three lines or three dots) at the top-left.

You'll know it worked when you see a clean, empty chat screen with a blinking cursor in the message box and no previous messages in the list.

Step 2 — Paste your rough description into the chat box

Click inside the message box at the bottom of the screen and type, or paste, one short sentence describing what you actually did at work. Don't worry about making it sound impressive yet — plain language is exactly what you want. You can write it the way you'd explain your job to a friend over coffee.

The text appears in the box as you type. If you paste from a notes app, you may see a short preview of how many characters you've added. Once you're happy with the description, you're ready for the next step — but don't hit send just yet.

Example: *"I helped customers with their bills and account questions at the front desk of a phone store for about two years."*

You'll know it worked when your sentence sits in the message box, ready to send, and the cursor is still blinking at the end.

Step 3 — Ask Le Chat to rewrite it as a resume bullet

Now add a short instruction below your description telling Le Chat exactly what you want, then hit **Enter** or the **send** icon — usually a small paper airplane or upward arrow on the right edge of the text box. Le Chat reads your description plus your instruction and replies with one or more polished bullet points.

If your app has separate buttons for **send** versus **new line**, look for the send icon — it's almost always on the right side of the message box. On a phone, send often lives at the far right of the keyboard area rather than as a separate button.

Example: full message to send — *"I helped customers with their bills and account questions at the front desk of a phone store for about two years. Please rewrite this as 3 strong resume bullet points a recruiter would notice. Start each one with a strong action verb."*

You'll know it worked when Le Chat's reply appears below your message with three (or more) bullet lines, each starting with a verb like *Resolved*, *Assisted*, *Handled*, or *Supported*.

Step 4 — Ask for a version with measurable impact

Strong resume bullets almost always include a number — how many, how much, how often. Send a short follow-up message asking Le Chat to add measurable impact, but with one clear honesty rule: tell it to use placeholders like **[X%]** or **[# of customers]** for anything you haven't actually measured. That way you don't accidentally paste in fake results that sound true.

Le Chat will write a new version of the bullets with placeholder numbers in square brackets. If the reply doesn't include placeholders, send a quick follow-up: *"Please show numbers in brackets like [X%] so I can fill in real numbers from my own records."*

Example: *"Now give me a stronger version with measurable impact. Put any numbers you don't know in brackets like [X%] or [# of customers] — I will fill in the real numbers myself."*

You'll know it worked when the new reply contains bullet points that include either real numbers or clear placeholders in brackets, never bare made-up figures you could mistake for real ones.

Step 5 — Pick your favorite and ask for one more variation

Read Le Chat's reply and pick the bullet that sounds closest to your real experience. Copy that one. Then send a short follow-up asking Le Chat to tighten the wording — make it shorter, swap a weak verb for a stronger one, or combine two ideas into one line.

If two bullets both sound good, ask Le Chat to merge them. If none of them sound like you, ask for a version with a different tone (for example, "more formal" or "more friendly"). Le Chat will respond with a fresh draft you can compare against the earlier one.

Example: *"I like bullet #2 the best. Please make it slightly shorter and swap 'helped' for a stronger verb. Keep it to one line under 25 words."*

You'll know it worked when you can clearly see one bullet point that sounds like *you* on a strong day, fits on one line, and would make a recruiter pause for half a second longer than the rest of your CV.

Step 6 — Save the final line and repeat with your next bullet

Copy your finished bullet into your resume document. Add a tiny note to yourself, like an asterisk * next to it, so you remember this one came from a Le Chat draft and still wants a quick personal read-through. Then start the cycle again with the next plain sentence from your work history.

If you're not happy with the final line, send one more message asking Le Chat to **combine** your favorite two drafts into one bullet. The trick is to keep iterating in small steps rather than asking for a "perfect" version in one go — small requests give cleaner results.

Example (next round): *"Now do the same for this bullet: 'answered phones and took messages for a real estate office'. Rewrite as one strong resume line with placeholders for any numbers."*

You'll know it worked when you have one polished bullet sitting in your resume, you've starred it for review, and you've already sent a fresh prompt to Le Chat for the next one — momentum is what makes this faster than staring at a blank page.

Common mistakes

- **Letting Le Chat invent real numbers.** It may try to fill in plausible figures (like "20% increase") that sound true but aren't. Always use the bracket-placeholder trick from Step 4, then add only numbers you can prove.
- **Writing bullets that sound nothing like you.** If the final line uses five-dollar words or industry jargon you never use at work, ask Le Chat to **rewrite in your own everyday voice** — recruiters can spot fake-sounding CVs in seconds.
- **Doing all your bullets in one giant prompt.** Long prompts give vague replies. One plain sentence plus one clear instruction per bullet is faster, cleaner, and far easier to fix when something is off.

What this means for you

- **In everyday life:** If you're applying for a new job, switching fields, or returning to work after a break, this turns the hardest part of resume writing — describing yourself — into a ten-minute habit instead of an evening of dread.
- **For your business or work:** If you help others with hiring, mentoring, or coaching, you can teach this exact loop to job-seeking friends, students, or clients, and skip the part where they freeze in front of an empty Word document.
- **If you're just getting started:** Don't try to rewrite your whole resume today. Start with the single bullet you've been avoiding the most — the one job you never know how to describe — and run it through Steps 2 to 5. That's more than enough.

Try it now

Open Le Chat, paste the roughest sentence you can write about one thing you did at your last job, and add this line below it: *"Rewrite this as one strong resume bullet a recruiter would notice, with action verb, task, and result."* Send it. You'll have your first improved bullet before your coffee gets cold.

Pull Names, Emails, and Phone Numbers Out of Messy Text with Cohere

6 steps · Read online

> **In short:** Cohere's Playground can read a messy block of contact info and pull out every name, email, and phone number as a clean list. Sign up free, paste your text into the chat box, and ask for structured output. The model can still miss details, so always review the result before you trust it.

If you've ever stared at a forwarded "who's coming Friday" email trying to copy ten contact details one by one, this workflow is for you. Cohere is an AI company whose Command model is good at reading tangled text and pulling structured data — meaning data organized into clear categories — out of it. By the end of this guide, you'll have a five-minute routine that turns a paragraph of mixed-up contact info into a tidy list you can paste straight into a spreadsheet.

Before you start

- A free Cohere account. Sign-up needs an email address and is available on the Cohere website. The free trial comes with enough credits to run this workflow many times.
- A block of messy text containing at least five contacts. Good sources: a forwarded meeting email, a pasted conference attendee list, or a stack of email signatures from a recent thread.
- A web browser on a laptop or desktop. The Playground is much easier to use on a larger screen than on a phone.
- About 10 minutes of unhurried time. No coding or technical background is needed.

Step 1 — Create your Cohere account

Open the Cohere website in your browser and look for a **Sign Up** or **Get Started** button, usually in the top-right corner of the homepage. Click it and register with your email address, or pick Google or GitHub sign-in if those options are visible. After you confirm your email, Cohere drops you on a dashboard. This is your control center: it shows your credit balance (how many free uses you have left), any API keys (codes developers use to call Cohere from their own software — you can ignore these for this task), and links to the tools. If you already have an account, just sign in. If a "Tell us about yourself" survey appears, look for a small **Skip** link at the bottom to bypass it.

If your home screen shows API Keys or Billing first instead of a welcome page, you're still in the right place — use the left sidebar to find the rest of the tools.

You'll know it worked when you land on a page that shows your name in the top-right corner and a sidebar with at least one tool link.

Step 2 — Open the Playground

From your dashboard, look for a link or tile labeled **Playground** — sometimes called **Try Cohere** or **Chat**. Click it. A new page opens with a large text box on the left (where you'll type instructions) and an empty response area on the right. At the top of the page there's usually a dropdown menu to pick which Cohere model you want to use. Think of the Playground as a sandbox — a safe testing space where you can experiment without affecting any real system.

If you don't see a Playground link in the sidebar, look for a top navigation bar with options like **Dashboard**, **Playground**, **API Keys**, and **Docs** instead. If you land on a page full of code, you've opened the API reference — click the **Playground** tab to leave it.

You'll know it worked when you see a big empty text box and a button labeled something like **Generate** or **Send**.

Step 3 — Pick the Command model

At the top of the Playground, find the model dropdown — a small clickable box that lists available AI models (an AI model is the "brain" that reads your text and writes a response). Choose **Command** or **Command-R** if those options are visible, since those are best at following instructions for structured tasks. If newer model names appear, any of the "Command" family should work fine. Leave the other settings (temperature, max tokens) at their defaults for now — **temperature** controls how creative or predictable the output is, and **max tokens** controls how long the response can be. Defaults are safe.

If the dropdown only shows older model names like "generate" or "summarize", Cohere may have redesigned this part of the Playground — pick whichever current model is highlighted or marked as recommended.

You'll know it worked when the model name at the top of the page shows a "Command" option, and the right-hand response area is empty and ready.

Step 4 — Paste your messy text and ask for a clean list

Click inside the big text box on the left and paste your block of messy contact info. Then, on a new line below the text, write a short instruction telling the model what you want. The instruction is called a **prompt** — a prompt is simply the question or task you give an AI. Be specific: name

the three things you want pulled out (names, emails, phone numbers) and the format you want them in.

💬 Example: Below the pasted text, add: `Extract every person's full name, email address, and phone number from the text above. Return the result as a plain list with one contact per line, in this format: Name — Email — Phone.`

If your Playground uses separate "System" and "User" boxes, put the instruction in the System box (which sets the AI's behavior for the whole conversation) and the messy text in the User box (which carries your actual request).

You'll know it worked when you press Send and the right-hand panel fills with text that lists each contact on its own line.

Step 5 — Read the answer and double-check it

Look at the response on the right. It should show one contact per line in the format you asked for. Read every line carefully — AI models can make small mistakes, especially with phone numbers that use unusual formatting or emails that look similar to each other. If anything is missing or wrong, that's normal; you can either fix it by hand or try the prompt again. To try again, click the small **retry** or **regenerate** button (often shaped like a circular arrow ↻) next to the response.

If the answer comes back as a paragraph instead of a list, your prompt wasn't specific enough — go back to Step 4, add the words "one contact per line" and "format: Name — Email — Phone", and try again.

You'll know it worked when the response shows every contact from your original text, neatly aligned, and any missing or wrong entries are easy to spot and correct.

Step 6 — Copy the clean list into your spreadsheet

Highlight the model's response with your mouse, copy it with **Ctrl+C** (Windows) or **Cmd+C** (Mac), then open your spreadsheet and paste it into the first cell with **Ctrl+V** or **Cmd+V**. If your pasted text uses dashes (—) as separators and your spreadsheet expects tabs or commas, you can split it into columns using your spreadsheet's "Text to Columns" feature — usually found under the **Data** menu.

If your spreadsheet pastes everything into a single column instead of three, run Text to Columns with the dash as the separator. If you don't have a spreadsheet, you can paste the list into a plain text file and still have a useful, organized record.

You'll know it worked when your spreadsheet shows one row per contact, with names, emails, and phone numbers sitting in their own neat columns.

Common mistakes

- **Asking for too many things at once.** If your prompt asks for names, emails, phones, job titles, companies, and addresses in the same go, the model may drop some details. Split big tasks into smaller prompts to keep results accurate.
- **Trusting the output blindly.** AI models can misread a "1" as an "l" or skip a contact entirely. Always skim the clean list against your original text before you send anything.
- **Pasting private info you wouldn't want stored.** Free AI tools may log your inputs to improve their service. For sensitive contacts, replace names with initials or use a paid enterprise plan that promises not to train on your data.

What this means for you

- **In everyday life:** When a club, volunteer group, or family thread scatters contact info across dozens of messages, you can turn the whole mess into one tidy list in under five minutes — no typing, no copy-paste errors.
- **For your business or work:** Receptionists, sales teams, and event coordinators can clean up leads, conference RSVPs, or business-card scans in bulk, freeing hours each week for higher-value work.
- **If you're just getting started:** Try it once with a small forwarded email first. Seeing ten contacts appear as a clean list in seconds is the fastest way to trust the workflow — and to start imagining the other messy texts in your life you could clean up the same way.

Try it now

Open the Cohere Playground, paste any short email containing three or more contacts, and run the prompt from Step 4. Watch the clean list appear — then paste it into a spreadsheet to see the time you just saved.

Use Alexa+ as a Hands-Free Meeting Note-Taker and Task Builder at Work

6 steps · Read online

> **In short:** Alexa+ can listen to your meeting on an Echo device, capture what was said, and turn the action items you call out into a personal to-do list. The notes usually live inside the Alexa app, not in a shared doc, so they're a private capture layer rather than a full meeting tool.

You're already comfortable talking to Alexa+ for music, timers, or quick questions. This guide teaches one more capable workflow: running Alexa+ alongside a meeting to capture notes by voice, then pulling a clean task list out of what was said. Your hands stay on the keyboard, and you don't have to switch windows to write things down.

Before you start

- An **Amazon account** with **Alexa+** access set up and an Echo device (or the Alexa mobile app) already signed in. You should already know how to say "Alexa, ..." and get a response.
- **Permission to capture**: let people know the meeting is being recorded by voice, or stick to using Alexa+ on your own desk for personal note-taking.
- Roughly **15 minutes** for your first full run; shorter after that.
- A quiet spot for the Echo within earshot of where you usually sit in meetings.
- A free account is enough for basic note-taking, but Alexa+ features and availability vary by region and device — some capabilities may require a paid Amazon plan, so check Amazon's official page if anything you need is locked.

Step 1 — Open a new note in the Alexa app

Say **"Alexa, create a note"** or **"Alexa, start a new note called Quarterly planning."** On the Alexa app on your phone, you can also open the **Notes** section from the menu and tap a button like **Create Note**. A small confirmation sound or a notification appears, and Alexa tells you the note title back so you can verify it was created.

If your device shows a different menu, look for a **Notepad**, **Lists and Notes**, or **Capture** section instead — the label varies by app version. If you don't see any notes option at all, your region or plan may not support it; double-check Amazon's help page.

Example: say *"Alexa, create a note called Q2 planning sync."*

You'll know it worked when Alexa replies with the note name and you see a new entry in the Alexa app's Notes section.

Step 2 — Add to the note as the meeting happens

During the meeting, dictate into the open note. Say **"Alexa, add to my note"** followed by your sentence, or just say **"Alexa, write this down: ..."** if you want to be explicit. Each spoken line gets added as its own bullet in the same note, so you can keep capturing without interrupting the conversation.

If Alexa+ doesn't seem to hear you across the room, move the Echo closer, or lower background music and open windows. The microphone on most Echo devices is good for a normal office, but heavy room echo or fan noise can confuse it.

Example: *"Alexa, write this down: Maria will finalize the budget by Friday. John is drafting the comms. Sarah is booking the follow-up."*

You'll know it worked when Alexa repeats a short summary of what it captured and the new bullets appear in the note on your phone.

Step 3 — Ask Alexa+ to pull out the action items

After the meeting, say **"Alexa, what are the action items from my note?"** or **"Alexa, summarize the tasks from my Q2 planning sync note."** Alexa+ will read back the sentences that look like tasks — anything starting with a name and a verb, deadlines, or commitments. The reply is spoken, and you can also see it in the app.

If the summary misses something, try asking more directly: **"Alexa, list everything in that note that has a person's name."** Different wording often pulls out different pieces. If Alexa+ asks a clarifying question instead of answering, give it the note name again to keep things on track.

Example: *"Alexa, pull the action items from my Q2 planning sync note and read them back to me."*

You'll know it worked when Alexa reads a short list of name-plus-task pairs and that list matches what you remember from the meeting.

Step 4 — Send those tasks to your personal to-do list

Say **"Alexa, add those to my to-do list,"** or specify the list name: **"Alexa, add the action items to my Work to-do list."** Each action item becomes its own to-do entry, which you can then check off from the Alexa app or by voice later. This is the moment the captured meeting actually turns into tracked work.

If you only want some of the items, ask **"Alexa, add the first two action items to my to-do list"** instead of dumping all of them. If your to-do lists don't show up, ask **"Alexa, show my lists"** to see what's available, then rename or create a new one.

Example: *"Alexa, add all of those action items to my Work to-do list."*

You'll know it worked when Alexa confirms each task was added, and you can see the new entries in the to-do list inside the Alexa app.

Step 5 — Ask Alexa+ to draft a follow-up email from the note

Say **"Alexa, write an email to my team recapping the meeting"** or **"Alexa, draft a short summary email from my Q2 planning sync note."** Alexa+ will compose a short message that names the key decisions and tasks, which you can review before sending. This step is optional, but it saves the manual rewrite most people do right after a meeting.

If you want a different tone, say so: **"Alexa, make it more formal"** or **"Alexa, make it three sentences."** You can also add a recipient by saying **"Alexa, send it to John"** if John is saved as a contact and your setup supports email.

Example: *"Alexa, draft a follow-up email from my Q2 planning sync note. Keep it under five sentences."*

You'll know it worked when Alexa reads back a complete draft email with the key points and tasks you dictated.

Step 6 — Review and clean up at the end of the day

Before you close your laptop, open the Alexa app and skim the note. Delete anything that was said off the record, correct any names Alexa misheard, and check the to-do list for the day. This is the small habit that makes the whole workflow trustworthy — without it, you end up with a pile of half-clean notes.

If you want to keep the note, say **"Alexa, save my note"** or simply leave it open. If you'd rather discard it, say **"Alexa, delete the note called Q2 planning sync."** Both work; pick one and stay consistent.

Example: *"Alexa, show my note called Q2 planning sync so I can review it."*

You'll know it worked when the note opens in the Alexa app with all the bullet points you spoke during the meeting, ready for a quick proofread.

Common mistakes

- **Forgetting to tell people the meeting is being captured.** Voice capture is still recording. A short "I'm using Alexa+ to take notes for myself" at the top of the meeting avoids uncomfortable surprises later.
- **Letting one Echo hear two people at once.** Alexa+ tries to capture everything in range. If your coworker is on a call next to you, expect their words in your note. Use a closer Echo or a headset-mic setup if you need cleaner separation.
- **Trusting the summary without checking.** Alexa+ usually picks up the right action items, but it can miss a deadline or mishear a name. A 60-second skim of the note at the end of the day catches almost every mistake.

What this means for you

- **In everyday life:** for personal calls or family planning — say, a phone call with a relative about a doctor's appointment — Alexa+ can capture the to-dos without you scrambling for a pen. The same note-and-list workflow works outside work.
- **For your business or work:** in a one-on-one or a small recurring meeting, Alexa+ becomes a private capture layer: you stay engaged in the conversation, and the action items are already waiting in your to-do list when the meeting ends. It is not a replacement for shared meeting tools — it complements them.
- **If you want to go further:** once the to-do list is captured, ask Alexa+ to schedule the deadlines you mentioned, or to set reminders that fire an hour before each due date. That turns a passive list into a workflow that nudges you, not just a record you have to remember to check.

Try it now

Pick your next meeting — even a 15-minute one — and run through Steps 1 to 4 only. Create the note at the start, dictate a couple of lines, then ask Alexa+ to pull out the action items and add them to a to-do list. You'll see in under ten minutes whether the workflow fits how you actually work.

How to build or improve your CV with AI

6 steps · Read online

You'll walk away able to turn a messy list of jobs into a clear, professional CV and tweak it for each role in under 20 minutes. No design skills required.

Before you start

- **A free AI chat tool account:** Like ChatGPT, Gemini, or Claude. You'll need to have an account set up and be logged in. We'll use **ChatGPT** for our examples, but the steps are similar for other tools.
- **Access to the AI tool:** Via a web browser on a computer or tablet is recommended, though a smartphone app can also work.
- **Your work history:** A simple list of your job titles, dates, and key tasks (you'll create this in Step 1).
- **One job description:** The description of a job you'd like to apply for, which you'll use to tailor your CV.
- **Rough total time:** Allow about 15-20 minutes for your first draft.

Step 1 — Gather your raw material

Open a digital note-taking app, a blank document in a word processor, or even a simple text editor on your device. This is where you'll jot down your work history without worrying about formatting.

You'll see a blank page or a new document ready for you to type. Once you start typing, your text will appear on the screen. If you're more comfortable writing by hand, that's fine too – just make sure it's clear enough to type into the AI later.

Example:

```
2022–2024  Café Manager, Brew House
- Led a team of 8 during peak hours
- Created weekly rosters and handled payroll
- Reduced waste by 15 % through better ordering
2020-2022 Barista, The Daily Grind
- Prepared coffee and served customers
- Managed cash register and end-of-day reports
- Maintained cleanliness and stock levels
```

You'll know it worked when you have a rough, unformatted list of your work history, including dates, job titles, companies, and 2-3 bullet points of what you did for each role.

Step 2 — Open your AI tool and start a new chat

Navigate to your chosen AI chat tool (like ChatGPT) in your web browser or open its app. Look for the option to start a "New Chat" or "New Conversation," which is typically a button in the top-left or bottom-left corner of the screen.

You'll see the main interface of the AI tool, often with a list of past conversations on one side and a large input box at the bottom. Clicking "New Chat" will clear any previous conversation. If you don't see a clear "New Chat" button, look for a "+" icon or a conversation bubble icon.

You'll know it worked when you see a clear, empty chat space ready for your message, usually with a prompt that says "Message ChatGPT..." or similar in the text box.

Step 3 — Paste your history and ask for a first draft

Now, copy the raw work history list you prepared in Step 1. Go back to your AI chat tool, click into the message input box, and paste your list. Then, add a clear instruction (known as a "prompt") for the AI to follow, asking it to create a professional CV.

You'll see your work history list appear in the chat input box. After you type your instruction and press Enter or click the "Send" button (often a paper plane icon), the AI will begin typing its response below your prompt. If nothing happens, make sure you've clicked the send button after typing.

Example:

```
List:
2022–2024  Café Manager, Brew House
- Led a team of 8 during peak hours
- Created weekly rosters and handled payroll
- Reduced waste by 15 % through better ordering
```

```
Turn the list into a clean, professional CV in American English.
Use a simple reverse-chronological layout.
Keep it to one page.
Use bullet points, not paragraphs.
Do not add any made-up details.
```

You'll know it worked when the AI starts typing out a formatted CV based on your input, usually presenting it in distinct sections like 'Experience' and 'Education'.

Step 4 — Polish the language

After the AI provides its first draft, you'll want to make the bullet points more impactful. Instruct the AI to refine your descriptions by using strong action verbs and including quantifiable results where appropriate, which makes your achievements stand out more.

You'll see the AI's previous CV draft above, and the current chat input box below it. The AI will then generate a revised version of the CV directly within the same chat window. If the AI seems confused, simply say "Referring to the CV you just wrote, please..." to guide it.

Example:

```
Rewrite every bullet point to start with a strong action verb and include a quantifiable result where possible.
```

You'll know it worked when the AI presents a new version of your CV, with bullet points starting with words like 'Managed,' 'Developed,' or 'Achieved,' and where numbers (like 'reduced waste by 15%') are clearly highlighted.

Step 5 — Tailor for the job you want

This is where AI truly saves time. Copy the entire job description for the specific role you're applying for. Then, paste it into the AI chat and instruct the AI to adjust your CV to match the keywords and skills mentioned in that job description.

You'll see the text of the job description in your device's clipboard. After you paste it into the chat and send your prompt, the AI will use this new information to adapt your CV. If the job description is an image, you'll need to manually type out the key requirements.

Example:

```
Please rewrite my CV bullet points to specifically match the keywords and skills mentioned in this job description:

[Paste the full job description here, e.g., "Seeking a café manager with proven leadership skills, experience in staff rostering, and a trac
```

You'll know it worked when the AI presents an updated CV where the skills and experiences mentioned closely mirror the language used in the job description, making your application feel more relevant.

Step 6 — Format and export

Finally, ask the AI to give you the finished CV in a plain text format. This makes it easy to copy and paste into any word processor like Microsoft Word or Google Docs, where you can then apply your preferred visual formatting (fonts, spacing, etc.) before saving or printing.

You'll see the complete, tailored CV as the AI's latest response. The AI will then reformat its output, stripping away any complex styling to present the text in a simple, unformatted block. If the text still has too much formatting, just ask it again more directly, for instance, "Just give me the raw text, no special characters or extra lines."

Example:

```
Give me the final CV in plain text so I can easily copy it into Word or Google Docs for final formatting.
```

You'll know it worked when the AI displays your CV in a simple, unformatted block of text, easy to select and copy directly into your chosen document editor.

Common mistakes

- **Not providing enough detail:** If your initial work history is too sparse, the AI can't create strong bullet points.
  - **Fix:** Go back to Step 1 and add more concrete examples, responsibilities, and any numbers you can remember (e.g., "managed a team of 5," "processed 20 invoices daily").
- **Relying on the AI for facts:** AI tools can sometimes "hallucinate," meaning they confidently make up information that isn't true.
  - **Fix:** Always double-check every date, job title, company name, and achievement the AI includes. The AI can draft and phrase, but it cannot verify your personal history.

- **Skipping the "plain text" export:** If you just copy the formatted text directly from the chat, you might bring over hidden formatting codes that make it hard to edit in Word or Google Docs.
    - **Fix:** Always follow Step 6 and request plain text, then apply your desired formatting (like bolding and bullet styles) *after* you've pasted it into your word processor.

What this means for you

- **In everyday life:** Update your CV once with the AI, then reuse it for every application by quickly tailoring it to each job description, saving you hours.
- **For your business or work:** Keep a master CV or professional bio in plain text and ask the AI to quickly adapt it for client proposals, tender applications, or conference speaker profiles.
- **If you're just getting started:** Start with your most recent job and build one page at a time. The AI can help you find the right words even if you're not sure what to write.

Try it now

Open your chosen AI chat, paste the example prompt from Step 3, and hit send. You'll have a first draft of your CV in under two minutes, ready for refinement!

Understand anything

How to Ask GitHub Copilot to Explain a Confusing Error Message

7 steps · Read online

> **In short:** GitHub Copilot's chat can turn a confusing error message into plain English and a short list of things to check — you paste the error into the chat, then ask a simple question. It works inside editors like VS Code and on github.com; some chat features may require a paid plan, so check the official pricing page for the latest details.

You don't need to be a developer to get value from this. Plenty of people run into error messages while learning to code, tweaking a website, or running a script at work. GitHub Copilot is an AI coding assistant made by GitHub — think of it as a patient tutor sitting next to you, reading the error out loud in normal English. Here's the beginner-friendly way to use it.

Before you start

- A free GitHub account, signed in. Some chat features may require a paid plan — if a feature asks you to upgrade, that's why.
- The error message itself, copied from wherever you saw it (your terminal, your browser console, a log file, or your code editor).
- A code editor with the Copilot extension installed (VS Code is the most common), **or** the Copilot chat panel on github.com.
- Roughly 5 minutes.

Step 1 — Install the Copilot extension in your editor (if you haven't yet)

Open VS Code and look at the left sidebar — the narrow vertical strip with several icons. Click the square **Extensions** icon (it usually looks like four small squares, with one square separated from the others). A search panel opens on the left. Type **GitHub Copilot** into the search box at the top of that panel. The top result is published by GitHub itself — click **Install** on that row. After a few seconds, a new icon appears in your sidebar (often a chat-bubble shape). If you use a different editor like JetBrains or Visual Studio, the steps are similar — open your editor's Extensions or Plugins menu, search for "GitHub Copilot," and click Install. On github.com, you can skip this step entirely and go straight to Step 2. Example: (no typing needed in this step.) You'll know it worked when you see the Copilot icon in your sidebar, or when github.com shows a "Copilot Chat" option in its top menu.

Step 2 — Open the Copilot chat panel

Look for the Copilot icon you just installed — it's usually in the left or right sidebar, near the top. Click it once. A new side panel opens up, with a text box at the bottom and a label like "Ask Copilot" or "Send a message" above it. If you don't see the icon, try the keyboard shortcut: on Windows or Linux, press **Ctrl + Shift + I**; on macOS, press **Cmd + Shift + I**. The chat panel pops up the same way. Some setups hide the icon inside a **⋯ (three dots)** menu instead — open that menu and pick "Copilot Chat." Example: (no typing needed in this step.) You'll know it worked when you see a blinking cursor inside a text box labeled something like "Ask Copilot" or "Send a message".

Step 3 — Paste the full error message

Click inside that text box and paste your error. To paste, use **Ctrl + V** on Windows/Linux or **Cmd + V** on macOS. The error text appears in the box, usually in red or in a monospaced font (where every letter takes the same width, like a typewriter). Try to copy the whole thing, not just the first line — most errors include a **stack trace** (a list of which files and lines were involved, like a breadcrumb trail showing where the problem happened) right under the main message, and that context is what helps Copilot pinpoint the problem. If the error is in your browser, open the **Console** tab first (usually under the browser's "More tools" or "Developer" menu), then copy from there. Example: `TypeError: Cannot read properties of undefined (reading 'map') at App.render (App.jsx:14)` You'll know it worked when the error text sits in the chat box, exactly as it appeared on screen.

Step 4 — Ask for a plain-language explanation

Now type one short sentence *after* the error, asking Copilot to translate it. This part matters: the question you ask shapes the answer you get. Keep it simple and direct — no jargon required. Press **Enter** to send. Within a few seconds, Copilot replies with a friendly explanation of what the error means, in everyday words. If you don't see a response right away, look for a small **stop** or **regenerate** icon below your message and click it. If nothing comes back at all, your plan may not include chat — the reply box will usually show a small note asking you to upgrade. Example: `Explain this error to me like I'm new to coding. What is it actually telling me to fix?` You'll know it worked when the reply describes the problem in normal English — something like "your code tried to use a list that doesn't exist yet" — instead of repeating the original error back at you.

Step 5 — Ask for a checklist of likely fixes

Don't stop at the explanation. Right below Copilot's reply, type a follow-up asking for a short list of the most likely causes and what to check first. Copilot usually responds with a numbered list — each item is a hypothesis to test against your own code. Treat the list as a **checklist**, not a verdict: not every suggestion will fit your situation, and that's normal. If the answer feels too generic, add one line about what you're building (for example: *"I'm using React"* or *"I'm running Python on Windows"*) and ask again. Example: `List the 3-5 most likely causes of this error.`

`For each one, tell me what to look at in my code.` You'll know it worked when you get a numbered list with at least three concrete things to check, like "the variable might be empty" or "the function might be called before the data loads".

Step 6 — Test the suggestions against your own code

Open the file mentioned in the error (in the example, `App.jsx` on line 14). Read the line Copilot pointed to and see which suggestion fits. If none of them seem to match, paste a small chunk of your actual code into the chat and ask Copilot which suggestion is most likely. This is where the AI becomes useful — instead of guessing, you're now comparing each idea to the real code in front of you. One honest note: Copilot can be confidently wrong, especially for unusual setups — its suggestions are a starting point, not a final answer. 💬 Example: `Here's the function that produces the error. Which of your suggestions fits this code? function render() { return items.map(...) }` You'll know it worked when at least one suggestion clicks — you think "oh, that's probably it" — and you have a clear next step to try.

Step 7 — Ask a follow-up when something still doesn't make sense

If a word in the explanation or the list is unfamiliar, just ask. Tell Copilot exactly which term confused you, and it'll rephrase that one part. This is the secret most beginners miss: the chat is a conversation, not a one-shot answer. Each follow-up makes the next answer tighter and more useful for your specific situation. If the answer goes off-track, click the small **refresh** or **try again** icon to get a new draft. 💬 Example: `You mentioned "undefined". I don't fully get what that means in this context — can you explain just that word?` You'll know it worked when you can say the idea back in your own words, and you have at least one concrete thing to change in your code.

Common mistakes

- **Pasting only the first line of the error.** The lines below it (the stack trace) usually point to the file and line where things went wrong. Copy the whole block.
- **Forgetting to say what language or library you're using.** "Null" in JavaScript and "NULL" in a database mean similar but different things. Add one short line: *"I'm using React."*
- **Trusting the very first answer as the final answer.** Copilot can be confidently wrong, especially for unusual setups. If two or three ideas don't match your code, don't force the fourth — go back to the error message itself and ask again.

What this means for you

- **In everyday life:** when any tool gives you a wall of red text — a script, a build tool, a deployment — the same trick works. Paste, ask for a plain-language read, then ask for a checklist.
- **For your business or work:** faster debugging, fewer interruptions to teammates, and a growing habit of reading errors instead of fearing them. Over time you'll spot patterns and need the chat less.
- **If you're just getting started:** pick the next error you see today and use it as practice. One error a day, explained, adds up quickly.

Try it now

Open Copilot chat, paste the last error message you ran into — even if it's from yesterday — and send this single line:

💬 Example: `Explain this error to me like I'm new to coding, then list the 3 most likely things to check.`

You'll have a clear answer in under a minute.

Compare two job offers side-by-side with Le Chat in plain English

7 steps · Read online

By the end of this guide you'll be able to paste the key points of two job offers into Le Chat and receive a tidy, easy-to-understand table that highlights the differences—no legal jargon required. This is for anyone juggling two offers and wants confidence before signing.

Before you start

- A free (or paid) Mistral account so you can log into Le Chat.
- Access to the two job offers – either as PDF files you can copy from, or as plain-text emails.
- A computer or smartphone with an internet browser.
- About 10 minutes total – the guide walks you through every click and prompt.

Step 1 — Open Le Chat and start a fresh chat

Open your web browser and go to the Le Chat website (chat.mistral.ai). Sign in with your Mistral email and password. You'll see a clean white chat window with a text box at the bottom and a "New chat" button near the top-right corner. Click the "New chat" button; a fresh conversation pane slides in and the cursor blinks in the text box. If the button looks different on your screen (for example, it might be a plus-sign icon or simply the words "Start new chat"), look for the same meaning. Example: *"Hello Le Chat, I need help comparing two job offers."* You'll know it worked when the chat area shows a greeting like "How can I help you today?" or "New conversation started."

Step 2 — Copy the key details from each offer into plain text

Open the first job offer in your email or PDF viewer. Select only the most important sections—typically base salary, bonuses, work hours, leave entitlements, notice period, and any special perks. Copy the selection (Ctrl+C on a computer, long-press → Copy on a phone). Repeat for the second offer, keeping the two chunks separate on your clipboard. Back in Le Chat you'll see the empty text box ready for your first prompt. If you cannot copy from a PDF, use a free PDF-to-text converter or type the details manually. Example: *"Offer A: $95,000 base salary; 15 days annual leave; 1 month notice; $5,000 signing bonus. Offer B: $92,000 base salary; 20 days annual leave; 2 weeks notice; 3% profit share."* You'll know it worked when the copied text appears in the box after you paste (Ctrl+V).

Step 3 — Ask Le Chat to build a side-by-side comparison table

With the two offer snippets pasted into the text box, type a clear instruction: *"Create a side-by-side comparison table in plain English for Offer A and Offer B. Include salary, annual leave days, notice period, bonuses, and any other notable perks. Avoid legal jargon."* Press Enter or tap the send arrow. Le Chat will start typing a response that usually begins with a markdown-style table and then bullet points summarizing the contrast. If the assistant replies with a single paragraph instead of a table, simply add "Show it as a table" to the prompt. Example: *"Compare these two offers side-by-side and explain the pros and cons of each in simple language."* You'll know it worked when you see a two-column table with headings like "Offer A" and "Offer B" and simple sentences underneath each row.

Step 4 — Review the table and ask for missing rows

Look at the table Le Chat has generated. Each row should correspond to a contract element (salary, leave, notice, bonuses, etc.) and each column should contain the figure or phrase from the respective offer. If a row is missing (for example, "profit share"), type a follow-up such as *"Add a row for profit share details for both offers."* The assistant will update the table accordingly. On some devices the table may appear as plain text rather than a rendered grid; if you see vertical bars ("|") and dashes ("—") instead of a visual grid, that's still the correct markdown format. Example: *"Can you explain in everyday terms what the different notice periods mean for my life?"* You'll know it worked when the assistant adds the requested row or provides a short, jargon-free explanation beneath the table.

Step 5 — Copy the final comparison into your notes

When the table looks complete, click the three-dot menu (or right-click) on the response and choose "Copy message" or "Download as text" depending on your browser. Paste the copied result into a notes app, email to yourself, or print it out. If the copy option isn't visible, you can manually select the whole response and press Ctrl+C (or long-press → Copy on a phone). Example: *No extra prompt needed – just use the copy command.* You'll know it worked when the entire table appears in your notes or email, ready for you to review side-by-side with a coffee.

Step 6 — Double-check the numbers against the original offers

Open the original offer documents again and compare the figures in your notes with the source. If anything looks different, paste the discrepancy into Le Chat and ask for clarification: *"The original Offer A PDF shows $95,000 base salary, but your table shows $92,000. Which is correct?"* Le Chat will usually point to the exact line in the text you pasted earlier. If the mismatch remains, paste the relevant paragraph again and ask for a corrected table. You'll know it worked when the numbers in your notes match the original documents.

Step 7 — Decide with confidence

With the plain-English table in hand and the numbers verified, you now have a clear snapshot of the two offers. Take five minutes to weigh salary, leave, notice, bonuses, and any other perks against your personal priorities. You'll know you're ready to decide when you can confidently say which offer better supports your lifestyle and goals.

Common mistakes

- **Copying the entire PDF instead of the key points** – the assistant may get overwhelmed and miss important details. Fix: only copy salary, bonuses, leave, notice, and perks.
- **Using legal-sounding wording in the prompt** – Le Chat might mirror that jargon. Fix: ask explicitly for “plain English” or “simple language” in your first instruction.
- **Forgetting to ask for a table** – you may receive a paragraph of text. Fix: add “show it as a table” or “format as a two-column chart” in a follow-up prompt.
- **Not verifying the numbers against the originals** – a small typo can change the outcome. Fix: always compare the final table with the source documents before deciding.

What this means for you

- **In everyday life:** You can instantly see which offer gives you more vacation days or a higher take-home pay, without needing a lawyer or calculator.
- **For your business or work:** HR professionals can use the same approach to benchmark internal promotions against external market offers in minutes.
- **If you're just getting started:** The easiest first move is to copy the salary and leave details from one offer and ask Le Chat for a quick plain-English summary.

Try it now

Open Le Chat, paste the salary and leave figures from one of your current job offers, and type: *“Summarize these numbers in simple language.”* In under two minutes you’ll see how Le Chat turns raw numbers into an easy-to-understand sentence.

Quickly Research Any Topic and Summarize Web Pages with Microsoft Copilot in Edge

7 steps · Read online

Imagine reading a long article or research page and getting a clear, concise summary in seconds—without leaving your browser. With **Microsoft Copilot in Edge**, you can do exactly that. Whether you're studying, working, or just curious, this tool helps you understand web pages quickly and answer your questions on the spot. And it's free.

Before you start

- A **computer, tablet, or phone** with internet access.
- The **Microsoft Edge browser** installed (free from microsoft.com/edge).
- A **Microsoft account** (free to create at account.microsoft.com).
- About **5 minutes** of your time.

> 💡 Tip: You don't need to sign in to use basic features, but signing in lets you save your chats and access them across devices.

Step 1 — Open Copilot in Edge

Open the **Microsoft Edge** browser on your device. Look at the top-right corner of the screen. You'll see a small **blue Copilot icon** (it looks like a sparkle or a chat bubble with a sparkle inside).

Tap or click the **Copilot icon**. A sidebar will slide open from the right side of your screen. This is where Copilot lives—it's like a helpful assistant inside your browser.

You'll know it worked when ... the Copilot sidebar appears with a message like "Ask me anything" or "What can I help you with?"

💬 Example: Just click the sparkle icon in the top-right corner of Edge.

Step 2 — Choose "Summarize" from the menu

Inside the Copilot sidebar, look for a dropdown menu at the top. It might say "Chat" or "Compose" by default. Click the dropdown and select **"Summarize"**.

Now, Copilot is ready to summarize the web page you're viewing. It will read the content of the page and create a short, clear summary for you.

You'll know it worked when ... the sidebar shows a message like "Summarizing this page..." and then displays a bullet-point summary below.

💬 Example: In the Copilot sidebar, click the dropdown and choose "Summarize".

> ⚠️ If you don't see "Summarize" in the dropdown, make sure you're on a web page with text (not a video or image-only page). Also, look for a button like "📄 Summarize" or "📝 Summarize page".

Step 3 — Read the AI summary

After a few seconds, Copilot will show you a **short summary** of the web page. It might include:

- The main points or arguments
- Key facts or data
- A conclusion or takeaway

You can scroll through the summary to get the gist of the page without reading every word.

You'll know it worked when ... you see a clear, bulleted list or short paragraphs that capture the most important information from the page.

💬 Example: You'll see something like:

- "This page explains how to compost at home."
- "It lists 5 easy steps and common mistakes to avoid."
- "It also includes a list of local composting services."

Step 4 — Ask follow-up questions

After the summary appears, you can ask Copilot to **clarify or expand** on any point. Just type your question in the chat box at the bottom of the sidebar.

For example:

- “Can you explain point 3 in more detail?”
- “What are the benefits of composting?”
- “Give me a shorter version.”

Copilot will respond based on the content of the page. This is like having a conversation with the article itself.

You’ll know it worked when … Copilot answers your question using information from the page, and you see a new message appear in the chat.

💬 Example: Type: “What are the top 3 benefits of composting mentioned here?”

> ⚠️ If Copilot gives an answer that doesn’t match the page, try rephrasing your question or check if the page content changed. Copilot only knows what’s on the current page.

Step 5 — Summarize a new page without leaving your tab

To summarize a **different web page**, just navigate to it in your browser. Copilot will automatically detect the new page and ask if you want to summarize it.

You can say “Yes” or click the “Summarize” button that appears in the sidebar. Copilot will generate a new summary for the new page.

You’ll know it worked when … the sidebar updates with a new summary of the page you’re now viewing.

💬 Example: Open a new article or research page, and Copilot will prompt: “Want me to summarize this page?”

Step 6 — Save or share the summary

If you want to save the summary for later, you can:

- **Copy the text** and paste it into a document or note app.
- **Take a screenshot** of the summary.
- **Ask Copilot to save it** by typing: “Save this summary to my notes.”

You’ll know it worked when … you have the summary text in your notes, or you see a confirmation like “Summary saved.”

💬 Example: Type: “Save this summary to my notes” to store it in your Copilot chat history.

> ⚠️ On the free plan, your chat history is saved for a limited time. If you want long-term access, sign in with your Microsoft account.

Step 7 — Use Copilot on mobile (optional)

If you’re using a phone or tablet, open the **Edge app**. Tap the **Copilot icon** (sparkle) in the bottom toolbar. Then tap the **“Summarize”** button at the top of the Copilot panel.

Copilot will summarize the page you’re viewing. You can scroll, ask questions, or copy the summary—just like on desktop.

You’ll know it worked when … the summary appears in the Copilot panel on your mobile device.

💬 Example: In the Edge mobile app, tap the Copilot icon, then choose “Summarize”.

> ⚠️ On mobile, the interface may look slightly different. Look for a button like “ Summarize” or “Summarize page”.

Common mistakes

- **Mistake:** Trying to summarize a video or image-only page. **Fix:** Copilot only works on pages with text. Make sure the page has readable content.
- **Mistake:** Expecting Copilot to understand pages you haven’t opened. **Fix:** Copilot only sees the page you’re currently viewing. Open the page first, then ask for a summary.
- **Mistake:** Not signing in and losing your chat history. **Fix:** Sign in with your Microsoft account to save your summaries and access them later across devices.

What this means for you

- **In everyday life:** Quickly understand long articles, research papers, or product pages without reading every word. Great for students, parents, or anyone short on time.
- **For your business or work:** Save hours reviewing reports, competitor websites, or industry updates. Ask Copilot to pull out key data or trends from any page.
- **If you're just getting started:** Start by summarizing one article you're already reading. Try asking a follow-up question to see how it works.

Try it now

Open **Microsoft Edge**, go to any article or research page, and click the **Copilot icon** in the top-right corner. Choose "Summarize" from the dropdown, and see your first AI summary appear in seconds.

Fact-Check Online Claims with Perplexity's Sources

6 steps · Read online

In our busy online world, information flies around faster than ever. Sometimes, a post on social media might sound convincing but isn't quite right. This guide will show you how to use Perplexity, an AI assistant, to quickly check the facts and understand where information truly comes from. By the end, you'll be able to confidently verify claims, helping you decide what's reliable to share. While Perplexity is excellent at finding and compiling information, remember it doesn't perform human critical analysis or judgment; it's always up to you to evaluate the original sources.

Before you start

- A device with internet access (computer, tablet, or smartphone).
- Access to Perplexity, either through its website (perplexity.ai) or its mobile app. A free account is sufficient for this guide.
- Roughly 5-7 minutes of your time.

Step 1 — Open Perplexity

First, let's get Perplexity ready. Open your web browser and type `perplexity.ai` into the address bar at the top of the screen, then press your keyboard's Enter key. If you're on a mobile device, find and launch the Perplexity app. Once loaded, you'll typically see a clean, simple screen featuring a prominent search bar right in the middle, often displaying placeholder text like "Ask anything..." or "Search the web". This is your main input area, waiting for your questions.

If it looks different: The specific layout might vary slightly depending on whether you're using the website or the app, or due to recent updates. However, you should always find a clear, large text box or search bar positioned centrally, as it's the most obvious place to type your query.

You'll know it worked when you see the Perplexity interface with a clear text input area, ready for you to type your question.

Step 2 — Ask about your claim

Now, let's ask Perplexity to investigate a social media claim. Carefully type the specific statement or question you want to check directly into that main search bar. Try to phrase your query clearly and directly, just as you would ask a person for information. For example, if you saw a post about a new animal, you might ask "Is the 'blue-spotted tree frog' a real animal?". Once you've finished typing your question, press the Enter key on your keyboard or tap the "Ask" button (which might appear as an arrow or a magnifying glass icon) usually located next to your input.

If it looks different: If there isn't a clearly visible "Ask" button, simply pressing Enter after typing your question usually works on most devices. Make sure your question is specific enough so Perplexity understands exactly what you're trying to verify.

Example: `Is it true that taking a cold shower every morning can cure the common cold?`

You'll know it worked when Perplexity starts to process your request, often showing a loading animation, a spinning icon, or a "Generating answer..." message as it looks for information.

Step 3 — Review the summary and its citations

After a brief moment, Perplexity, acting as an **AI assistant** (a computer program designed to help with tasks and provide information), will present you with a **summary answer**. This is the AI's best attempt to directly address your question, drawing information from various online sources. Read through this summary carefully. You'll typically find a concise answer, along with key facts or explanations. Notice that within the summary, there will be small **numbered references**, like [1], [2], or [3], next to specific pieces of information. These numbers are crucial because they point directly to the original sources Perplexity used.

If it looks different: The summary might be formatted as a series of bullet points, a few longer paragraphs, or a mix of both. Regardless of the exact appearance, the core information and the numbered references should be clearly visible, usually appearing immediately below your initial question.

You'll know it worked when you see a generated answer to your query, complete with small, numbered citations embedded within the text.

Step 4 — Examine the source list

This is where Perplexity truly shines for transparency and fact-checking. Scroll down slightly below the main summary answer. You'll find a dedicated section, typically labeled "Sources" or "References." Here, you'll see a list of the actual websites or articles that Perplexity consulted to formulate its answer, corresponding to the numbers you saw in the summary. For each source, there will usually be a title, a brief description of the content, and the website's address. This **transparency** (showing you exactly where the information came from) is vital because it lets you see *where* the information originated, not just *what* the AI says.

If it looks different: The sources might be displayed as small, clickable cards, or as a simple list. On a mobile device, you might need to tap an "Expand Sources" or a similar button to reveal the full list. Just look for a distinct section that clearly lists external website links.

You'll know it worked when you see a list of external website links, each labeled with a number that corresponds to the citations in the AI's summary.

Step 5 — Dive into an original source

To truly verify a claim and understand its context, it's essential to go directly to the original source. Click on one of the numbered source links from the "Sources" section that seems most relevant to the claim you're checking or appears to be from a **reputable** (trustworthy and respected) organization. This action will open that original webpage in a new browser tab or sometimes within Perplexity itself (depending on your device settings). Now, you can read the information in its original context. Check if the information presented by Perplexity accurately reflects what the source says, look for publication dates, the author, and whether the website itself appears trustworthy (e.g., is it a known news organization, a research institution, or a personal blog?). Comparing the AI's summary with the actual source helps you confirm its accuracy.

If it looks different: If clicking a link doesn't open a new tab, check your browser's pop-up blocker settings, as they can sometimes interfere. You might also have the option to right-click (on a computer) or long-press (on a phone) the link to copy its address, then paste it directly into a new browser tab yourself.

You'll know it worked when you are actively viewing the actual website or article that Perplexity used, allowing you to read the original information.

Step 6 — Refine your query or ask follow-ups

Sometimes, your first question might not give you the complete picture, or the sources provided might be unclear. If you need more information or a different angle, you can ask a follow-up question. Below Perplexity's initial answer, you'll usually see an "Ask a follow-up" prompt or a new blank input box. Type your refined question there, focusing on clarifying details or asking for more specific types of sources. For example, if the initial sources were mostly opinion pieces, you might ask "Are there any scientific studies on this topic?". This iterative process helps you narrow down your search and get more precise, verifiable information.

If it looks different: Instead of an explicit "Ask a follow-up" button, you might simply see another blank input box directly below the answer, allowing you to continue the conversation. If you want to start a brand-new search entirely, look for a "New Thread" or "New Search" button, usually found near the top of the screen or within a side menu.

Example: `Can you find government health organisation statements about cold showers and colds?`

You'll know it worked when Perplexity processes your follow-up query, generating a new answer and set of sources based on your refinement.

Common mistakes

- **Not checking the sources thoroughly**: It's tempting to just read Perplexity's summary and assume it's entirely correct. The AI compiles information, but it can sometimes misinterpret or give too much weight to less reliable sources. Always click through to the original sources to verify the context and content yourself.
- **Being too vague with your question**: Asking broad questions like "Tell me about this news" might not yield precise fact-checks. Instead, provide the specific claim or statement you want to investigate. Being specific helps Perplexity pull more relevant information.
- **Ignoring source reliability**: While Perplexity provides sources, it doesn't always rate their trustworthiness. A quick check of a source's website (e.g., is it a well-known news organization, a research institution, or just a personal blog?) and its publication date is crucial for deciding if the information is truly reliable.

What this means for you

- **In everyday life:** You can quickly become a more informed social media user. Before sharing a sensational post or falling for misinformation, you can take a minute to verify it, protecting yourself and your friends from spreading inaccurate content.
- **For your business or work:** If you work in communications, marketing, or any field where accurate information is key, Perplexity can be a fast way to cross-check industry news, claims from competitors, or statistics before you use them in a report or presentation.
- **If you're just getting started:** The easiest first move is to pick one social media post you're curious about today and run its exact claim through Perplexity. Focus on just reading the summary and then clicking through one source link to see the original content.

Try it now

Find a recent claim or statistic on your social media feed that caught your eye, perhaps something surprising or highly shared. Type that exact claim into Perplexity's search bar and hit Enter. Once you have the summary, scroll down and click on the first source link provided to see where the information originally came from.

How to Summarize a Long Article or PDF into Key Bullet Points

6 steps · Read online

> **In short:** Paste the text (or upload the PDF) into any AI assistant, then use a short "summarize into bullet points" prompt. You'll get the main argument and the key takeaways in under five minutes — works in ChatGPT, Gemini, Claude, Copilot, Le Chat or DeepSeek.

You've just opened a 20-page report, a long news feature, or a 40-page PDF someone sent you. You don't have time to read all of it, but you do need to know what it actually says. This guide walks you through a simple workflow that turns a long document into a tight set of bullet points you can scan in a minute, using a prompt you can reuse on any AI assistant.

Before you start

- **An AI assistant** you can type into — ChatGPT, Gemini, Claude, Copilot, Le Chat or DeepSeek all work. Pick the one you already have an account with.
- **A free plan is fine** for short articles. Uploading large PDFs (say, over 50 pages) may need a paid plan or a paid tier with bigger file limits — check the assistant's current limits if your file is very long.
- **The article or PDF you want summarized.** Have it open in another tab, or saved on your device.
- **Roughly 5 minutes** from opening the assistant to having your bullet-point summary ready.

One honest note before you start: an AI summary is a starting point, not a replacement for reading. If the document is something important (a contract, a medical paper, a policy that affects you), treat the summary as a map of the territory, not the territory itself. Use it to decide what to read carefully.

Step 1 — Open your AI assistant and find the chat box

Open the website or app of your chosen AI assistant in a browser, or launch the mobile app. Look for the chat box — the long text field at the bottom of the screen, usually with a placeholder line like *"Message [assistant name]..."* or *"Ask anything..."*. Tap or click inside it so your cursor is blinking in the field. If you're on a phone and don't see a text field, look for a round + or **new chat** button on the home screen, then start a fresh conversation. A new, blank chat is ideal — older messages in the thread can confuse the summary, and the assistant only has a limited "memory" of what you paste in (this is called the **context window** — think of it as the assistant's short-term working memory, the amount of text it can actively read in one go). Example: type nothing yet — just confirm the chat is open and ready. You'll know it worked when you see an empty text field with a send button (usually a paper-plane icon or an up-arrow) next to it.

Step 2 — Get the text you want summarized ready

Open the article or PDF in another window. You have two options, and which one is easier depends on the assistant and the file type. **Option A (easiest for articles):** select the text in the article, copy it (Ctrl+C on Windows, Cmd+C on Mac, or long-press and "Copy" on a phone), then switch back to the chat. **Option B (easiest for PDFs):** look for a **paperclip icon** , an **attachment icon**, or a + button near the chat box — many assistants let you upload a PDF directly. If you don't see a paperclip, look for a menu labeled "Add files", "Upload" or similar. If neither option works, fall back to opening the PDF in a reader, selecting the text, copying it, and pasting it into the chat as plain text. Example: highlight the first few paragraphs of a long article, copy them, and have your clipboard ready. You'll know it worked when you either have text highlighted and copied, or a file ready to upload in the chat.

Step 3 — Paste the text or upload the PDF

Click into the chat box and paste the text (Ctrl+V / Cmd+V, or long-press and "Paste" on a phone). The text will appear inside the chat field as a block of words. If you're uploading a PDF, tap the paperclip or + button, choose your file from the list, and wait for the upload progress bar to finish — a small file thumbnail or filename usually appears in the chat field once it's uploaded. If the assistant says the file is "too large", try copying just the first half of the text and pasting it as plain text instead. On the free plan, file size limits are common, and they change often. Example: paste a 3,000-word article into the chat, or upload a 12-page PDF called *"Q3-strategy-review.pdf"*. You'll know it worked when you can see the pasted text block (or the uploaded file name) sitting in the chat field, ready to send.

Step 4 — Add a "summarize into bullet points" prompt

This is the most important step. The quality of your summary depends almost entirely on the instruction you give the assistant. A vague request like "summarize this" usually gives you a waffly paragraph. A specific request gives you clean, scannable bullets. Type the following prompt directly under the pasted text or uploaded file, then send the message (press the paper-plane or up-arrow button). The assistant will start "reading" the content, and within a few seconds you'll see bullet points appear. Example: *"Please summarize this into bullet points. I want: 1) the main argument or thesis in one sentence, 2) the 5–7 key takeaways as short bullets, 3) any important numbers, dates or names mentioned, 4) one sentence on what the author is asking the reader to do or believe. Keep it under 250 words."* You'll know it worked when the assistant replies with a structured list of bullets (not a long paragraph), and you can identify the main point within the first three bullets.

Step 5 — Check the summary and ask for a tweak

Read the bullet points. Ask yourself: do I now know what this document is actually about? If the summary is too long, too vague, or missing the main point, send a short follow-up. The assistant remembers the original document in the same chat, so you can refine without re-pasting anything. If the summary missed something, name what you wanted. If it's too long, ask for it shorter. If you want a different angle — say, the parts that affect your job — ask for that view. 💬 Example: *"That's good. Now give me just the 3 bullets most relevant to a small business owner."* or *"Cut this down to 5 bullets maximum."* You'll know it worked when the second reply gives you a summary that you can read in under a minute and confidently explain to someone else.

Step 6 — Save or share the result

Once you're happy with the summary, copy it out of the chat (select the assistant's reply, copy, and paste into a note, email or document). If your assistant has a **share** or **copy** button near the reply, use that — it's usually faster. Saving the bullets somewhere means you don't have to come back to the chat later, and you can paste them into a meeting agenda, a study note, or a message to a colleague. Over time, building a small folder of "bullet summaries" is a quiet superpower for keeping up with reading you don't have time to do properly. 💬 Example: paste the bullets into a note called *"Reading — June 2026"*, with a one-line note above them saying where they came from. You'll know it worked when the summary is saved in a place you can find again, and you can skim it in under a minute the next time you need it.

Common mistakes

- **Asking a vague question.** "Summarize this" usually gets you a wobbly paragraph. Always specify the format: bullet points, a number of bullets, and what you want covered (main argument, key takeaways, numbers). The more specific the instruction, the better the output.
- **Pasting a 50-page PDF as plain text.** If the document is very long and you paste it all in, the assistant may quietly drop the middle of the text because of the context window limit. For long files, use the PDF upload option if your assistant offers one, or paste just the introduction and the conclusion first to get a quick orientation.
- **Trusting the summary without a sanity check.** AI assistants sometimes invent a number, a name, or a quote that isn't in the original. If a bullet contains a specific fact that matters to you, scroll back to the original document and confirm it's there. Treat surprising claims as a prompt to verify, not as confirmed truth.

What this means for you

- **In everyday life:** Got a long newsletter, a council report, or a terms-of-service update you don't want to read in full? Run it through this workflow in five minutes and you'll know whether it's worth your time.
- **For your business or work:** Use it on meeting prep notes, industry reports, competitor blog posts, or long client briefs. A bullet summary you can re-read in 60 seconds is far more useful than a document you read once and never open again.
- **If you're just getting started:** Pick one document you've been meaning to read — an article someone sent you, a report you bookmarked — and run it through Steps 1–6 today. The first time feels slow; the second time is genuinely five minutes.

Try it now

Open the assistant you use most, find a long article you've been putting off, and run it through this exact prompt: *"Please summarize this into bullet points. I want: 1) the main argument in one sentence, 2) the 5–7 key takeaways as short bullets, 3) any important numbers or dates, 4) one sentence on what the author wants the reader to do or believe. Keep it under 250 words."* Five minutes, one document, one useful habit started.

Translate a product description into Spanish and French with Cohere

6 steps · Read online

> **In short:** Use Cohere's free Playground to turn one English product description into Spanish and French in a single prompt. You'll need a Cohere account (free to start), and the result is a strong first draft — you should still read it through before publishing.

By the end of this guide, you'll have two ready-to-use product descriptions — one in Spanish, one in French — sitting in your store or in a backup doc. This walkthrough is for small business owners and marketers who want a fast, good-enough first draft without hiring a translator for every product.

Before you start

- **A free Cohere account.** Sign-up takes about 2 minutes at the official Cohere site.
- **Your English product description.** Copy it from your store or product sheet and have it ready to paste.
- **Roughly 10 minutes.** Most of that is reading and double-checking the AI's answer.
- **A free plan is enough** for this task. Cohere may limit how many messages you can send per minute, but a single product description easily fits inside the free tier.

A quick note on what Cohere **cannot** do for this task: it won't catch every specialist word (a fabric name, a regulatory phrase, a brand-specific term) and it isn't a substitute for a sworn legal translation. Treat its output as a polished first draft, not a final one.

Step 1 — Create your free Cohere account

Open your browser and go to Cohere's homepage. Click the button in the top-right corner — it's usually labeled "Sign up", "Get started", or "Try for free". You'll land on a sign-up page. You can register with your email and a password, or click the Google or GitHub button to skip the form. After signing up, Cohere may send a confirmation email; open it and click the verification link inside before you continue.

Example: use your real work email so future receipts, API keys, and account-recovery messages all reach you.

If the button wording looks different ("Get API key" or "Start building"), click that instead — it leads to the same dashboard. Cohere updates its labels often.

You'll know it worked when you see a screen that greets you by name, or a dashboard page that lists "Playground" as one of the options.

Step 2 — Open the Playground

Once you're signed in, look at the left-hand navigation menu and click the item labeled "Playground" (it may also be called "Try it now" or sit under a "Tools" header). A new screen opens with two big text boxes side by side: a prompt box on the left where you type your request, and a response box on the right where the AI's answer will appear. At the top of the panel there's a dropdown menu for choosing which AI model to use — leave it on the default for now.

Example: Cohere's "Command" family of models handles everyday translation tasks well; no need to change it for a product description.

If you don't see a left-hand menu, look in the top-left corner for a "hamburger" button — three horizontal lines — and click it. On a smaller screen or a narrow window, the menu hides behind that icon.

You'll know it worked when you see two empty boxes labeled something like "Prompt" (left) and "Response" or "Output" (right), with a "Generate" button near the bottom of the prompt box.

Step 3 — Paste your English copy and write a clear instruction

Click inside the prompt box on the left. Press Ctrl+V (Windows) or Cmd+V (Mac) to paste your English product description directly into the box — don't retype it, that introduces typos. Now add a short instruction above or below the pasted text, telling the AI exactly what you want: which languages, what tone, and how to label the answer.

Example prompt: "Translate the following product description into natural, shopper-friendly Spanish and natural, shopper-friendly French. Label each version clearly. Match the warm, persuasive tone of the original. Here is the description: [your English copy here]"

If your description has unusual characters (curly quotes, em dashes), don't worry — they paste fine. If the box only shows part of your text, that's normal; the AI can still read the whole input.

You'll know it worked when the pasted English text shows up inside the prompt box, with the same line breaks as in your original.

Step 4 — Generate both translations

Click the button at the bottom of the prompt box — look for "Generate", or a play-arrow icon (▶). The right-hand box will start filling with the AI's answer. Most product descriptions finish in under 30 seconds; longer ones may take up to a minute. When it's done, you'll see two clearly separated blocks: one Spanish version, one French version, each labeled.

Example of what a good output looks like: "Spanish: ... [paragraph] ... French: ... [paragraph] ..."

If nothing appears after a minute, click the button once more — the first click sometimes doesn't register. If you only get one language, the prompt may have been cut off; see Step 5.

You'll know it worked when you can read two separate paragraphs, one Spanish and one French, each with its own label.

Step 5 — Ask for a redo if anything feels off

If a translation sounds too stiff, too long, or just doesn't read like something your customers would say, you don't have to start over. In the same playground panel, click at the end of the previous answer (or in a new prompt box that appears below it) and type what you'd like changed. The model treats the new message as a follow-up, so it knows what you're reacting to.

Example follow-up: "Make the Spanish version about 20% shorter and a bit warmer — like a friend recommending the product. Keep the French version unchanged."

If the playground doesn't keep a chat history, paste the previous answer back into the prompt box and ask for the rewrite from scratch — slower, but it still works.

You'll know it worked when the new reply shows the changed version, and any language you asked to keep untouched is still there in the right language.

Step 6 — Copy each version into your store

Highlight the Spanish text in the response box, copy it (Ctrl+C or Cmd+C), and paste it into your store's Spanish-language product page or into a translations spreadsheet. Then go back and do the same with the French text. Save the originals somewhere you'll find again — a Google Doc, a Notion page, or a spreadsheet with one row per product and one column per language.

Example spreadsheet columns: English | Español | Français — one row per product, so any future edit stays in sync across all three languages.

If your store only has a single language field, save the translations in a notes app or a Google Doc titled with the product name, so you can find them later.

You'll know it worked when both Spanish and French versions sit in your store (or in your backup doc), readable, with no garbled characters, no English words left over, and a clear label on each.

Common mistakes

- **Pasting the description alone.** If you only paste the description with no instruction, you may get stiff, word-for-word translations. Always tell the AI the tone (friendly, professional, playful) and who's reading (shoppers, parents, professionals).
- **Trusting the output blindly.** Cohere is strong at everyday language but can mistranslate specialist words (a specific fabric, a regulatory term, a brand name). Read both versions out loud — anything you wouldn't say yourself is worth fixing.
- **Forgetting to save your English source.** If you close the playground and come back later, the original prompt may be gone. Keep your best English wording in a doc alongside the translations, so future edits have a single source of truth.
- **Mixing up the labels.** The AI sometimes merges Spanish and French into one block, or swaps the order. Always scan for the "Spanish:" and "French:" headers before you copy — that's how you know which version is which.

What this means for you

- **In everyday life:** If you ever need to email a French-speaking friend or read a Spanish recipe at a friend's house, the same Playground prompt works — paste the text, ask for a translation, done.
- **For your business or work:** A team of one can keep a small product catalog in three languages without paying a translator per product. You stay in control of the wording, and you ship faster.
- **If you're just getting started:** Pick your best-selling product today. Translate just that one. See how the result reads, then decide whether to roll the workflow out across the rest of your catalog.

Try it now

Open Cohere's Playground, paste this exact prompt, and replace the bracketed text with your own description: "Translate the following product description into natural Spanish and natural French. Label each version clearly. Here is the description: [paste your English product copy here]." Hit Generate, read the result out loud, and save the two versions somewhere you'll find them again.

How to compare two rental listings side by side with weighted pros and cons

6 steps · Read online

> **In short:** Open any AI assistant, paste both listings, list your priorities with weights, and ask for a scored comparison table. The assistant won't visit the apartment for you, but it will turn a wall of listing text into a clear trade-off table in minutes.

When you're choosing between two rentals, the listings blur together. This five-minute workflow works in any AI assistant — ChatGPT, Gemini, Claude, Copilot, Le Chat, DeepSeek, or whichever you already have — and scores both places against the priorities you set yourself.

Before you start

- An account with any AI assistant you already use (free or paid — both work for this).
- The full text of both rental listings, copied from the listing page or app. Screenshots alone won't work — only pasted text.
- A short list of what actually matters to you (commute, noise, utilities included, lease length, pets, parking — three to six items is enough).
- About 15 minutes.

Step 1 — Open a new chat and paste both listings

Start a fresh conversation — don't drop this into an old chat, or the assistant will weigh your rentals against whatever you were talking about before. In the message box at the bottom of the screen, paste the first listing, leave a blank line, then paste the second one. End the message with a one-line note telling the assistant what it's looking at. Example: "Listing A — 2-bed in Eastside, $X/month, 12-month lease. — — Listing B — 1-bed + den in Northgate, $Y/month, month-to-month. Compare these two rentals." If the listing is mostly images or a PDF, you'll need to type out the key facts by hand — assistants can't reliably read screenshots unless your plan includes file upload or vision, and even then the result varies. You'll know it worked when you can scroll up and see both listings sitting in the chat, and the assistant acknowledges them in its first reply.

Step 2 — Tell the assistant your priorities and how much each matters

Send a second message that lists the things you care about and gives each one a weight — how much it should count in the final score. Weights are how you stop the assistant from treating "has a dishwasher" as equally important as "20-minute commute." The format that works best is a numbered list with a 1–5 or 1–10 weight next to each priority. Example: "Here are my priorities, weighted 1–5: 1) Commute under 30 min by transit (5). 2) Quiet street, no main road (4). 3) Utilities included in rent (3). 4) Allows my cat (5). 5) In-unit laundry (2). Score both listings against these." If your assistant has a "memory" or "custom instructions" setting (different apps call it different things), you can save the weights there once and reuse them for future comparisons. You'll know it worked when the assistant restates your priorities back to you without being asked — that means it actually loaded them into the conversation.

Step 3 — Ask for a clean side-by-side table first, before any scores

Now ask for a comparison table with one row per feature and one column per listing. Don't ask for scores yet — you want a neutral snapshot of what each listing actually says, so you can spot places where the assistant fills in gaps with guesses. The clearer this table is, the more honest the scoring will be in the next step. Example: "Make a table with these rows: rent, lease length, square feet, utilities included, parking, laundry, pet policy, commute time to downtown, nearest noise source. One column for Listing A, one for Listing B, one for 'verified or missing.' If a fact isn't in the listing, write 'not stated.'" If the table renders as plain text instead of a real grid (some mobile apps do this), ask the assistant to "rebuild it with pipes and dashes" and it will produce a text-only version that still reads clearly. You'll know it worked when every row is filled with text from the listing — or says "not stated" — instead of the assistant inventing something to fill the gap.

Step 4 — Add weighted scores and a recommendation

Send a third message asking the assistant to score each listing on your weighted priorities and add a one-line recommendation at the bottom. This is the step that turns the table into a decision. The math is simple: multiply each priority's weight by the listing's score on that priority, then add them up. The listing with the higher total wins — *on your priorities*, not on someone else's. Example: "Now score both listings 1–10 on each priority from Step 2, multiply by the weight, and show the totals. Add a two-sentence recommendation naming which one wins on my priorities and why." If the assistant hedges with "it really depends on you," that's a sign it didn't read your weights carefully — nudge it with "use the weights literally, don't substitute your own." You'll know it worked when you see a final number for each listing and the recommendation names a clear winner with a stated reason.

Step 5 — Push back on vague answers and list what to verify

Before you act on the recommendation, ask the assistant to flag anything it had to guess at, plus every "not stated" cell from the table. This step is the safety net — assistants can sound confident about facts that aren't actually in the listing, and a lease you signed based on a wrong assumption is an expensive lesson. The honest answer often involves several things you need to ask the landlord directly. Example: "List every fact in the table you marked 'not stated' and anything you had to infer. For each one, write a one-line question I can email the landlord." If the assistant invents a detail (like "utilities are included" when the listing didn't say so), that's a hallucination — when an AI confidently states something that

isn't actually there. Call it out with: "You made that up; the listing didn't say it. Replace with 'unverified, ask landlord.'" You'll know it worked when you have a short, concrete list of questions to send before you sign anything.

Step 6 — Save the prompt structure as a reusable template

Copy the three prompts you sent (Steps 1, 2, and 4) into a notes app and label it "Rental comparison template." Next time you're weighing two places, open a new chat, paste the template, swap in the new listings and weights, and you'll get the same scored comparison in minutes — no setup, no re-thinking the structure. 💬 Example: "PROMPT 1: [paste listings + intro line]. PROMPT 2: [paste weighted priorities]. PROMPT 3: [paste score-and-recommend request]." If your assistant supports a way to save reusable prompts (different apps call this Projects, Gems, custom instructions, or similar — check the menus if you're not sure), you can store the template there so it loads automatically next time. You'll know it worked when the template is sitting in your notes ready to go the next time a listing site or a landlord's email hands you two options.

Common mistakes

- **Letting the assistant pick your priorities.** It will default to generic things like "location" and "price." Insist on *your* weights, even if they feel quirky — a quiet street might matter more to you than square footage, and that's the whole point of doing this yourself.
- **Pasting only the headline summary.** The first paragraph of a listing is marketing. The fine print about utilities, parking, and pet fees is usually further down — paste the whole thing, or you'll be comparing ad copy instead of facts.
- **Trusting the score without checking the table.** A 9/10 on "utilities included" means nothing if the listing never actually said that. Always read the table before you read the score, and send the "not stated" list to the landlord before you commit.

What this means for you

- **In everyday life:** When you're weighing two rentals (or two cars, or two phone plans), this same three-prompt pattern turns a messy comparison into a clear trade-off table in minutes — without you having to argue with yourself about which factor matters more.
- **For your work or business:** The same structure works for choosing between vendors, job offers, software plans, or any two options where you want *your* priorities to drive the score — not the assistant's defaults.
- **If you want to go further:** Try extending the prompt to three listings instead of two, or add a "dealbreaker" rule that automatically disqualifies any option missing a non-negotiable (like "no cat-friendly place counts, period").

Try it now

Pick one priority that matters to you (commute, pets, quiet, anything), find two short rental listings — even made-up ones work for practice — and run Steps 1 to 4 in a fresh chat. You'll see the decision become obvious in under five minutes.

How to use Gemini Deep Research to plan a big purchase decision

6 steps · Read online

> **In short:** Gemini Deep Research browses dozens of websites for you and returns a structured report with source links, so comparing big purchases like cars, appliances, or laptops stops being a full weekend of manual Googling.

Buying a new laptop, a refrigerator, or a used car usually means opening forty browser tabs and reading half of them. Gemini's **Deep Research** feature does that scrolling for you: it reads many sites, summarizes what it found, and hands you a tidy report with links back to where each claim came from. This guide walks you through using it, step by step, so the report actually answers *your* question.

Before you start

- **A Google account** signed in on the Gemini website or the Gemini app.
- **Access to Deep Research.** It may be limited to paid Google AI plans on some accounts, so check the plan page for the latest availability before relying on it.
- **A rough idea of what you're shopping for** (a category, a budget range, and 2–3 must-haves work best — more on this in Step 2).
- **About 15–25 minutes.** Deep Research usually takes a few minutes to run, and you'll want time to read the report it returns.
- An honest expectation: the report gives you a strong starting point, but it does **not** buy the thing for you, and prices or availability can be outdated by the time you read them.

Step 1 — Open Gemini and pick the Deep Research mode

Go to the Gemini website or open the Gemini app, then look in the chat box for a model or tool picker — usually a small dropdown above the text field. From that picker, choose the option labeled **"Deep Research."** On some accounts the option appears as a separate **Tools** button or a dedicated tab at the top of the chat. If you don't see it, it might not be enabled on your plan yet, or it could live behind a different name like **"Research"** or **"Pro."**

Example: when the picker appears, click the dropdown that currently shows "Fast" or "Pro" and switch it to **Deep Research**.

You'll know it worked when the text field gains a small brief outline or "Plan" panel, and the placeholder text inside the box mentions research or sources.

Step 2 — Describe your purchase in concrete detail

This is the most important step: the more specific your prompt, the more useful the report. Type a short paragraph that names the **item, your budget range, your must-haves, your deal-breakers, and how you'll use it.** Think of it like a one-page brief you'd give a friend who's shopping on your behalf — vague briefs produce vague reports. Avoid jargon the AI won't understand; instead, describe the thing in plain words.

Example: "I need a quiet, energy-efficient dishwasher for a small kitchen, budget roughly $700–$1,000 USD, must fit a 24-inch opening and run under 45 dB. I cook most nights for two, sometimes four. Compare the top 5 models available in my country, with pros, cons, and where each one tends to be cheapest."

You'll know it worked when you see a clear "Research plan" or outline appear below your message — Gemini usually shows what it intends to search for before it starts.

Step 3 — Approve the plan and let it run

Below the research plan, look for a button like **"Start Research"** or **"Begin."** Click it, then step away. Deep Research browses many websites in the background and writes as it goes, so the response can take **several minutes** (sometimes five, sometimes fifteen) — far longer than a normal chat reply. A progress bar or running indicator usually appears. **Do not close the tab** while it's running, or you'll lose the work.

Example: hit the "Start Research" button, then make a coffee; the first draft of the report appears inline as it finishes each section.

You'll know it worked when a long, structured report lands in the chat, with section headings, bullet points, and small numbered or linked citations you can click.

Step 4 — Read the report and click at least three sources

Don't trust any AI research report blindly. Scroll to the end of the report first and look for a **Sources** section, or for inline footnote-style numbers next to specific claims. Click at least three of those links and read the original page for a minute each, especially anything related to price, warranty, or safety. The sources tell you whether the report is repeating marketing copy or actually weighing real reviews.

Example: open the report, jump to the "Sources" list at the bottom, and click the first three links — if any of them are dead, irrelevant, or from a brand's own press release, treat the matching claim with caution.

You'll know it worked when you can point to a specific source link and say, out loud, "this is where that fact came from."

Step 5 — Ask two or three follow-up questions to narrow the field

Once the first report lands, treat it like a starting point, not the final answer. Ask Gemini to compare your top two candidates head-to-head, to look for known reliability complaints, or to translate the report into a simple "if you care most about X, choose Y; if you care about Z, choose W" decision guide. One question per message works best. If an answer feels too generic, push back: ask for specific model numbers, specific dates, or specific failure modes.

Example: "Now compare models A and B from your report on noise, energy use, and reliability over five years — which one wins for a household of two who cooks daily?"

You'll know it worked when Gemini's reply is sharper and more specific than the original report, and ideally points back to the same source list.

Step 6 — Turn the report into a personal decision checklist

Before you close the chat, ask Gemini to summarize the entire research session into a short **decision checklist** you can actually use: 3–5 finalist products, the deciding criteria for each, and the questions you still need to answer before you spend money. Copy that checklist into a note on your phone or print it. The point of the report is to leave your head clear, not crammed with new jargon.

Example: "Summarize this whole conversation into a checklist: top 3 dishwashers I should see in a store, what to test in person, and what to double-check on the manufacturer's site before I buy."

You'll know it worked when you have a one-page summary in your hand that you could read in two minutes in a showroom.

Common mistakes

- **Mistake:** Treating the report as a buy recommendation. **Fix:** It's a starting brief — the final choice is still yours, and any purchase should be confirmed against the official product page and the live price.
- **Mistake:** Skipping the sources. **Fix:** Always click at least three source links. If you can't find them, or they look unreliable, ask Gemini to redo the research with more trusted sites.
- **Mistake:** Closing the tab while research is running. **Fix:** Deep Research takes several minutes; leave the tab open until the full report appears, or you'll lose the work.
- **Mistake:** Asking a vague prompt like "best laptop under $1,000." **Fix:** Add your must-haves (size, weight, battery life, ports) and your main use (work, gaming, school) — the more detail, the better the report.

What this means for you

- **In everyday life:** A big purchase — a car, a fridge, a laptop — stops being a weekend of forty open tabs. You get a sourced brief instead, then spend your time on the decisions that actually need a human.
- **For your business or work:** The same workflow works for buying equipment, software subscriptions, or vendor comparisons at work: ask Deep Research to compare three vendors against your checklist, then read the sources yourself before signing anything.
- **If you're just getting started:** Pick one purchase you've been putting off, write a two-sentence brief in Step 2, and run the report end to end. The first run is the slowest; the second one feels easy.

Try it now

Open Gemini, switch to **Deep Research**, and paste this single sentence as your first prompt: *"I'm buying a [item] with a budget of [low]–[high], and my must-haves are [X, Y, Z]. Compare the top 5 options and end with a short decision checklist."* Fill in the brackets, hit Start Research, and read the sources before you trust any of it.

Catch Up Fast: Summarize Long Chats in WhatsApp and Messenger with Meta AI

6 steps · Read online

By the end of this guide you'll know how to ask Meta AI (a built-in helper inside WhatsApp and Messenger) to read through a long chat and give you a short, useful summary in plain language. This is for anyone who has ever scrolled back through hundreds of messages thinking *"what did I miss?"* — whether it's a family group chat or a busy work thread.

Before you start

- A phone with **WhatsApp** or **Messenger** installed, updated to the latest version (Meta AI is rolled out in stages, so update to be sure).
- A personal Meta account already logged in on that app.
- Meta AI available in your country or region — if you cannot see it after updating, it may not have reached your area yet.
- Around **5 to 10 minutes** of your time, plus the chat you want summarized open in another window or conversation.

One honest note up front: Meta AI **cannot read old messages you have not shared with it**. You will need to either paste the messages into a chat with Meta AI or @-mention it in a recent thread. It is free to use, but anything you send is processed by Meta — so avoid pasting very sensitive personal details.

Step 1 — Open WhatsApp or Messenger and look for Meta AI

Open the app you usually use (WhatsApp or Messenger) and look at the top of your chat list or your contacts. You should see Meta AI listed as an option, sitting alongside your usual contacts. Tap it once. A new chat window opens, just like starting a conversation with a friend. You'll know it worked when you see a message from Meta AI indicating it's ready to help. **If it looks different:** some users see a different icon or a "New chat" option with an AI symbol, or a search bar at the top — any of these can lead to Meta AI. If you see nothing related to AI, your app may need an update, or Meta AI may not yet be available in your region.

Step 2 — Open the chat you want summarized, side by side

Leave the Meta AI chat open, then open the long chat you want to catch up on. On most phones you can switch between apps with the square "recent apps" button, or use your phone's split-screen feature if available. You will see the full message history of your long chat on one screen and the Meta AI chat on the other. This makes copying easy. **If it looks different:** on tablets or split-screen devices you can arrange both windows side by side so both sit next to each other at the same time.

Step 3 — Copy the messages you want summarized

In the long chat, press and hold the first message you want Meta AI to read, then choose **Copy**. For a range of messages, keep tapping each one until you have selected the section you missed (often the last 20 to 50 messages is plenty). Then tap **Copy** again. The text now sits on your phone's clipboard, ready to paste — like copying a paragraph from a website. **If it looks different:** some Android phones show a small menu with "Copy", "Forward" and "Share"; on iPhone you'll see "Copy" appear above the selected message. If selecting one-by-one is too slow, try **Forward** instead, then paste into the Meta AI chat in one go.

Example: *"I copied messages from a 40-person family group covering about a day's worth of chat — roughly 30 messages with photos, links and a few voice notes."*

You'll know it worked when your phone briefly shows a small "Copied" label or a tick next to the selected messages.

Step 4 — Paste the messages into Meta AI and ask for a summary

Switch back to the Meta AI chat. Tap inside the message box, hold for a moment, and choose **Paste**. The copied messages appear as one long block. On a new line underneath, type a clear, kind request that tells Meta AI exactly what you want — think of this as giving instructions to a helpful assistant, not just dumping text on it. Then tap the send arrow. After a few seconds, Meta AI replies with a summary. **If it looks different:** if the paste does not appear, try tapping once inside the text box first, then holding your finger down again. If Meta AI replies in a language you did not expect, add "Please reply in English" at the end of your message.

Example: *"Here are the messages from my family group from yesterday. Please summarize them in 5 short bullet points, focus on any decisions or plans people made, and ignore the photos. Reply in plain English."*

You'll know it worked when Meta AI replies with a summary rather than echoing the messages back at you.

Step 5 — Refine the summary if it is too long or off-target

Read Meta AI's reply. If the summary is too vague, too long, or misses what you actually care about, reply with a short follow-up message asking for a change. The assistant remembers the conversation, so you can be casual about it. Each new request usually takes a few seconds. **If it looks**

different: if Meta AI seems stuck or repeats the same answer, try starting a brand-new chat and pasting the messages again with a clearer instruction — sometimes a fresh start produces better results.

Example: *"Great, but please make it 3 bullet points instead of 5, and add a one-line 'action items' section at the end showing what I might need to reply to."*

You'll know it worked when the new summary matches what you actually wanted to know — short, focused and useful.

Step 6 — Save or share the summary

Once you are happy with the summary, you can save it for later or send it to someone else. Tap and hold the assistant's message, then choose **Copy** to put it on your clipboard, or look for the forward option to send it to another chat or contact. Many phones also let you tap **Share** and save it straight to Notes, email or a document. **If it looks different:** if you don't see forward options, look for a menu icon or arrow next to the message. From there you can usually forward, share or save the message.

Example: *"I forward the summary to a private chat with myself called 'Catch-up notes' so I can search it later without scrolling through the original group."*

You'll know it worked when the summary now lives somewhere easy to find — your notes app, a saved message, or a chat you control.

Common mistakes

- **Pasting months of messages at once.** Meta AI has a limit to how much text it can read in one go, and very long pastes often get cut off or produce a generic answer. **Fix:** paste the last 20 to 50 messages only, or split a longer chat into two requests.
- **Sharing sensitive personal information.** Anything you paste into Meta AI is processed by Meta's servers, so names, addresses, phone numbers or financial details should be removed first. **Fix:** change real names to "Alex", "Sam" and "the host" before pasting, especially for group chats with people you do not know well.
- **Expecting Meta AI to know your whole chat history without pasting.** Unlike asking a friend, the assistant only sees what you have shared with it in that conversation. **Fix:** always paste the messages you want summarized, or @-mention Meta AI inside a recent thread so it can see the surrounding context.

What this means for you

- **In everyday life:** catch up on noisy family or friend group chats before replying, instead of scrolling through 200 messages on the train.
- **For your business or work:** turn long WhatsApp or Messenger threads with clients, teammates or event groups into a short briefing you can skim in 30 seconds and forward to a colleague.
- **If you're just getting started:** pick one group chat you fall behind on, copy the last 20 messages, paste them into Meta AI and ask for 3 bullets — that single step is the whole skill in miniature.

Try it now

Open your busiest group chat right now, copy the last 20 messages, paste them into a Meta AI chat, and ask: *"Summarize these in 3 bullet points and list anything I need to reply to."* It takes under two minutes and you'll have the skill for life.

Summarize a long PDF into 5 bullet points with Claude

6 steps · Read online

Ever faced a huge PDF document and wished you could get the gist of it without reading every single page? This guide will show you exactly how to use Claude, an AI assistant, to summarize any PDF file into just five easy-to-read bullet points. By the end, you'll be able to quickly understand important documents for work, study, or personal life.

Before you start

- You'll need a free Claude account. If you don't have one, it only takes a few minutes to sign up with an email address or Google account.
- Access to a computer or mobile device with internet.
- A PDF file you want to summarize.
- This guide assumes you are using the free version of Claude, which allows ample usage for this task.
- **Rough total time:** 6 minutes.

Step 1 — Get Ready for Claude

First, you need to open Claude in your web browser. Claude is an **AI assistant** – think of it like a very smart digital helper you can chat with, designed to understand your requests and provide helpful responses.

Open your web browser (like Chrome, Firefox, or Safari) and go to **claude.ai**. You'll see either a login screen or a chat interface if you're already logged in. If prompted, either sign in with your email and password or use a linked account like Google.

Once you're successfully logged in, you'll see Claude's main chat interface, which usually has a clear area to type messages and a history of past conversations on the left. You'll know it worked when you see the main Claude chat screen, ready for a new conversation.

Step 2 — Start a Fresh Conversation

To get the best results, it's always a good idea to start a new, clean conversation for each task. This ensures Claude focuses solely on your current request without being influenced by previous discussions.

Look for a button that says **"Start a new chat"** or **"New Chat"**. This is usually found in the top-left corner of the screen, often next to a Claude logo. Click on this button.

After clicking, the main chat area will clear, and you'll see a fresh, empty conversation window. A text box at the bottom will say something like "Message Claude..." or "Type your message here...". If it looks different, look for an icon like a plus symbol (+) or a speech bubble to begin a new discussion.

You'll know it worked when the chat panel is empty, ready for your first message.

Step 3 — Attach Your PDF Document

Now it's time to give Claude the document you want summarized. A **PDF** (Portable Document Format) is a common file type for documents, reports, and books, designed to look the same on any device.

In the message box at the bottom of the Claude chat window, look for an **attachment icon**. This usually looks like a **paperclip** or a **plus symbol (+)**. Click on this icon. A file explorer window will pop up from your computer. Navigate to where your PDF document is saved, select the PDF file, and then click **"Open"** or **"Choose"**.

You'll see the name of your PDF file appear above the message box, sometimes with a small preview or an "X" icon next to it, indicating it's ready to be sent. If you can't find the paperclip, try dragging and dropping your PDF file directly into the message box. You'll know it worked when your PDF's filename appears within the Claude chat interface, ready for your instruction.

Step 4 — Tell Claude What You Need

This is where you give Claude your **prompt** – which is simply the instruction you want it to follow. The clearer your prompt, the better Claude's response will be. You'll ask it to summarize the document into bullet points.

In the message box, where it says "Message Claude...", type your request. Make sure to specify the format (bullet points) and the number of points you want.

Example: `Summarize this document into 5 key bullet points. Please use simple, plain English.`

Once you've typed your prompt, double-check it to make sure it's clear and asks for exactly what you want. You'll know it worked when your text appears in the message box, alongside the attached PDF file, ready to be sent.

Step 5 — Send Your Request to Claude

With your PDF attached and your prompt typed, the next step is to send it to Claude so it can start working its magic.

Look for the **send button**, which is typically on the right side of the message box. It often looks like a **paper airplane** icon or a **right-pointing arrow**. Click on this button to send both your PDF and your instruction to Claude.

Claude will then begin processing your request. You'll usually see a message like "Claude is thinking..." or "Generating response..." while it works. This might take a few moments, especially for longer documents. If the send button seems unresponsive, ensure you have both the PDF attached and text in the message box. You'll know it worked when Claude starts typing out a response in the chat window, directly below your sent message and the attached PDF.

Step 6 — Review and Refine Claude's Summary

Claude will now present you with the summary based on your prompt. It's important to review this response to ensure it captures the essence of your document accurately.

Read through the 5 bullet points Claude has provided. Check if they cover the main ideas and if the language is indeed simple and plain as you requested. While Claude is very capable, it can sometimes **hallucinate** (confidently make up information that isn't in the document) or miss a subtle nuance, especially with very complex or highly technical texts.

If the summary isn't quite right, or if you'd like more detail on a specific point, you can ask Claude to refine it. For example, you could type: `Can you expand on the third bullet point?` or `Please rephrase the first point to be even simpler.` You'll know it worked when Claude presents a clear list of 5 bullet points in response, ready for your review or further instructions.

Common mistakes

- **Forgetting to attach the PDF or sending a blank message:** Claude can only summarize what you provide. Always make sure the document name appears above the message box before you hit send. If you sent a blank message, simply attach the PDF and send your prompt in a new message.
- **Using a vague prompt:** If you just say "Summarize this," Claude might give you a very general overview or a different format than you wanted. Be specific about the number of points and the desired language, like "Summarize into 5 bullet points, in simple English."
- **Expecting perfect accuracy without review:** While Claude is powerful, it's an AI and not infallible. It might occasionally misinterpret information or omit crucial details. Always read the summary carefully and compare it to your understanding of the document, especially for critical information.

What this means for you

- **In everyday life:** Instead of spending hours reading lengthy terms and conditions, user manuals, or community notices, you can quickly get the essential information in a few bullet points, saving you time and effort.
- **For your business or work:** Instantly summarize research papers, market reports, client proposals, or meeting minutes. This helps you and your team quickly grasp key takeaways, make faster decisions, and stay informed without getting bogged down in dense text.
- **If you're just getting started:** The easiest first move is to try it with a PDF you've already read. This way, you can easily compare Claude's summary to your own understanding and see how well it works.

Try it now

Find a PDF document on your computer – perhaps an old instruction manual, a news article you saved, or a school assignment. Head over to **claude.ai**, start a new chat, attach your document, and use the example prompt to get your first 5-point summary! You'll be amazed at how quickly you can distil information.

Summarize Long Articles into Key Bullet Points with Grok on X

6 steps · Read online

In today's fast-paced world, it's easy to feel overwhelmed by the sheer volume of information available online. This guide will show you how to use Grok, an AI assistant integrated into the X platform, to quickly condense long news articles or reports into concise bullet points. By the end, you'll be able to get the essential information from any article in moments, letting you stay informed without needing to read every word. While Grok can summarize the content of an article, it cannot verify the accuracy of the original article's claims or perform actions like sending emails for you.

Before you start

- An active X Premium+ subscription, as Grok is only available to subscribers of X's highest tier.
- The X app on your mobile device or access to the X website on a computer.
- Roughly 5-10 minutes of your time.

Step 1 — Access Grok on X

First, open the X app on your phone or visit the X website on your computer. Once the X feed has loaded, look for a special icon that represents Grok. On mobile, this is usually a diamond-shaped icon found in the bottom navigation bar. On a computer, you might find a "Grok" option listed in the left-hand sidebar menu, often near your profile or notifications. Tap or click on this icon to enter the Grok interface. If it looks different: The X interface can change, so if you don't see a diamond icon or a clear "Grok" label, look for an icon that stands out, or check the 'More' menu (often three lines or dots) for an entry labeled "Grok". You'll know it worked when you see Grok's chat interface, which usually includes a welcome message or a list of previous conversations.

Step 2 — Copy the Article Link

Now, you need to find the online article you want to summarize. Open your web browser (like Chrome, Safari, or Edge) and navigate to the news article, blog post, or report you're interested in. Once you're on the page, look for the web address (URL) at the top of the browser window. Tap or click on this address to highlight it, then use your device's copy function (e.g., right-click and 'Copy' on a computer, or long-press and 'Copy' on a phone) to save the link to your clipboard. If it looks different: Some apps have a 'Share' button (often an arrow pointing upwards or a three-dot menu) that includes a 'Copy Link' option. If you can't easily highlight the URL, look for that share option. You'll know it worked when the link is saved and ready to be pasted elsewhere.

Step 3 — Start a New Chat with Grok

With your article link copied, return to the Grok interface within the X app or website. To begin a fresh summarisation task, you'll need to start a new conversation. Look for a button that typically indicates starting something new, like a '+' symbol, a pencil icon, or a "New Chat" button, usually located at the top-right or bottom-right corner of the Grok chat screen. Tapping this will clear any previous messages and give you a clean slate. If it looks different: If you don't see an obvious 'new chat' icon, you can often just type your request directly into the existing message box. Grok is usually smart enough to treat a new request as a new topic. You'll know it worked when the chat area is empty and a message input box is clearly visible at the bottom of the screen, ready for you to type.

Step 4 — Paste Your Link and Request a Summary

Now it's time to tell Grok what you want. Click or tap inside the message input box at the bottom of the Grok chat screen. Paste the article link you copied in Step 2 into the box. After the link, type a clear instruction asking Grok to summarize the content. Be specific about the format you prefer, like asking for bullet points, to help Grok deliver exactly what you need. If it looks different: If pasting doesn't work, try typing the full URL manually if it's short, or check your device's clipboard settings to ensure the link was copied correctly. Grok might indicate if it's having trouble accessing the link. 💬 Example: `Please summarise this news article into 5 clear bullet points, focusing on the main findings: [paste your article link here]` You'll know it worked when Grok displays a message like "Generating response..." or "Thinking...", followed shortly by the summary appearing in the chat.

Step 5 — Review Grok's Summary

Once Grok has finished processing your request, its summary will appear in the chat window, presented in the bullet-point format you asked for. Read through these points carefully. This is your chance to quickly understand the core message, key arguments, and important facts from the article without having to read every sentence of the original text. Take a moment to ensure the summary captures what you expected and is easy to understand. If it looks different: Sometimes, Grok might ask a clarifying question if the article is too long, behind a paywall (which Grok cannot bypass), or contains ambiguous content. If this happens, try rephrasing your request or finding a more accessible article. You'll know it worked when you have a concise list of key information that helps you quickly grasp the article's main content.

Step 6 — Refine the Summary (Optional)

Grok's initial summary might be great, but you can always ask for more details or a different perspective if needed. If a particular bullet point sparked your interest, or if you feel something important was missed, simply type a follow-up question into the message box. You can ask Grok to expand on a specific detail, explain a complex term, or even rephrase the summary for a different audience (e.g., "Summarize this for a high school student"). This allows you to dig deeper without returning to the original lengthy article. If it looks different: If Grok says it can't find the information, it might be that the detail wasn't explicitly in the article, or it's too specific for a general summary. Try rephrasing your follow-up in simpler terms. 💬 Example: `Can you elaborate on the economic implications mentioned in bullet point 2?` You'll know it worked when Grok provides a more detailed or adjusted response, showing it's continuing to process the article's information based on your new prompt.

Common mistakes

- **No Grok access:** Grok is a premium feature, so if you can't find it or get an error, ensure your X subscription is the highest tier (X Premium+). The fix is to upgrade your X subscription.
- **Pasting a broken or inaccessible link:** If Grok says it can't access the article, check that the link is correct and that the article isn't behind a paywall or login that Grok can't bypass. The fix is to find a publicly accessible link or try a different article.
- **Vague summary request:** If Grok gives a less-than-useful summary, your request might have been too general. The fix is to be more specific in your prompt, like "Summarize into 5 bullet points" or "Focus on the environmental impact."

What this means for you

- **In everyday life:** Quickly catch up on news headlines or understand long reports from community groups without spending hours reading. It's perfect for when you're short on time but want to stay informed.
- **For your business or work:** Efficiently digest industry news, competitor analysis, or lengthy internal reports, allowing you to extract critical insights faster and make more informed decisions.
- **If you're just getting started:** Practice by picking a news article you'd normally skim, then use Grok to get a bullet-point summary. Compare it to what you would have highlighted yourself.

Try it now

Open the X app, navigate to Grok, and paste a link to any recent news article. Ask Grok to summarize it for you in three bullet points. See how quickly you can get the main gist!

Translate and Explain Any Restaurant Menu with Qwen on Your Phone

6 steps · Read online

> **In short:** Snap a photo of a foreign-language menu and Qwen will translate each dish into English, list the likely ingredients, and flag common allergens. The free version handles photos and text, but it cannot guarantee allergen safety — always confirm with staff if you have a serious allergy.

Here's how to turn a confusing wall of unfamiliar words into a clear, ordered list of dishes you can actually understand and order from, even when you don't speak the language.

Before you start

- A phone with a working camera (iPhone or Android).
- The free Qwen app installed, or a browser open to Qwen's website if you'd rather not install anything. No paid plan is required for basic menu translation.
- A clear, well-lit photo of the menu, or the ability to type or paste the menu text. Total time: about 5 minutes for a single menu.

Step 1 — Open Qwen on your phone

Tap the Qwen app icon on your home screen. If you haven't installed it yet, open your phone's app store, search for **"Qwen"**, and download the official app from the developer listed as Alibaba. When the app opens, you'll see a chat screen with a text box near the bottom and a short welcome message above it. If you're using a browser instead, the same chat layout appears once the page loads. If you don't see a chat box, look for a button labeled **"New chat"** or a + icon in the corner of the screen. You'll know it worked when the chat box is visible, empty, and ready for you to type or attach something.

Step 2 — Start a fresh chat

Tap **"New chat"** (or the + / pencil icon) so you begin with a clean conversation. A blank screen appears with just the input field at the bottom. This matters because earlier messages in a long chat can confuse the assistant when you switch topics. If you already have an old chat open, look for a small + or **"New"** label near the chat history list, sometimes hidden inside a side menu (a ☰ or ≡ icon). A fresh chat also means the response won't be colored by whatever you asked Qwen yesterday. You'll know it worked when the screen shows no previous messages and only an empty input field at the bottom.

Step 3 — Add the menu (photo or text)

Tap the **paperclip / attachment / + icon** next to the chat box and choose **"Camera"** to snap a fresh photo of the menu, or **"Photos"** to pick one you already took. Hold the phone steady, frame the dish list, and capture one section at a time if the menu is long — sharp photos translate far better than blurry ones. If you have the menu as text (maybe you copied it from a website), tap the chat box, **long-press**, and choose **Paste**. The image preview or pasted text appears inside the message bubble before you send. If you don't see a paperclip, look for an icon that looks like a **plus sign** or a **picture frame**. 💬 Example prompt to use after pasting a text menu: "Here's a menu from a small restaurant in Lyon, France." You'll know it worked when the photo or text shows up inside your message bubble, ready to send.

Step 4 — Ask for translation with ingredients and allergen notes

Tap the text field and type your request, then send it with the arrow or paper-plane button. In a few seconds, Qwen replies with a translated, dish-by-dish breakdown. Be specific in your ask: a one-line "translate this" usually gives less useful answers than a request that names what you want to know. If the response is too brief, tap **"Regenerate"** or send a follow-up asking for more detail. 💬 Example: "Translate this menu into English. For each dish, give a one-sentence description, list the main ingredients, and flag any common allergens like dairy, gluten, eggs, nuts, shellfish, soy, or pork." You'll know it worked when Qwen returns a clear list where each dish has a name, a description, ingredients, and an allergen note.

Step 5 — Ask follow-up questions about specific dishes

Read through the translation and pick anything that sounds unusual or risky for you. Tap the chat box, type a question about a single dish, and send it. Qwen keeps the menu in memory, so you don't need to re-upload or re-explain. This is the moment to ask about taste, preparation style, or whether something is vegetarian, spicy, or contains a hidden ingredient. If the answer is fuzzy, ask the same question again with sharper wording — for example, instead of "is this safe?", try "does this dish contain any peanut products?" 💬 Example: "What does 'confit de canard' usually taste like, and is it ever made with something other than duck?" You'll know it worked when Qwen gives a confident, specific answer that helps you decide whether to order it.

Step 6 — Save or share the translation

When the translation is complete and you have what you need, keep a copy for the table. On most phones you can long-press a Qwen reply and choose **Copy**, or take a regular screenshot by pressing the side buttons together. You can also tap **"Share"** (often a small arrow or three dots near the reply) and send the text to a travel companion, your own notes app, or a messaging group. If you'd like the menu translated into another language too (for example, to show a server in their own tongue), ask Qwen — it can switch languages without re-uploading the photo. You'll know it worked when you can scroll through the translation offline and show a single dish to your server without any guesswork.

Common mistakes

- **Forgetting to mention allergens.** If you only ask for a translation, Qwen won't flag allergens by default. Always include the line "and flag any common allergens" in your prompt, and if you have a serious allergy, double-check with restaurant staff — Qwen's guesses are helpful, not medical advice.
- **Sending a blurry photo.** Small, dim menu photos produce incomplete or wrong translations. Tap to focus before you shoot, hold steady, and capture one section at a time instead of the whole spread.
- **Asking for "the best" dish.** Qwen doesn't know your taste, your budget, or what your travel companions feel like eating. Ask for short descriptions instead, then decide yourself.

What this means for you

- **In everyday life:** the same trick works for grocery labels, medicine packaging instructions, and street signs when you travel. Once you've done a menu, the workflow is identical.
- **For your business or work:** if you host international guests or clients, you can use the same steps to translate a sample menu, brochure, or schedule from a foreign partner — and ask Qwen to flag anything culturally sensitive, like dishes that include alcohol or non-halal meat.
- **If you're just getting started:** don't worry about the prompt being perfect. Even a rough "translate this menu and tell me what's in each dish" gets you surprisingly far. You can always add detail in a second message.

Try it now

Open Qwen, snap a photo of the first menu you see today — even an English one — and ask it to list the ingredients of one dish you've never tried. You'll get a feel for the workflow in under two minutes, and you'll be ready for the real thing on your next trip.

Quickly gauge public mood with Hugging Face Sentiment Analysis

5 steps · Read online

Understanding the overall feeling behind customer feedback or social media posts can be a huge advantage for your business. This guide will show you how to use a tool called Hugging Face to quickly identify whether a piece of text is positive, negative, or neutral. By the end, you'll be able to get a rapid "mood check" on any comment, helping you respond effectively or spot trends.

Before you start

- An internet connection and a web browser.
- A piece of text (like a customer review or a social media comment) you want to analyze.
- A free Hugging Face account (sign-up is free and takes just a minute).
- Roughly 5-7 minutes to follow the steps and try it out.

Step 1 — Head to the Hugging Face website

Open your web browser and go to the official Hugging Face website. This is a popular online hub for AI models, like a big library where different AI "tools" (called **models**) are stored and shared. We're heading there to find a specific type of model that understands sentiment.

- **What you see:** You'll land on the Hugging Face homepage, which usually features a search bar, trending models, and links to various parts of their platform.
- **What happens next:** The page will load, ready for you to explore.
- **If it looks different:** The exact layout of homepages can change, but look for a central search bar or navigation menu (often at the top) to help you find your way.
- **You'll know it worked when:** You see the "Hugging Face" logo and a clear homepage in your browser window.

Step 2 — Find the Text Classification tool

On the Hugging Face website, we'll look for a specific type of AI tool that can categorise text. This is often called **Text Classification**, and it includes figuring out the emotional tone, or sentiment, of your words. Think of this tool as a smart assistant that sorts your text into "positive," "negative," or "neutral" piles.

- **What you see:** Look for a prominent search bar (usually at the top or center of the page) or a navigation menu option like "Tasks" or "Models."
- **What happens next:** If you use the search bar, type "Text Classification" and press Enter. If you're looking through menus, click on "Tasks" or "Models" and then locate "Text Classification" from the list of options. You'll be directed to a page dedicated to this function.
- **If it looks different:** If you can't find a direct "Text Classification" link or search result, try typing "sentiment analysis" into the main search bar. Then, look for a result that clearly offers an interactive "pipeline" or "widget" where you can input text.
- **You'll know it worked when:** You arrive on a page that features a clear text input box, usually with an example sentence already loaded, and a button like "Compute" or "Run inference" nearby.

Step 3 — Input your text for analysis

Now that you've found the interactive tool, it's time to put in the text you want to analyze. This text is called your **input**, and the AI model will "read" it to understand its emotional tone. Make sure the text is a sentence or a short paragraph for the best results from this online demo.

- **What you see:** A text box or field, typically labeled something like "Input text," "Enter your text," or similar. There might be some example text already there.
- **What happens next:** As you type or paste, the text will appear clearly in the input box. Any existing example text will be replaced by your new entry.
- **If it looks different:** The box might be larger or smaller, or appear in a different spot on the page, but its function is always to accept the text you want to process. Clear any existing example text before adding your own.
- Example: `I absolutely love this new coffee shop! The staff are so friendly and the flat whites are perfect every time.`
- **You'll know it worked when:** Your chosen text is clearly visible within the input box, ready for analysis.

Step 4 — Run the sentiment analysis

With your text entered, it's time to let the AI do its work. This step involves telling the AI model to process your input and predict its sentiment. The AI uses its internal "knowledge" (which it gained from being trained on vast amounts of text) to categorise your input.

- **What you see:** A prominent button, often labeled "Compute," "Run inference," "Submit," or simply "Analyze." It's usually located right below or next to the text input box.

- **What happens next:** The system will process your text. You might see a brief loading indicator (like a spinning wheel or a progress bar) while the AI model is thinking. This usually takes just a few seconds.
- **If it looks different:** The button's label might vary, but its purpose is always to trigger the analysis. If there's no visible button, sometimes hitting the 'Enter' key after typing in the text box can trigger it, or the analysis might run automatically after a short pause.
- **You'll know it worked when:** The loading indicator disappears, and new information appears on the screen, showing the results of the sentiment analysis.

Step 5 — Interpret the sentiment results

After the AI runs the analysis, it will show you the results. This typically includes a "label" (like POSITIVE, NEGATIVE, or NEUTRAL) and a "score" or "confidence level" (a number between 0 and 1, or 0% and 100%) indicating how sure the AI is about its prediction. This score tells you how much the AI "believes" its own classification.

- **What you see:** Below the input box, or in a separate results panel, you'll see a sentiment label (e.g., "POSITIVE") and a number next to it (e.g., 0.98 or 98%). This score tells you the **confidence** the model has in its prediction – closer to 1 (or 100%) means very confident.
- **What happens next:** You can now read and understand the AI's assessment of your text's mood. You'll instantly see if the AI thinks your comment is generally positive, negative, or neutral, along with its certainty.
- **If it looks different:** The results might be displayed as a bar chart, a simple list, or a phrase. Look for keywords like "label," "score," "positive," "negative," or "neutral" to find the key information.
- **You'll know it worked when:** You see a clear sentiment label and a confidence score for your input text, giving you an immediate understanding of its emotional tone.

Common mistakes

- **Too much text at once:** The interactive widgets on Hugging Face are designed for short snippets. Trying to analyze an entire document or a very long review might lead to an error or incomplete results. Stick to a few sentences or a short paragraph at a time for these online demos.
- **Misinterpreting the score:** A score like "0.55" for "POSITIVE" means the AI is only 55% confident it's positive, which isn't very strong. Aim for scores above 0.75 (75%) for reliable sentiment. Lower scores indicate ambiguity, meaning the text is hard to classify, or that the text is truly neutral.
- **Expecting perfect understanding:** AI models aren't human. They can struggle with sarcasm, irony, or highly nuanced language. If a review says "Great customer service, if you like being ignored for an hour," the AI might just pick up "Great customer service" as positive, missing the sarcastic context. The tool helps, but a human eye is still valuable. Also, this tool can draft the sentiment analysis, but it cannot respond to the customer for you.

What this means for you

- **In everyday life:** If you're sifting through comments on a news article or a product you're considering buying, a quick sentiment check can help you grasp the general public opinion without reading every single word in detail. It's like having a quick emotional barometer for online discussions.
- **For your business or work:** You can quickly scan customer reviews for new products, analyze social media mentions of your brand, or get an instant feeling from survey feedback. This helps you identify common pain points (if negative sentiment is high) or celebrate successes (if positive sentiment is dominant), allowing for faster decision-making about what actions to take.
- **If you're just getting started:** Try taking a few recent comments from a local business's social media page or website reviews. Paste them into the sentiment analysis tool one by one and see how accurate the AI's prediction feels to you.

Try it now

Copy a recent comment from a local cafe's Facebook page or a quick customer review for a product you recently bought. Go back to the Hugging Face Text Classification tool (Step 2), paste your text (Step 3), and hit 'Compute' or 'Analyze' (Step 4) to see its sentiment.

Unpack Complex Ideas with DeepSeek's Simple Explanations

5 steps · Read online

Ever felt lost trying to understand a tricky topic? DeepSeek can be your personal explainer, helping you grasp even the most complicated subjects by breaking them down into simple, easy-to-digest language. This guide will show you exactly how to ask DeepSeek to simplify anything, from science concepts to historical events.

Before you start

- You'll need a device connected to the internet (computer, tablet, or smartphone).
- Access to the DeepSeek website. It's a free-to-use **AI assistant** (a computer program designed to understand and respond to human language, like a very clever digital helper) for this task, though you may need to create an account for full access.
- This guide will take about 5-10 minutes to follow.

Step 1 — Go to the DeepSeek Website

Open your preferred internet browser (like Chrome, Safari, or Firefox) and navigate to the DeepSeek chat website. Create an account if prompted, then you're ready to start asking questions.

You'll see a clean interface, usually with a large chat box in the center and perhaps a few example questions. This chat box is where you'll type your requests, known as **prompts** (the instructions or questions you give to an AI). If it looks different, don't worry – most AI chat interfaces have a clear text box ready for your input.

You'll know it worked when you see a blank text area or a message like "Ask me anything" in the main part of the screen, ready for you to type.

Step 2 — Ask for a Simple Explanation

Now, it's time to ask DeepSeek to explain something for you. Think of a topic you find confusing. Type your request directly into the large text box at the bottom of the screen. It's best to be clear that you want a simple explanation.

Once you've typed your request, you'll usually see a "Send" button (often an arrow or a paper plane icon) next to the text box. Click this button. If you don't see a "Send" button, pressing the `Enter` or `Return` key on your keyboard usually works too. DeepSeek will then process your request, which might take a few seconds.

You'll know it worked when you see DeepSeek starting to type its response directly above your input box.

Example: "Explain quantum physics in simple terms for someone who knows nothing about science."

Step 3 — Review DeepSeek's Explanation

DeepSeek will present its explanation in the chat window. Read through it carefully. The goal is for you to understand the core concepts without getting bogged down in jargon. Pay attention to any terms that still seem unclear, even after DeepSeek's attempt at simplification.

After DeepSeek finishes typing, its full response will appear in the chat history. If the explanation seems too complex or uses words you don't understand, don't worry – we'll address that in the next step. If it looks different, perhaps the text might be in a slightly different font or layout, but the content will be clear.

You'll know it worked when you have a complete explanation from DeepSeek displayed in the chat window, ready for you to read.

Step 4 — Request Further Simplification or Clarification

If DeepSeek's first explanation still left you scratching your head, or if there's a specific term you'd like it to break down even further, just ask! You can follow up in the same chat. DeepSeek remembers your previous questions and its answers, so you don't need to re-explain the whole topic.

Type your follow-up question in the text box and click the "Send" button again. DeepSeek will then generate a new response, focusing on your specific area of confusion. If the AI seems to forget the context, you might need to briefly remind it of the topic in your follow-up question.

You'll know it worked when DeepSeek provides a refined or more detailed explanation based on your specific request, appearing below the previous response.

Example: "That's helpful, but can you explain 'wave-particle duality' from the previous explanation even more simply? Use an everyday analogy."

Step 5 — Ask for a Summary or Different Perspective

Sometimes, understanding comes from seeing the same information presented in a different way or in a very concise summary. You can ask DeepSeek to rephrase its explanation, provide a summary, or even give you a different analogy. This is a great way to solidify your understanding.

Type your request into the chat box and send it. DeepSeek will then process your request and present the information in a new format. If you don't like the new format, you can always ask it to try another, or go back to the original explanation in the chat history.

You'll know it worked when DeepSeek displays a summary, a rephrased explanation, or a new analogy for the topic, helping you to look at the information from a fresh angle.

Example: "Can you summarize the main points of quantum physics in three bullet points?"

Common mistakes

- **Using too much jargon in your prompt:** If you use complex terms in your initial question, DeepSeek might assume you already understand them and use similar language in its answer. **The fix:** Use simple, everyday words in your request. If you want it explained for a beginner, say so clearly.
- **Expecting a deep, personalized teaching experience:** DeepSeek is an AI assistant, not a human tutor. It can't sense your frustration or adapt its teaching style based on non-verbal cues. **The fix:** Be specific in your follow-up questions. If you're still confused, point to the exact part of the explanation that's unclear.
- **Not specifying the level of explanation:** If you just say "Explain X," DeepSeek might give a general answer that's still too advanced or too basic. **The fix:** Always include phrases like "in simple terms," "for a beginner," or "as if I'm 10 years old" to guide its response.

What this means for you

- **In everyday life:** Instead of getting stuck on a complex news article or a tricky instruction manual, you can quickly get a simple breakdown from DeepSeek. This helps you understand everything from how your new gadget works to the latest scientific discovery.
- **For your business or work:** If you come across industry jargon or a complex report, DeepSeek can help you understand the core concepts quickly. This boosts your comprehension without needing to spend hours researching every single term, helping you make quicker, more informed decisions.
- **If you're just getting started:** The easiest first move is to pick one confusing topic you've heard about recently and ask DeepSeek to explain it as simply as possible. See how quickly you can grasp the basics!

Try it now

Open DeepSeek and ask it to explain "blockchain technology" in a way that someone who's never heard of it can understand.

How to Ask Gemini to Explain Your Bank Statement in Plain English

6 steps · Read online

> **In short:** You can upload a bank statement into Google's Gemini and ask plain-English questions to surface real spending, recurring charges and hidden fees. Gemini can't read your bank account — think of it as a translator for the numbers on your statement.

A monthly bank statement is full of small numbers, merchant codes and fee names that rarely mean anything at first glance. With Google's Gemini, you can hand it that document and ask, in normal sentences, "what am I actually paying for?" — and get a readable answer back. This guide walks you through it, one step at a time.

Before you start

- A free Google account (the same one you use for Gmail or YouTube).
- Access to Gemini in a browser (gemini.google.com) or the Gemini mobile app — both work for this. The free version is enough for reading a single statement.
- A copy of your monthly bank statement, saved as a PDF or a screenshot — most banks let you download it from their website or app.
- About 10 minutes the first time; once you've done it once, the same questions work on next month's statement.

Step 1 — Open Gemini and start a fresh chat

Go to **gemini.google.com** in your browser, or open the **Gemini app** on your phone. If you are not signed in, the page will ask you to sign in with your Google account — the same email and password you use for Gmail. Once you are in, look for a button or link along the left side or top of the screen that says something like **"New chat"** or **"Start a new chat"** and click or tap it. A blank conversation opens. If you don't see that exact label, look for a **pencil icon**, a **"+" sign**, or a text box that says **"Ask Gemini"** — any of those will start a fresh conversation. Example: open the Gemini app, tap "New chat," and you should see an empty text box with the prompt "Ask Gemini" at the bottom. You'll know it worked when you see an empty text field with no previous messages above it.

Step 2 — Upload your bank statement

Look for a small **paperclip icon**, a **"+" button**, or an **"Upload"** option near the text box — every version of Gemini has a way to attach a file. Click it, then choose the **PDF** of your statement from your computer or phone. If your statement is a stack of screenshots instead, you can either upload them one at a time or, on the desktop site, drag and drop them all into the chat window together. The file name appears in your message once it uploads. If you can't find an upload option at all, your version may only accept **pasted text** — open your PDF, select all the text (Ctrl+A or Cmd+A), copy it (Ctrl+C or Cmd+C), and paste it (Ctrl+V or Cmd+V) directly into the chat box instead. Example: "I'm pasting my bank statement below — it's about 3 pages of transactions." You'll know it worked when a small thumbnail or file name shows up next to your message, and Gemini replies with something like "Got it, I'm reading the statement now."

Step 3 — Ask for a plain-language summary

Now type a first question in the same chat. Don't ask for anything fancy yet — just ask Gemini to read the statement and tell you, in normal words, what happened in the month. Think of it as asking a friend who is good with numbers to glance at the page and explain it. Be specific about the period, and tell Gemini how you'd like the answer — short bullets, full sentences, whatever you prefer. If Gemini answers with generic budgeting advice instead of reading your actual statement, it probably didn't see the file — re-upload it or paste the text directly into the chat. Example: "Please read the statement I just uploaded and give me a plain-English summary of what happened in November. List the top 3 things I spent money on and any fees I was charged. Keep it under 200 words." You'll know it worked when Gemini replies with a readable summary that mentions real names and amounts from your statement — not a generic answer about budgeting.

Step 4 — Ask about recurring charges and small fees

Now go one level deeper. Small monthly fees — a $4.99 streaming service, a $12 "account maintenance" charge, a $2 ATM surcharge — are easy to miss when you scroll quickly. Ask Gemini to look specifically for anything that appears more than once, or anything that looks like a fee rather than a purchase. This is the step where most people discover charges they've been paying for months without noticing. If Gemini says it can't see any recurring charges, the file may have been only partially read — ask it to look at every page, or paste the full transaction list into the chat. Example: "Now look only at the charges that happened more than once this month, and anything that looks like a fee or service charge. For each one, tell me the name, the amount, and how many times it appeared." You'll know it worked when you get a short list of repeated merchants and fee names you can scan in 30 seconds — and probably one or two that surprise you.

Step 5 — Ask about your real spending patterns

The next question is about categories. Gemini can group your transactions by type — food, transport, subscriptions, cash withdrawals — and show you where most of the money went. This is helpful even if your statement already has a "category" column, because Gemini's version often groups things more sensibly for everyday reading, and you can ask follow-up questions about any category that looks too big or too small. If Gemini

lumps almost everything into "other," your statement may have lots of unfamiliar merchant names — ask it to list the biggest "other" items so you can see what got hidden. Example: "Group all my spending from this statement into simple categories like food, transport, subscriptions, bills and other. Show me the total for each category and the percentage of my total spending. Then tell me which category grew the most compared to a normal month." You'll know it worked when you see a small table or list with named categories, dollar totals and clear percentages — something you could read once and remember.

Step 6 — Save your insights or turn them into questions for your bank

The last step is making sure you don't lose what you just learned. If something in the summary or the fee list looked wrong, copy Gemini's exact wording of that charge — the merchant name, the date, the amount — and use it next time you contact your bank. You can also ask Gemini to rewrite the whole summary into an email or a checklist you can save. To come back to this conversation later, look in the left sidebar for your recent chats — your version of Gemini should keep a running list you can reopen. If you can't find it, bookmark the page in your browser before you close the tab. Example: "Turn the summary from Step 3 and the fee list from Step 4 into a short email I can send to my bank asking about the $12 monthly fee. Keep the tone polite and include the date and amount for each charge." You'll know it worked when you have either a saved conversation, a copyable email draft, or a clear list of questions to bring to your bank's customer service.

Common mistakes

- **Uploading a statement from a bank that masks real names.** Some banks replace merchant names with codes (like "POS DEBIT 8842"). Gemini can't decode those on its own — you'll get a vague answer. If this happens, ask your bank for an itemized version.
- **Asking one giant question instead of several small ones.** A request like "tell me everything about my statement" usually produces a fuzzy answer. Three or four short, focused questions (like in Steps 3, 4 and 5) get sharper results.
- **Forgetting to remove sensitive info if you share the chat.** A bank statement has your account number, sometimes your address and always your full name. If you ever screenshot a Gemini chat to share it, blur those parts first.
- **Treating Gemini's answer as a final verdict.** It can misread a number, miss a charge or guess wrong about a fee. Always check the original statement before calling your bank — Gemini helps you ask better questions, it doesn't replace them.

What this means for you

- **In everyday life:** the next time a charge looks unfamiliar on your statement, you can paste it into Gemini and ask, in one sentence, "what is this $9.50 charge from XYZ?" before you spend 20 minutes on hold with customer service.
- **For your business or work:** if you run a small business and review statements monthly, this same set of questions can become a 5-minute weekly check instead of a 90-minute pile of receipts.
- **If you're just getting started:** try it on a statement you already understand well — your last month, perhaps. That way you can quickly see whether Gemini's summary is accurate, before you trust it on a more confusing one.

Try it now

Open Gemini, upload this month's bank statement, and paste this single prompt: "Read this statement and list every charge I don't recognize, with the date and amount. Keep it to 10 items or fewer." Ninety seconds, one question — and you'll likely learn at least one thing you didn't know you were paying for.

How to Research a Medical Procedure or New Medication with Perplexity

6 steps · Read online

> **In short:** Perplexity can turn confusing medical jargon into plain English and help you build a question list for your doctor. It is not a doctor, and its answers can be wrong — always confirm what you find with a healthcare professional before acting on it.

When your doctor mentions a procedure you've never heard of, or hands you a prescription for a medication whose name you can't pronounce, the first thing most people do is type it into a search engine. The trouble is, the results are usually full of scary forum stories, ads for law firms, or pages of medical jargon that make your eyes glaze over. This guide shows you how to use Perplexity — an AI answer engine that pulls from current web sources and writes back in plain language — to research a medical topic calmly and walk into your next appointment actually prepared.

Before you start

- A free Perplexity account (you can use it without signing in, but logging in saves your searches). A paid plan adds deeper reasoning and more sources, but is not required for this guide.
- A web browser or the Perplexity mobile app (free on iOS and Android).
- The name of the procedure or medication you want to research — copy it from your doctor's note, prescription label, or patient portal.
- About 15–20 minutes for your first try.

Step 1 — Open Perplexity and start a fresh thread

Open **perplexity.ai** in a browser, or launch the Perplexity app on your phone. You'll see a large search bar near the top of the screen, often with a row of suggested questions underneath it. Type your first prompt into that bar and press **Enter**, or tap the arrow icon, to send it.

Example: "Explain laparoscopic gallbladder removal in plain English. I'm not a doctor."

If you're on a phone and the app opens to your old searches, look for a **"New Thread"** or **"+"** button — usually at the top or bottom of the screen — to start clean. Perplexity lets you choose how it answers; on the free plan a fast default model handles your question, while paid plans offer stronger reasoning options. For a first pass, the default is plenty.

You'll know it worked when the search bar is empty and a fresh page is ready for your question.

Step 2 — Ask for a plain-English explanation, not a textbook one

Type the medical term or medication name into the search bar and add a phrase that asks for plain language, then hit **Enter**. After a few seconds, an answer appears, with small footnote numbers next to sentences that link to sources listed at the bottom of the page.

Example: "What is metformin? Explain it like I'm smart but not a doctor. What does it actually do in the body, and what are the most common side effects people notice in the first month?"

The trick is the second half of your prompt — "explain it like I'm smart but not a doctor" tells the tool what tone you want. Without that hint, you often get a Wikipedia-style summary. With it, you get something closer to a friend's explanation. If the answer still feels dense, look for a button to regenerate or simplify the response — different versions of the same answer can read very differently.

You'll know it worked when you can read the answer through without reaching for a dictionary.

Step 3 — Ask for a short list of questions to bring to your doctor

In the same chat thread (don't start a new one), type a follow-up asking for questions to bring to your next appointment. The new answer appears right below your previous question. This is where Perplexity really earns its keep — it can play the role of a thoughtful friend who's already done a lot of research.

Example: "Give me 6 smart questions I should ask my doctor before agreeing to this surgery. Make them practical — about recovery time, what could go wrong, and what I'd need to arrange at work and at home."

If the thread doesn't continue, or you accidentally closed it, click **"New Thread"** and paste a one-line summary of your previous answer as context. Perplexity reads the whole thread, so follow-up questions can reference what came before.

You'll know it worked when you have at least five specific questions you would not have thought of on your own.

Step 4 — Switch to a deeper source mode for medical topics

Look near the search bar for a small **"Focus"** dropdown or mode selector — it usually looks like a tiny filter or label. Perplexity offers different source modes; for medical research, look for one labeled **"Academic"** or **"All"** if available, since the **"Social"** mode pulls from forums and is

rarely what you want for health questions. Select one, then ask your next question.

💬 Example: Pick the Academic or All focus, then type: "What does recent research say about metformin and weight loss? Are there any big studies from 2024 or later?"

The mode affects which sources Perplexity pulls from. The standard web search is fine for general explanations, but for medication questions, academic sources and reputable health sites tend to be more reliable than random web pages. If you don't see a Focus selector on your version of the app, the default still pulls from real web sources with citations — keep going.

You'll know it worked when the cited sources include recognizable medical publishers, hospital systems, or government health sites.

Step 5 — Sanity-check the most important claim

Before you take anything to your doctor, pick the single most worrying claim from your research — a side effect, a recovery time, or a risk — and ask Perplexity to confirm or contradict it. The point is not to find the perfect answer; it's to see whether multiple sources agree.

💬 Example: "You said recovery from this surgery is usually 1–2 weeks. Does that match what major hospital systems say, or is it more commonly 3–4 weeks?"

If two searches from different angles return conflicting answers, that's a signal the topic is genuinely uncertain — and a perfect question to bring to your doctor. If they mostly agree, you can be more confident about that point and focus your appointment on what still feels unclear.

You'll know it worked when you can point to at least two sources saying similar things about the most important point.

Step 6 — Save your research before the appointment

Before you close the tab, look for a **"Share"** button, a link icon, or simply copy the page URL from your browser. Perplexity threads have their own web address, and if you're logged in they're saved to your account automatically. On mobile, tap **"···"** or **"Share"** and send the thread to your notes app or email.

If you can't find a share option, the simplest fallback is to copy-paste the question list from the answer into your phone's notes app, or take a screenshot. The goal is to have the questions ready in your hand when you sit down with your doctor — not buried in a search history you can't find.

You'll know it worked when you can open your notes app and see your questions waiting for you, no internet required.

Common mistakes

- **Trusting the first answer without checking sources.** Perplexity cites its sources at the bottom of every answer — scroll down and glance at them. If none of the sources look reputable (a major hospital, a government health site, a medical journal), treat the answer with caution.
- **Asking Perplexity to diagnose you.** This tool can explain terms and list questions, but it cannot examine you, run tests, or know your full medical history. If you describe symptoms, treat any answer as a starting point for a conversation with a real doctor, never as a diagnosis.
- **Forgetting to bring your notes to the appointment.** A question list only helps if it's in front of you during the visit. Save it somewhere you'll actually see it — your phone's notes app, a printed page, or a photo on your lock screen.

What this means for you

- **In everyday life:** The next time a doctor mentions something you don't recognize, you'll have a 5-minute way to understand it before you walk into the office — instead of nodding along and searching it later at 11pm.
- **For your business or work:** If you're helping an older relative, a friend, or even managing your own healthcare decisions, this same workflow can save you from feeling pressured into agreeing to something you don't fully understand.
- **If you're just getting started:** Don't try to research everything at once. Pick the one thing that's bothering you most — a side effect, a recovery time, a cost — and run that single question through the steps above. You'll get comfortable with the rhythm quickly.

Try it now

Pick one medical term from a recent doctor visit, prescription, or patient message. Open Perplexity, type it in with the phrase *"explain it like I'm smart but not a doctor"*, and see what comes back. Read the answer, scroll to the sources, and ask yourself: *"Is this something I'd actually want to ask my doctor about?"* That's your first follow-up question — and your guide is already working.

Research and Draft Short Social Media Captions with Perplexity

6 steps · Read online

By the end of this guide, you can turn a fuzzy idea into a ready-to-post social media caption backed by real sources — whether you manage a personal account or post on behalf of a small business. This is for beginners who want to save time without sounding robotic.

Before you start

- **A free Perplexity account** — sign up at the Perplexity homepage with an email or social login. A free plan is enough for this guide; Perplexity Pro (paid) gives more daily searches and access to advanced AI models.
- **A device with a browser** — Perplexity works on desktop, mobile web, and the iOS/Android apps. Screens look similar across all three.
- **A topic in mind** — even something rough, like "sustainable coffee" or "weekend hiking near Nashville".
- **Rough total time: 15–20 minutes.**

> **One honest note:** Perplexity can research facts, summarize articles, and draft captions, but it can't post to social media for you — you'll still copy the caption into Instagram, LinkedIn, X, or your scheduler.

Step 1 — Open Perplexity and start a new thread

Open Perplexity in your browser or app, then sign in if prompted. Look for a **"New Thread"** button or a clean search box at the top of the screen — it usually has placeholder text like *"Ask anything..."* or *"What do you want to know?"*. Click or tap inside the box. The cursor blinks, ready for typing.

If your screen looks different, look for a **pencil icon**, a **"+" button**, or a **"Home"** tab that returns you to a blank search box. Each Perplexity conversation is called a "thread", so starting fresh keeps your research tidy.

You'll know it worked when the page reloads with an empty box and no previous answers above it.

Step 2 — Ask Perplexity to research trending angles for your topic

Type a clear, specific question into the search box. Beating your hands at the keyboard here won't help — give Perplexity a **topic**, an **audience**, and a **goal** in plain English. After you press Enter (or tap the arrow icon), Perplexity starts writing a short answer with **numbered footnotes** — those tiny superscript numbers are links to the sources it used. Below the answer, you'll often see a list of **related questions** you can click for a deeper dig.

If the answer feels too generic, try the **"Focus"** dropdown (look for buttons labeled **Web**, **Social**, **Academic**, or **All**) — switching to **Social** tells Perplexity to lean on posts and trending discussions rather than formal articles.

Example: *"I'm writing a 3-sentence Instagram caption for a small American coffee roastery. What are 5 trending topics in sustainable coffee right now in the US, with sources from 2025 or 2026?"*

You'll know it worked when Perplexity returns an answer with at least 3 numbered footnote links you can tap to read the underlying articles.

Step 3 — Pick one angle and ask for supporting facts

Scroll through the answer and choose **one** angle that excites you — don't try to cover all five. Click the related question that matches your pick, or type a new follow-up question into the same box. A new answer appears below the first, still with its own numbered sources. The conversation history stays visible, so you can scroll up and compare.

If a footnote link looks interesting, click it — Perplexity opens the source in a new tab. Don't worry about reading every word; skim for a quote, a stat, or a phrase you might want to mention.

Example: *"Tell me more about the rise of compostable coffee pods in the US. Give me 3 specific statistics or quotes from American sources I could mention in a post."*

You'll know it worked when you have one sentence or stat from a source you'd be happy to credit in your caption.

Step 4 — Switch into "drafting mode" with a clear instruction

Now you've done the research — time to write. Stay in the same Perplexity thread and type a new prompt. This time, you're telling Perplexity to **write**, not research. Be specific about the **platform**, **length**, **tone**, and what facts to include. Perplexity drafts a caption right inside the chat.

If the tone feels off, try adding words like "warm", "playful", "matter-of-fact", or "professional" to your instruction. Perplexity works best when you tell it exactly how the caption should sound.

Example: *"Write a 3-sentence Instagram caption for our compostable coffee pods. Tone: warm and friendly, not salesy. Mention that 1 in 3 people now look for compostable packaging, according to a 2025 survey. End with a question to invite comments."*

You'll know it worked when Perplexity produces a short caption (3 sentences or so) that includes your chosen fact and ends with a question or call-to-action.

Step 5 — Refine the draft in the same conversation

Read the caption out loud — does it sound like you? If a sentence feels clunky, type a quick refinement request right below the draft. Perplexity updates the answer in place, and the old version scrolls up so you can compare. This back-and-forth is normal and usually takes 2–3 quick rounds.

If Perplexity suggests adding emojis and you'd rather skip them, just say so. The same goes for hashtags — you can ask it to suggest 3–5 relevant ones or to leave them out entirely.

Example: *"Make sentence 2 shorter and a bit punchier. Keep the question at the end. Suggest 5 relevant hashtags."*

You'll know it worked when the caption sounds natural when you read it aloud, includes your key fact, and fits the platform's vibe.

Step 6 — Save your sources and copy the caption

Before you leave Perplexity, capture the **footnote links** from Steps 2 and 3 — you may want to credit a source in your post, or come back to read more later. Look for the small numbered superscripts next to sentences; click them to copy the URL from your browser's address bar.

Then copy the caption itself. Look for a small **copy icon** (two overlapping squares) near the draft, or simply highlight the text with your mouse or finger and copy it. Paste it straight into Instagram, LinkedIn, X, or your scheduling tool.

If Perplexity's Pro plan would help you here, note that the **free plan has a daily cap on Pro searches**, but regular searches are unlimited for most users — you can complete this whole guide on the free tier.

You'll know it worked when you have the final caption ready to paste, plus 2–3 source links saved in case you want to reference them later.

Common mistakes

- **Mistake:** Asking one giant question like *"Write me a caption about coffee."* **Fix:** Always include the **topic**, the **audience**, the **platform**, and the **length** — specificity saves you a round-trip.
- **Mistake:** Trusting Perplexity's footnotes without checking one or two. **Fix:** Click at least one source to confirm the stat or quote actually says what Perplexity implied. AI can misread a source.
- **Mistake:** Forgetting to add a human touch. **Fix:** After copying the caption, change one small thing — a word, an emoji, your own opener — so it sounds like *you*, not a template.

What this means for you

- **In everyday life:** Posting about a hobby, a trip, or a family milestone becomes a 10-minute task instead of a 30-minute staring-at-the-screen session.
- **For your business or work:** You can keep a steady posting rhythm without paying an agency, and every caption can be backed by a real source — handy if you ever get asked "where did you get that number?"
- **If you're just getting started:** Treat the first caption as practice, not a masterpiece. The second one will already feel easier.

Try it now

Open Perplexity, type this exact prompt, and watch what comes back:

> *"I'm writing a 2-sentence LinkedIn post for a small accounting firm in Portland. Find 1 recent statistic about AI tools being used by American small businesses in 2025 or 2026, with a source. Then draft the post in a confident but friendly tone, ending with a question."*

You should have a researched, source-backed draft within 2 minutes.

Compare Products with Perplexity Before You Buy

6 steps · Read online

Want to make sure you're getting the best value or the right features before spending your hard-earned cash? This guide will show you how to use Perplexity, an AI-powered search engine, to quickly compare products, understand their differences, and even check user opinions, all in one place. By the end, you'll be able to compare anything from new phones to kitchen appliances, helping you shop smarter and more confidently.

> **In short:** Perplexity lets you compare products by asking direct questions, then shows you AI-generated summaries with source links — useful for research, though the free version has usage limits.

Before you start

- An internet-connected device (computer, tablet, or phone).
- Access to the Perplexity website (perplexity.ai) or the Perplexity app.
- Perplexity offers both free and paid tiers — the free tier has daily search limits, while paid plans (Pro, Enterprise) offer higher usage. This guide works with either, but you may hit limits on the free version.
- The whole process should take about 5-7 minutes.

Step 1 — Open Perplexity

Start by opening Perplexity in your web browser or launching the app on your device. When you first arrive, you'll see a clean interface designed for asking questions. The main feature is usually a large text box in the center, often labeled "Ask anything..." or "Search here." This is your gateway to getting information. If you're on a mobile device, the layout might be slightly more compact, but the central search bar will still be prominent. You'll know it worked when you see the main Perplexity screen with a clear input box, ready for your query.

Step 2 — Enter your comparison query

Now it's time to tell Perplexity what you want to compare. Click or tap inside the main search box and type a clear question asking for a comparison. Think about what specific products you're interested in and what aspects you want to know about. After typing, you might see a "Search" or "Ask" button appear, or you can simply press the 'Enter' key on your keyboard or the 'Go' button on your mobile device's keyboard. If it looks different, some versions might have a microphone icon if you prefer to speak your query. You'll know it worked when the AI starts processing your request, often showing a spinning icon or a "Generating answer" message. 💬 Example: `Compare the Google Pixel 8 vs Samsung Galaxy S24 on camera quality and battery life.`

Step 3 — Review the AI's summary

Once Perplexity has finished processing your request, it will present a summary of the comparison. This usually appears as a concise, easy-to-read answer, often using bullet points, short paragraphs, or even a simple table to highlight key differences and similarities between the products. Below the summary, you'll typically see numbered source links – these are the websites and articles Perplexity used to gather its information. If it looks different, the exact layout can vary slightly, but always look for a clear answer section at the top of the page. You'll know it worked when you see a well-organized summary directly comparing the products you asked about.

Step 4 — Explore the sources

To ensure the information is reliable and to dig deeper into specific details, it's a good idea to check Perplexity's sources. Each piece of information in the summary often has a small, numbered link next to it. Click on these numbers, or find the "Sources" section usually located below the main answer. What happens next is that clicking a source link will open the original webpage in a new tab in your browser. If it looks different, sometimes the sources are just listed with their titles, and you click the title rather than a number. You'll know it worked when a new browser tab opens to an external website, letting you see where the AI got its information from, like a review site or a manufacturer's product page.

Step 5 — Ask follow-up questions

Perplexity isn't a one-shot tool; you can continue the conversation to get more detailed insights. Below the initial answer, you may find a section with suggested "Related Questions" or a new input box where you can type your own follow-up. This allows you to ask for clarification, or to explore other aspects of the products you're comparing. For example, if the initial summary didn't cover specific features you care about, now's your chance to ask. If it looks different, you might need to scroll down to find a new prompt area or simply type directly into the main chat box again. You'll know it worked when Perplexity provides a fresh answer, refining its previous information based on your new question. 💬 Example: `Which one is generally considered better for low-light photography according to user reviews?`

Step 6 — Refine your query for specific details

Sometimes, you need to go beyond a general comparison and focus on very specific features or scenarios. You can start a completely new query (by clicking "New Thread" or clearing the search bar if available) or continue in your current conversation. For this step, try asking Perplexity a

highly specific question about a particular product feature or how it performs in certain conditions. This is where Perplexity shines for deep dives, pulling targeted information from various sources. If it looks different, some versions might automatically clear the search bar for a new query, while others require you to click a "Start a new query" button. You'll know it worked when Perplexity gives you a highly focused answer, directly addressing your specific detail. 💬 Example: `What are the pros and cons of an OLED TV screen for watching sports, compared to a QLED screen?`

Common mistakes

- **Being too vague with your query:** If you just type "compare phones," Perplexity won't know what you're looking for.
  - **The fix:** Be specific! Name the products and the aspects you care about (e.g., "Compare iPhone 15 vs Samsung Galaxy S24 for camera and battery").
- **Not checking the sources:** Perplexity is excellent, but like any AI, it draws from various sources, and some might be outdated or biased.
  - **The fix:** Always glance at the numbered sources. If you see a review site you trust, or the manufacturer's official page, that's a good sign. If something seems off, check another source.
- **Expecting live pricing or stock information:** Perplexity excels at product specifications, features, and reviews, but it's not a live shopping tool.
  - **The fix:** Use Perplexity for product research, then check specific retailer websites for the latest prices, deals, and stock availability.
- **Hitting free tier limits:** If you're on the free version, you may reach daily search limits quickly.
  - **The fix:** If you need unlimited comparisons, consider upgrading to a paid plan. Otherwise, space out your queries or focus on the most important comparisons.

What this means for you

- **In everyday life:** Imagine you're buying a new fridge or a pair of running shoes. Instead of trawling through countless review sites, you can ask Perplexity to summarize the key differences, pros, and cons of your top choices. This saves you time and helps you feel more confident about your purchase.
- **For your business or work:** If you're deciding between two different project management software tools, or researching the best laptop for your team, Perplexity can provide a quick breakdown of features, pricing tiers (if publicly available), and user feedback, helping you make an informed decision for your budget and needs.
- **If you're just getting started:** Try comparing two items you know well, like different brands of coffee machines or popular car models, focusing on just one or two aspects you care about most.

Try it now

Open Perplexity and ask it to compare two household items you might consider buying, like "Compare features of the Dyson V15 Detect vs Samsung Bespoke Jet for pet hair removal." See how quickly you get a concise answer!

How to Compare Phone Plans and Pick the Right One Without the Jargon

6 steps · Read online

> **In short:** Pick a plan by looking at your real monthly data use, the all-in cost (not just the advertised price), and the coverage in your area — most people overpay because they pick by brand instead of fit.

Picking a phone plan can feel like decoding a foreign language. Carrier websites throw around words like "unlimited," "deprioritization," "throttling," and "MVNO" without explaining any of them. This guide walks you through a simple way to compare plans on your own terms: what you actually use, what you actually pay, and what you actually need.

Before you start

- Your last 2-3 phone bills (paper, email, or in your carrier's app)
- About 30-45 minutes of quiet time
- A notebook, spreadsheet, or a notes app
- A friend or family member on a different plan (optional, for coverage feedback)
- Roughly 7 days to think it over — no need to decide in one sitting

Step 1 — Look at your real monthly usage

Open your last 2-3 phone bills — paper, email, or the carrier's app. Find the "data used" line. Data means the internet your phone uses over the cell network, not your home Wi-Fi. If your bill just says "unlimited," check your phone's settings: most phones show usage under **Settings > Network & internet > Data usage**, and you'll see a number like "8.2 GB used this month." If you can't find it, menus change between phone brands — look for "Cellular" or "Mobile data" instead. No number at all? A rough rule: light users (email, maps, some video) use 5-10 GB; heavy users (lots of streaming, video calls) use 20+ GB. Example: "I used 12 GB last month — mostly maps, Spotify, and one Zoom call a day." You'll know it worked when you have one number written down that you trust.

Step 2 — Decide which features actually matter to you

Make a short list of must-haves and nice-to-haves. Must-haves are things you'd cancel the plan over if missing. Common features to consider: **hotspot** (using your phone as a Wi-Fi source for a laptop — some plans limit this to 5-10 GB even on "unlimited" tiers), international calling or roaming, talk and text (most plans include unlimited, but double-check), family or multi-line discounts, and 5G access. Don't trust the word "unlimited" without reading the fine print — many "unlimited" plans slow your speeds after 20-50 GB, a practice called **throttling**, or push you behind other users in crowded areas, called **deprioritization**. Example list: "Must-haves: at least 15 GB, hotspot, 5G. Nice-to-haves: international roaming." You'll know it worked when you have a clear top-3 list of features you actually care about.

Step 3 — Learn the four main plan types

There are roughly four kinds of phone plans, and knowing which one you're looking at saves real money. **Postpaid** = you pay at the end of the month, usually through a big carrier, often with a phone installment plan attached. **Prepaid** = you pay upfront for 30 days, no contract. **MVNO (Mobile Virtual Network Operator)** = smaller brands (think Mint, Consumer Cellular, Cricket) that rent coverage from the big carriers (Verizon, AT&T, T-Mobile) and resell it cheaper. **Family plan** = multiple lines on one account, usually with a per-line discount. The catch: MVNOs are usually much cheaper but may run slightly slower in busy areas because the host carrier prioritizes its own customers first. Example: "I'll compare a Verizon postpaid plan, an AT&T-based MVNO, and a T-Mobile prepaid plan." You'll know it worked when you can tell which type a plan is just by reading its website.

Step 4 — Calculate the real total monthly cost

The advertised price is almost never what you actually pay. Add up: the plan's base price, taxes (often 10-20% extra, varies by region), any **device installment** (paying off a phone in monthly chunks — say $25 for a phone on a 24-month plan), and any "regulatory" or "administrative" fees. Then check for auto-pay or paperless billing discounts — many carriers knock off a small amount if you let them charge your card automatically. Also watch for **introductory pricing** that jumps up after 12 or 24 months. The real question isn't "what does it cost now" but "what will it cost in year two?" A $30 plan that becomes $70 after a year is the more expensive plan. Example: "$40 plan + $6 taxes + $25 phone = $71/month total. After promo ends: $95/month." You'll know it worked when you can see — in one number — what this plan will cost you a year from now.

Step 5 — Check coverage where you actually live and travel

A cheap plan is worthless if it doesn't work in your apartment, your office, or your parents' house. Coverage maps on carrier websites are optimistic — they're the best case. For a real check, ask friends and neighbors on different carriers how their service is at your home address. If you commute, ask about the route. If you travel, look for the carrier's international options or roaming add-ons. Some MVNOs have slightly weaker indoor signal in rural areas, even though they use the same towers as the host carrier. If you live in a city, almost any plan will work fine. If you live in a small town or rural area, this is the most important step in the whole guide. Example: "I asked my neighbor: 'How's your T-Mobile

signal in our building?' He said: 'Two bars, calls drop.'" You'll know it worked when you can name two people whose coverage matches your daily life.

Step 6 — Compare 2-3 plans side by side

Pick your top three plans from the previous steps. Make a simple table — paper, notes app, or spreadsheet — with these columns: plan name, monthly cost in year two (including everything), data allowance, your must-have features (yes/no), and a "gut feeling" line. The gut-feeling column matters: a plan that ticks every box but has awful customer service is a plan you'll hate in six months. Read recent reviews, but focus on patterns, not single angry comments. Once your table is filled in, the right plan usually jumps out: it hits your must-haves, stays affordable in year two, and has good coverage where you live. If two plans tie, pick the one with the friendlier cancellation policy — flexibility is worth real money. 💬 Example row: "Plan A: $55/mo, 25 GB, hotspot yes, support fair. Plan B: $45/mo, 15 GB, no hotspot, support good. Pick A — hotspot is a must-have." You'll know it worked when you can point to one plan and say, "That's the one — and here's why."

Common mistakes

- **Picking by brand loyalty instead of fit.** The cheapest "unlimited" plan is rarely the cheapest plan for you specifically.
- **Forgetting the year-two price.** Promotional rates that double after 12-24 months are the most common cause of bill shock.
- **Trusting "unlimited" without reading the fine print.** Most "unlimited" plans slow down your speeds or deprioritize your traffic after a set amount of data.

What this means for you

- **In everyday life:** a $20-40 monthly difference over two years is $500-1,000 — real money for a vacation, an emergency, or a small home repair.
- **For your work or business:** if your job depends on reliable calls or hotspot for a laptop, the cheapest plan is rarely the right one — downtime costs more than the savings.
- **If you're just getting started:** pick one plan from a major carrier or well-known MVNO with no contract, then reassess in 3 months once you know your real usage. No decision has to be permanent.

Try it now

Pull up your last phone bill right now. Find the data usage number. Write it down somewhere you can find again. That's the single most important number in choosing a plan — and you just got it in under a minute.

Use Copilot to explain code you found online in plain English

6 steps · Read online

> **In short:** Paste any code snippet you don't understand into Copilot Chat, ask it to explain each line in plain English, and follow up with safety questions before you run the code. You need a GitHub account and a Copilot plan (free tier, paid, or a free education/open-source seat). Copilot can describe what code is *meant* to do, but it can't actually run it for you.

Found a code snippet in a tutorial and have no idea what it does? Copilot Chat can translate it into plain English, line by line, so you can decide whether to actually run it. This guide walks you through the exact questions to ask, in the order that works best for a complete beginner.

Before you start

- A **GitHub account** (free to create at github.com — just an email and a password).
- A **Copilot plan**. There's a free tier with a limited number of chats, paid plans for heavier use, and free seats for verified students and open-source maintainers. If you're unsure, sign in and look for the "Try Copilot" prompt — it will tell you which plan you're on.
- A **web browser** (Chrome, Firefox, Edge, or Safari). If you prefer, you can do the same thing in VS Code with the Copilot extension installed, but this guide sticks to the web.
- A **code snippet to test with** — keep a tutorial, blog, or Stack Overflow answer open in another tab.
- About **5–10 minutes**.

Step 1 — Open Copilot Chat in your browser

Head to github.com and sign in. Once you're on the home page, look for a **Copilot icon** in the top navigation bar — it usually looks like a small spark or chat bubble. Click it, and a chat panel slides open from the side of the page. If you don't see the icon, try the **profile menu** (your avatar) in the top-right and look for a "Copilot" or "Try Copilot" link. On a paid plan you get generous chat limits; on the free tier you may get a small number of chats per month. Example: (no typing yet — just opening the chat) You'll know it worked when a chat box appears at the bottom of the panel with a prompt like "Ask Copilot" or "Type a message."

Step 2 — Paste the code snippet into the chat

Go back to the page where you found the code. Select the snippet with your mouse, copy it (Ctrl+C on Windows, Cmd+C on Mac), then click into the Copilot chat box and paste it (Ctrl+V / Cmd+V). The text appears as a **code block** — a grey box with monospaced font — so it's easy to read. If the snippet is very long, paste only the first 10–20 lines first; you can always add more later. Don't paste any secrets like API keys, passwords, or private tokens — replace those with placeholders such as `"YOUR_API_KEY"` first. Example: a real-looking snippet you might paste: `import requests; r = requests.get("https://api.example.com"); print(r.json())` You'll know it worked when the code shows up formatted in the chat, with the assistant waiting below for your next message.

Step 3 — Ask for a plain-English explanation

Now type a clear instruction in the chat box. The trick is to say exactly what you want and how you want it explained — a line-by-line breakdown is usually the most useful starting point. After you press Enter, Copilot writes back a multi-paragraph explanation right in the same chat. If instead you get more code, rephrase your prompt to start with the word "Explain..." rather than "What is this..." Example: "Explain this code line by line, like I'm 12 years old and have never programmed before." You'll know it worked when Copilot describes what each line does in everyday words, not technical jargon — and ideally gives a one-sentence summary at the top.

Step 4 — Drill into the parts you still don't get

Read Copilot's answer and pick the one sentence that's still confusing. Type a new question that points at exactly that part — the more specific you are, the better the answer. If the explanation still uses words you don't know, ask for an analogy. Each follow-up sits in the same chat, so you can scroll up to see the full conversation and Copilot remembers what you already asked. Example: "Can you explain `requests.get(...)` like I'm ordering food at a restaurant?" You'll know it worked when you can summarize the snippet in your own words to a friend without looking back at the chat.

Step 5 — Ask about safety before you run it

Now the important question: should you actually use this code? Type a new prompt asking Copilot to flag risks. It can't actually run the code for you, but it can spot common things to check — hardcoded URLs, missing error handling, dependencies you might not have installed, or commands that delete files. Treat its answer as a starting checklist, not a guarantee; you still need to test in a safe place like an online playground or a throwaway folder. Example: "Are there any security risks, missing dependencies, or things that could break if I run this on my computer?" You'll know it worked when Copilot names at least one concrete thing to check, like a library you'd need to install or a permission it requires.

Step 6 — Save the explanation for later

Hover over Copilot's answer and look for a small **copy icon** (two overlapping squares) at the top-right of the message. Click it to copy the whole reply, then paste it into a note app, a Google Doc, or an email to yourself. If there's no copy button, just select the text with your mouse and copy it the normal way. Saving now means you can come back to the explanation later without re-asking the same questions — and you have a record of what the snippet does in case you forget. 💬 Example: (no typing — just clicking the copy icon and pasting into a note) You'll know it worked when you can open your notes app and read the explanation back, with the original code snippet included.

Common mistakes

- **Asking without context.** If you just paste code and type "what is this?", Copilot guesses at your skill level and may give an answer full of jargon. The fix: tell it up front that you're a beginner — the example phrase in Step 3 works well.
- **Trusting the explanation as proof of safety.** Copilot can confidently misread code, and it can't actually test it for you. The fix: always run unfamiliar snippets in a sandbox, a throwaway folder, or a free online playground like replit.com before using them on real files.
- **Pasting private information.** API keys, passwords, customer names, or your company's internal code shouldn't go into any AI chat. The fix: replace sensitive values with placeholders like `"YOUR_API_KEY"` or `"<redacted>"` before you copy the snippet.

What this means for you

- **In everyday life:** You'll be able to follow programming tutorials you find online without getting stuck at the third line. That opens up small automations — renaming a folder of photos, formatting a spreadsheet, pulling data from a public webpage.
- **For your business or work:** When a colleague drops a "quick script" into a chat channel, you can read it yourself instead of just asking "is this safe?" — and you can reply with a short summary, not a shrug.
- **If you're just getting started:** Find a beginner Python tutorial, copy the very first code block, and paste it into Copilot Chat. The first time you read a plain-English explanation of code, the whole programming world becomes a little less mysterious.

Try it now

Open github.com, find any beginner Python tutorial (the official Python wiki has a "Hello, World" example), copy the first line of code, and ask Copilot: "What does this do, in one sentence?" That's it — one snippet, one sentence, one answer. You've just made Copilot your tutor.

Summarize a long chat thread into 3 bullet points with Copilot

6 steps · Read online

Ever feel lost trying to find the key takeaways from a really long email chain or group chat? This guide is for anyone who wants to quickly cut through the noise and get straight to the important decisions and actions. By the end, you'll be able to use Microsoft Copilot to summarize any long text into three clear bullet points, making it easy to catch up or share concise updates.

Before you start

- A free Microsoft account (you can sign up for one easily if you don't have it).
- A web browser like Microsoft Edge, Google Chrome, or Mozilla Firefox.
- A long piece of text you want to summarize, such as an email conversation or a chat history.
- Roughly 5-7 minutes to follow the steps.

Step 1 — Open Copilot in your browser

Your first move is to open Copilot in your preferred web browser. Simply type `copilot.microsoft.com` into your browser's address bar at the top of the window, then press `Enter` on your keyboard. The Copilot homepage should load, presenting you with a clear chat input box usually located at the bottom of the screen, waiting for your commands. If the page looks a little different, don't worry – websites update often. Just look for a prominent chat area or a button that says "Start chat" or "Ask me anything." You'll know it worked when you see a clear input box at the bottom of the screen, probably saying "Ask me anything" or similar.

Step 2 — Sign in to your Microsoft account

To get the most out of Copilot, it's best to sign in. Look for a "Sign in" button, typically positioned in the top-right corner of the Copilot page, and click it. A new screen will appear, prompting you to enter your Microsoft account details, like your email address or phone number, and then your password. If you're already signed in, you'll see your profile picture or initials in that corner and can simply skip this step. If it offers choices like "work or school account," usually you'll select "personal account" for this task. You'll know it worked when your profile picture or initials appear in the top-right corner, indicating you are successfully logged in.

Step 3 — Craft your summary request

Now it's time to tell Copilot exactly what you want it to do. In the chat input box at the bottom of the screen, type out your specific instruction. Be clear that you want a summary, how many bullet points you need, and what you want it to focus on – in this case, decisions and action items. This "instruction" is often called a **prompt** (which is just the message you send to an AI). Your typed message will appear in the input box, ready for you to add the content you want summarized. 💬 Example: `Summarize the following chat thread into exactly three short bullet points, focusing on key decisions and action items:` You'll know it worked when your complete instruction is clearly visible in the chat box, ready for the content you're about to paste.

Step 4 — Paste your long text

Next, you need to provide Copilot with the text to summarize. Go to the original source of your long chat thread or email chain (like your email client or messaging app), select all the text you want, and copy it. Then, click directly into the Copilot chat input box, making sure your cursor is positioned right after the instruction you typed in Step 3. Now, paste the copied text. The entire block of text will appear in the input box after your instruction. Copilot has a "context window" (the amount of text it can "remember" and process in one conversation), but for most chat threads, it will fit easily. You'll know it worked when all your copied text is visible in the input box, ready to be sent.

Step 5 — Send your request

With your instruction and the long text in place, it's time to send it to Copilot. Press the `Enter` key on your keyboard, or click the "Send" button that looks like a paper plane icon, usually located next to the input box. Once sent, your entire message will move from the input box into the main chat history. Copilot will then show a "thinking" indicator, meaning it's processing your request. After a short moment, it will generate its response, providing your summary. If the "Send" icon looks a little different, just know it's always near the input box. You'll know it worked when Copilot starts typing its response, and you see the summary appear as three bullet points in the chat window.

Step 6 — Review and refine your summary

Read through the three bullet points Copilot has generated. Does it capture the essence of the conversation? Does it highlight the key decisions and actions as you intended? If it's not quite perfect, you can ask Copilot to refine it. Simply type a follow-up request in the input box, explaining what changes you'd like. Copilot will then generate a revised summary based on your new feedback. This helps you get precisely what you need from the conversation. 💬 Example: `Please make the bullet points even shorter, just focusing on who needs to do what.` You'll know it worked when you have a concise, useful three-point summary of your original long text that meets your needs.

Common mistakes

- **Pasting text without a clear instruction:** If you just paste a long thread into Copilot without telling it what to do, it might simply repeat parts of the text or give a generic response. Always start with a specific instruction, like "Summarize this into three bullet points."
- **Assuming the AI knows context:** Copilot only "knows" what you provide in the current conversation. If your chat thread refers to external documents or past discussions, Copilot won't have that information unless you paste it in. Make sure all necessary information is included for the best summary.
- **Too much text for one go:** While Copilot has a large capacity, extremely long documents (tens of thousands of words) might exceed its "context window" (the amount of text it can process at once). If Copilot seems to miss the end of your content, try breaking the original text into two or three smaller chunks and summarizing them separately.

What this means for you

- **In everyday life:** Imagine catching up on a flurry of family messages about a holiday plan or quickly understanding the core points of a long community group discussion. Copilot helps you get the gist without having to read every single message, saving you precious time.
- **For your business or work:** Instantly extract the crucial decisions, assigned tasks, and next steps from lengthy meeting transcripts, project update emails, or internal team chats. This ensures you're always informed and can provide concise updates to colleagues without sifting through pages of text.
- **If you're just getting started:** A great way to practice is to take a long news article or a detailed product review you're interested in, copy the text, and ask Copilot to summarize it into three bullet points. This helps you see how it condenses information.

Try it now

Find a recent long email or a lengthy article online that you've been meaning to read. Copy the main body of text. Then, open Copilot, paste your copied text, and simply add the instruction: "Summarize this into three clear bullet points." See how quickly you can get the core information!

Read the Room on X: Summarize Public Opinion with Grok

6 steps · Read online

By the end of this guide, you'll be able to open Grok, ask it to summarize what people on X are saying about a trending topic, and walk away with the main viewpoints, common arguments, and overall mood in just a few minutes. It's for anyone who wants to understand public reaction to a news story without scrolling through hundreds of posts by hand.

Before you start

- **An X account** (the one linked to your Grok access). Free X accounts can use Grok with daily limits; X Premium and Premium+ give you more messages and access to newer Grok models.
- **Access to Grok** — through the X app on iOS or Android, the X website, or **grok.com** in a browser. They're the same AI underneath.
- **A trending topic in mind** — a news event, a product launch, a sports moment, anything that's getting a lot of conversation right now.
- **About 5 minutes** to run through the steps the first time.

One honest note before we start: Grok reads what's publicly posted on X. It won't show you private accounts, deleted posts, or reactions from other platforms like TikTok or Reddit. Treat its summary as a starting point, not the final word.

Step 1 — Open Grok on the place you already use

Open the **X app** on your phone and tap the center button in the bottom navigation bar — it's labeled **Grok** and usually shows a small sparkle or wave icon. If you're on a computer, go to **grok.com** in your browser, or open X on the web and click the **Grok** option in the left-hand menu. A chat screen appears with a friendly greeting and a text box at the bottom.

If you don't see Grok in the bottom bar, your X app may need an update — open the App Store or Google Play and update the X app, then look again. On older versions, the button may be under the **More** menu.

You'll know it worked when you can see a chat-style screen with a blinking cursor in a message box, ready for you to type.

Step 2 — Turn on real-time access to X

In the chat box, look for a small **globe icon**, a **"Search" toggle**, or a label that says **"Use X posts"** or **"Real-time"**. Tap or switch it on so Grok reads fresh posts from X rather than only its general knowledge. On **grok.com** this is usually a toggle just above the text box. Inside the X app, it may appear as a small chip you can tap before sending.

If you can't find the toggle, don't worry — simply ask Grok in your message to use recent X posts and it will usually do so on its own.

You'll know it worked when the toggle is highlighted or colored in, and Grok's reply mentions specific posts or accounts from X.

Example toggle behavior: you turn on **Use X posts**, and the icon turns from gray to blue.

Step 3 — Ask Grok to summarize the conversation

Tap the message box and type a clear, plain-English question. Be specific about the topic, the time frame, and what you want to know. The clearer your instruction (prompt), the better the summary. Press the **arrow** or **send** button to send it.

Grok will think for a few seconds, then write back a structured answer. Because you've turned on X access, it'll pull from recent posts rather than guessing.

You'll know it worked when Grok replies with a written summary that mentions posts, accounts, or quotes from X — not a generic definition.

Example: *"Summarize what people on X are saying about the new AI safety rules announced this week. Give me the main viewpoints, the most common arguments for and against, and the overall mood."*

Step 4 — Ask for the three things you actually want

Most readers want three things from a summary: the **main viewpoints**, the **common arguments**, and the **overall mood**. If Grok's first reply doesn't cover all three, ask a follow-up in the same chat. Grok remembers the conversation, so you can refine instead of starting over.

Press **enter** or tap **send** again. A new reply appears underneath the previous one, building on the same topic.

You'll know it worked when you can see clearly labeled sections or bullet points for viewpoints, arguments, and mood — or a single paragraph that covers all three.

Example: *"Good. Now add the most common arguments people are making on each side, and tell me the overall mood — positive, negative, or mixed — and roughly what share of posts fall into each."*

Step 5 — Quote and verify before you share

Scroll through the reply and look for **quotation marks**, **post handles**, or **small links** next to claims. Tap any link to open the original X post in a new tab so you can read it yourself. If a number or quote feels surprising, ask Grok to back it up.

If a link doesn't open, or you can't see any links at all, ask Grok directly to show its sources. This is the most important habit: a summary is only useful if you can check it.

You'll know it worked when you can click at least one link and see the original post that Grok is referring to.

Example: *"Show me 3 real X posts that support the negative mood you described, with the usernames and links."*

Step 6 — Save the summary for later

When the reply looks good, look for a **copy icon** (two overlapping squares) or a **share icon** at the bottom of Grok's reply. Tap it to copy the text to your clipboard, then paste it into a note, email, or document. On grok.com you can also click the **download** or **bookmark** icon to save the whole conversation.

If you want to come back to it later, your chat history with Grok is available on the same screen — just open the **sidebar** or **history** list and tap the conversation title.

You'll know it worked when the text is in your clipboard, your notes app, or your saved conversations list, ready to re-read.

Common mistakes

- **Asking a fuzzy question.** "What's happening on X?" is too vague. Add the topic, the date, and what you want — Grok performs much better with specific prompts.
- **Forgetting to turn on X access.** Without it, Grok may answer from general knowledge and miss the live conversation. If the reply feels generic, go back to Step 2 and check the toggle.
- **Trusting the summary without checking.** Grok can misread sarcasm, mix up sides, or pick unrepresentative posts. Always click at least one source link before you repeat the claim to someone else.

What this means for you

- **In everyday life:** When a big story breaks — a sports result, a celebrity moment, a weather event — you can get a balanced feel for what people are actually saying in under five minutes, instead of doom-scrolling.
- **For your business or work:** If you're in marketing, comms, journalism, or customer experience, this is a quick way to do a morning scan of public sentiment around your brand, your industry, or a news event before a meeting.
- **If you're just getting started:** Pick one topic you already follow on X, run through the steps, and compare Grok's summary to what you saw yourself. That's the fastest way to learn how much to trust it.

Try it now

Open Grok, turn on X access, and ask: *"Summarize what people on X are saying about [one trending topic you already know about]. Give me the main viewpoints and the overall mood."* Then click one source link and see if the summary matches reality.

Ask Copilot to compare two product specs and list the differences

6 steps · Read online

You've got two product pages open and need to decide which one fits your needs. Instead of reading both from top to bottom, give Copilot the key specs and let it list the differences for you. In minutes, you'll have a clear side-by-side summary—no PhD in spreadsheets required.

Before you start

- A **GitHub account** (free is fine)
- A **computer or phone** with a web browser
- Two product pages open in separate tabs (for example, two software pricing pages or two laptop spec sheets)
- Rough time needed: **5 minutes**

Step 1 — Open Copilot in your browser

Go to github.com/github-copilot and sign in with your GitHub account. Alternatively, you can find Copilot through the GitHub homepage by clicking the Copilot option in the navigation menu. You'll see a chat box ready for your question.

If you already use Copilot in another way (like in your IDE or on Windows), you can also access the web-based chat from the main GitHub interface.

Example: Type or paste: "Hello"

You'll know it worked when Copilot replies with something like "Hi there! How can I help you today?"

Step 2 — Share the first product's specs

Since Copilot can't automatically read external web pages, you'll need to copy the key specs from the product page yourself. Go to your first product page, select and copy the main specifications (like processor, RAM, storage, screen size, weight, and price), then paste them into the chat box.

Example: Paste: "Here are the specs for the first laptop: Processor: Intel Core i7-13700H, RAM: 16GB LPDDR5, Storage: 1TB SSD, Screen: 14" 2.8K OLED 120Hz, Weight: 1.4kg, Price: $2,199"

You'll know it worked when Copilot acknowledges the specs and confirms it's ready to compare.

Step 3 — Ask Copilot to format the specs

Now, tell Copilot how you want the specs organized. Type: "Format these specs as a clean bullet-point list."

Copilot will return a neat list of the specs you shared, making them easy to compare with the second product.

Example: Type: "Format these specs as a clean bullet-point list."

You'll know it worked when you see a clear list like: • Processor: Intel Core i7-13700H • RAM: 16GB LPDDR5 • Storage: 1TB SSD • Screen: 14" 2.8K OLED 120Hz • Weight: 1.4kg • Price: $2,199

Step 4 — Repeat for the second product

Go back to your browser and copy the key specs from the second product page. Return to Copilot and paste those specs into the chat.

Example: Paste: "Now for the second laptop: Processor: AMD Ryzen 9 7940HS, RAM: 32GB LPDDR5, Storage: 512GB SSD, Screen: 16" 2.5K IPS 165Hz, Weight: 1.8kg, Price: $1,799"

You'll know it worked when Copilot acknowledges the second set of specs.

Step 5 — Ask Copilot to compare the two lists

Now ask Copilot to compare the two sets of specs. Type: "Compare these two products and list the differences in a clear table or bullet list."

Copilot will take the two lists and highlight what's different—like RAM upgrades, storage options, screen refresh rates, or price changes.

Example: Type: "Compare these two products and list the differences in a clear table or bullet list."

You'll know it worked when you see a concise comparison like: • RAM: 16GB vs 32GB (second product has more) • Storage: 1TB vs 512GB (first product has more) • Price: $2,199 vs $1,799 (second product is cheaper) • Screen: 14" vs 16" (second product has a larger screen)

Step 6 — Ask for a recommendation (optional)

If you're still unsure, ask Copilot to give a simple recommendation based on your needs. Type: "Which of these two products is better for video editing on a budget?"

Copilot will weigh the specs against your stated use case and suggest the better option.

Example: Type: "Which of these two products is better for video editing on a budget?"

You'll know it worked when Copilot gives a concise recommendation with reasons.

Common mistakes

- **Pasting the wrong specs**: Double-check what you're copying from the product page. If you include the wrong specs, just paste the correct information and ask Copilot to start over.
- **Not including enough detail**: If you only paste partial specs, Copilot can't give you a full comparison. Include the key specs: processor, RAM, storage, screen, weight, and price.
- **Assuming Copilot remembers**: Copilot doesn't keep memory between chats. If you open a new chat, you'll need to paste both sets of specs again. Bookmark your chat or keep it open to avoid redoing steps.

What this means for you

- **In everyday life**: Next time you're choosing between two gadgets, appliances, or software tools, use Copilot to do the heavy lifting. Copy the specs, paste them in, and get a clear comparison in seconds—no more squinting at tiny font on a phone screen.
- **For your business or work**: When your team debates which tool to adopt, run this comparison in a shared chat. Everyone sees the same facts, and you avoid long email threads full of conflicting opinions.
- **If you're just getting started**: Start with two simple product pages you already have open. Try the steps above, and you'll see how easy it is to turn messy specs into a clean comparison.

Try it now

Open two product pages in your browser right now. Copy the first set of specs, go to github.com/github-copilot, paste them in, and ask Copilot to format them. Do the same for the second page, then ask it to compare them. You'll have your answer in under five minutes—no spreadsheets, no PhD required.

Summarize Long Research PDFs with Cohere’s Clever AI

5 steps · Read online

Imagine turning a dense, fifty-page research paper or a lengthy report into a crisp, one-page summary in under ten seconds. Whether you're a student tackling academic journals or a professional keeping up with industry reports, this guide will show you how to use Cohere to extract key insights without the headache, using a simple online tool.

Before you start

- **A Cohere account:** You'll need a free developer account to access the tools. Sign up on Cohere’s website.
- **A research PDF:** Have a PDF document ready on your computer that contains selectable text (not scanned images).
- **A web browser and computer:** You'll use the Cohere website on a desktop or laptop, as the interface might look different on a mobile device.
- **Time needed:** This guide should take about 10-15 minutes to follow.

Step 1 — Log in and find your API Key

Before we can ask the AI to read our document, we need a digital passkey to access its brain. This passkey is called an **API Key (Application Programming Interface Key)**, which is a unique password that securely connects your computer or web browser directly to the AI engine.

Once you log into your new Cohere dashboard, you'll usually see a menu on the left side of the screen. Look for a tab or section labeled "API Keys" or "Developers." Click on this to reveal your keys. Your active "Trial" API key will be displayed there; make sure to copy it. Keep this key private, just like you would with a bank password, to protect your account. You'll know it worked when you see a long string of letters and numbers (your unique key) next to "Trial API Key" or similar wording on the screen.

Example: Your key will look like a long string of random letters and numbers, such as `a1B2c3D4e5...`

Step 2 — Open the Cohere Playground

Writing code can be intimidating if you're just starting out, but Cohere provides a **Playground** (an online sandbox or testing space) where you can experiment with AI settings using simple buttons instead of writing code.

From your Cohere dashboard, look for a prominent "Playground" button, which is often found at the top or in the sidebar. Clicking this will take you to a page where you can choose different AI "templates." Select the "Summarize" option from the list of templates. After selecting, a new interface will load, typically showing a large text box on the left for your input and various setting sliders on the right. If you don't immediately see "Summarize," look for a section like "Templates" or "Use Cases." You'll know it worked when you see a main text area on the left and settings options on the right, with "Summarize" active or prominently displayed.

Step 3 — Extract and paste your PDF text

Currently, most AI tools work best when you feed them raw text directly. We need to copy the text out of our PDF and give it to the tool. Remember, the Cohere Summarize tool processes text you provide; it can't directly *open* your PDF file itself.

Open your research PDF on your computer using any PDF reader. Carefully highlight the main sections of the paper you want summarized – the introduction, key findings, and conclusion are usually the most important for a comprehensive overview – and copy them (using Ctrl+C on Windows or Command+C on Mac). Switch back to your Cohere Playground window. You'll see a large input box on the left, usually labeled "Input" or "Text to Summarize." Paste the text you copied from your PDF directly into this box. Don't worry if the formatting looks a bit messy after pasting; the AI is smart enough to read past odd spacing and line breaks. You'll know it worked when your text appears in the input box, ready for the AI to process.

Example: Paste the text you copied from your PDF into the large input box on the left side of the Playground.

Step 4 — Customize your summary settings

On the right side of the screen, you'll see a panel of options. These sliders and dropdowns let you tell the AI exactly how you want your information presented, giving you control over the summary's output.

You'll typically find settings for "Length," "Format," and "Extractiveness."

1. **Length:** Choose between "Short," "Medium," or "Long." For a quick overview of a dense paper, "Short" is usually best.
2. **Format:** Select "Bullets" (perfect for quick scanning of key points) or "Paragraph" (better for a narrative, flowing summary).
3. **Extractiveness:** This slider controls how much the summary sticks to the exact wording of your original text versus rephrasing it in simpler, more general terms. Set it to "Low" if you want a plain-English translation that simplifies academic jargon, or "High" if you need the exact academic terms and phrases preserved.

If these options aren't immediately visible, look for a "Settings" or "Options" panel on the right side of the screen that you might need to expand. You'll know it worked when your chosen settings (e.g., "Short," "Bullets," "Low") are displayed next to their respective labels.

Step 5 — Generate your summary

Now it's time to let the AI do the heavy lifting and transform your long text into a concise summary.

Once you've adjusted your settings, scroll to the bottom of the right-hand panel or the Playground interface. Look for a button labeled "Generate" or "Run." Click this button to begin the AI summarisation process. Within a few seconds, a polished, easy-to-read summary will appear in a separate output window or section on the right side of the screen. If the "Generate" button is greyed out, double-check that you have pasted some text into the input area. You'll know it worked when a new section on the right-hand side of the screen populates with your AI-generated summary, which you can then copy and save for your study notes or work files.

Example output:

> **Key Findings:**
>
> - The study tested 500 participants over twelve months.
> - Using AI tools reduced administrative tasks by 40%.
> - Participants reported higher satisfaction when using clear, visual interfaces.

Common mistakes

- **Pasting scanned images instead of text:** If your PDF is a scanned photocopy of an old book, you won't be able to highlight and copy the text.
  - **The fix:** Ensure you use "selectable" PDFs, where you can actually highlight individual words. If you only have a scanned image, you'll need to run the file through an Optical Character Recognition (OCR) tool first to convert the images into real text before copying.
- **Hitting the text limit with very long documents:** While Cohere can handle substantial text, free accounts often have limits on the total number of words (or "tokens") you can process in one go. If you paste an entire book, it might fail or only summarize a portion.
  - **The fix:** For extremely long documents, try breaking them into key sections (e.g., summarize chapter by chapter) or focus on only the most critical parts like the introduction, methods, results, and conclusion.
- **Leaving your API key visible:** Your API key is like a personal password; sharing it publicly could allow others to use your Cohere account, potentially incurring costs or misusing the service.
  - **The fix:** Never share screenshots that show your API key, and never paste it into public forums or messages. If you accidentally share it, simply click "Regenerate Key" in your Cohere dashboard to disable the old one and get a new, secure key.

What this means for you

- **In everyday life:** You can use this method to quickly understand long product manuals, complicated rental agreements, or lengthy council planning reports before making big decisions, saving you hours of reading.
- **For your business or work:** You can easily digest competitor research, market studies, or complex industry regulations without spending hours reading every page, helping you stay informed and make quicker decisions.
- **If you're just getting started:** Try starting with a short news article (about 500-1000 words). Copying and pasting a single page of text into the Playground is the gentlest way to see how the AI structures its thoughts before tackling bigger documents.

Try it now

Open a browser tab, sign up for a free Cohere account, and copy just three paragraphs from any online news article. Paste it into the Cohere Summarize Playground, hit **Generate**, and watch how quickly it extracts the main point for you!

Automate the boring bits

Create a tiny AI weather-helper widget for your blog with Hugging Face widgets

7 steps · Read online

Ever wanted to add a touch of AI magic to your personal blog or website without needing to write any code? This guide will show you how to find a simple AI weather tool on Hugging Face and embed it as a "widget" – a small, interactive gadget – directly onto your site. By the end, you'll have a custom AI weather helper for your readers.

Before you start

- **A free Hugging Face account:** You'll create this in the first step.
- **A web browser:** Like Chrome, Firefox, Safari, or Edge.
- **A blog or website:** Where you have permission to edit the raw HTML code (for example, a WordPress blog, a personal site built with simple HTML, or a similar platform).
- **Time:** This will take around 15-20 minutes.

Step 1 — Sign up for a free Hugging Face account

First, let's create your free account on Hugging Face. This platform is like a big online library and playground where people share AI models and interactive demonstrations. Open your web browser and go to the Hugging Face website. You'll see a homepage often featuring popular models or news. Look for a "Sign Up" or "Join" button, usually located in the top-right corner of the screen. Click on it to start the registration process. You'll then be prompted to enter your email address and choose a password, or you might have options to sign up using your Google or GitHub account. Follow the on-screen instructions to confirm your email address if required. **If it looks different:** If you don't immediately see a "Sign Up" button, look for a "Log In" button first, as signing up is usually an option from there. You'll know it worked when you are logged into your new Hugging Face account, and you see your chosen username or profile icon in the top-right corner. 💬 Example: `[Enter your Email Address]`, `[Choose a Password]`, `[Click Sign Up]`

Step 2 — Discover a weather model

Now that you have an account, let's find an AI weather model to turn into a widget. Hugging Face is home to "Spaces," which are interactive, web-based demos of AI models. On the Hugging Face website, look for a navigation menu at the top or side of the page, and click on "Spaces." This will take you to a page listing many different AI demos. Once on the Spaces page, find the search bar. This is where you'll tell Hugging Face what kind of AI tool you're looking for. Type "weather" into the search bar and press Enter or click the search icon. This will filter the vast collection of Spaces to show you only the ones related to weather. Browse through the results and choose one that looks simple and functional, usually with a clear input field for a city or location. **If it looks different:** If the "Spaces" link isn't immediately visible, try looking for a "Models" link first, as Spaces are often linked from there, or a general "Explore" section. You'll know it worked when you see a list of AI Spaces related to "weather" appear on your screen. 💬 Example: `weather` (typed into the search bar)

Step 3 — Interact with the model

Before embedding the widget, it's a good idea to test the chosen weather model to make sure it works as expected. This step helps confirm that the AI model (the computer program trained to perform a specific task, like predicting weather) is functional and useful for your blog. Click on one of the weather Spaces you found in the previous step. You'll be taken to its dedicated page, where you'll usually see the "User Interface (UI)" – the visual layout with buttons, text fields, and displays – of the demo. Look for an input box, often labeled "City" or "Location." Type a city name into this box and click the "Submit" or "Generate" button, if there is one. The AI model will then process your request and display the weather information. Make sure the output makes sense, like current temperature or conditions. **If it looks different:** Some Spaces might have slightly different layouts, or they might run slowly on the free tier if they're very popular. Try a different weather Space if your first choice doesn't respond or seems too complex. You'll know it worked when the model successfully processes your input and displays weather information for the city you typed in. 💬 Example: `Chicago` (typed into the city input box)

Step 4 — Generate the embed code

Now that you've found a working weather model, let's get the special piece of computer code (called "HTML") that allows you to put content from one website onto another. This is known as the embed code. On the specific Hugging Face Space page for your chosen weather model, look for a button or icon that suggests embedding or sharing. This is often a `</>` icon, a "Share" button, or sometimes an "Embed" option within a menu — if it's available, it may be located near the top-right of the demo frame or page. Click this option. A small window or panel should appear, providing you with different sharing options, one of which may be "Embed" or "Embed Code." **If it looks different:** If you can't find a `</>` icon, look for a three-dot menu `...` or a "Settings" gear icon, as embed options are sometimes hidden there. Some older Spaces might not offer a direct embed option, in which case you might need to choose another Space. You'll know it worked when a new window or panel pops up, clearly displaying a block of HTML code, often within a text box. 💬 Example: `[Click on the </>]`

Step 5 — Copy the embed code

With the embed code generated, your next step is to copy it so you can paste it onto your blog or website. This code is the key to displaying your interactive weather widget. In the pop-up window or panel that appeared in the previous step, you will see the HTML embed code displayed. There is usually a "Copy" button right next to the code, or sometimes below it. Click this "Copy" button. This action will automatically copy the entire block of code to your device's clipboard, ready for you to paste it elsewhere. You don't need to understand what all the code means; just copy it as it is. **If it looks different:** If there's no dedicated "Copy" button, you can manually highlight all the text in the code box using your mouse (click and drag) and then right-click and choose "Copy" (or use keyboard shortcuts like Ctrl+C on Windows/Linux or Cmd+C on Mac). You'll know it worked when a small confirmation message like "Copied!" appears, or you can paste the code into a simple text editor to see it. 💬 Example: `[Click the "Copy" button next to the HTML code]`

Step 6 — Embed the widget on your blog

Now for the exciting part: adding the weather widget to your own blog! You'll need access to your blog's editing interface where you can add or edit content. Log in to your blog's administration area. Navigate to the post or page where you want to add the weather widget. Most blog platforms, like WordPress, offer two views when editing: a "Visual" editor (which shows you what your post will look like) and a "Text" or "Code" editor (which shows the raw HTML). You need to switch to the "Text" or "Code" editor. Find the exact spot in your blog post or page where you want the weather widget to appear, then paste the embed code you copied earlier. **If it looks different:** The name of the code editor might vary. Look for tabs or buttons labeled "HTML," "</>" (angle brackets), or "Source." If you're using a website builder with a drag-and-drop interface, look for an "Embed Code" or "Custom HTML" block you can add to your page. You'll know it worked when the embed code appears in your blog's text editor. Do not switch back to the "Visual" editor yet, as the widget might not display correctly there. 💬 Example: `[Switch to "Text" editor, paste HTML code]`

Step 7 — View your new weather widget

The final step is to save your changes and see your new AI weather widget live on your blog or website. Once you've pasted the embed code into your blog's HTML editor, save or publish the changes to your post or page. The exact button will depend on your blog platform, but it's usually "Update," "Publish," or "Save Changes." After saving, visit the live version of your blog post or page in a new browser tab. You should now see the Hugging Face weather widget displayed directly on your site, fully interactive and ready for your readers to use! **If it looks different:** If the widget doesn't appear, or you just see the raw code, double-check that you pasted it into the "Text" or "Code" editor, not the "Visual" one. Sometimes a browser refresh (F5 on Windows, Cmd+R on Mac) can help. You'll know it worked when you see the fully functional weather input box and display on your live blog page, allowing you to type a city and get weather results. 💬 Example: `[Click "Update" or "Publish" in your blog editor, then open your live blog page.]`

Common mistakes

- **Pasting into the wrong editor:** Many blog platforms have a 'Visual' editor and a 'Text' or 'HTML' editor. Pasting the embed code into the 'Visual' editor will just display the code as text, not the widget. **The fix:** Always switch to the 'Text' or 'HTML' view before pasting any embed code.
- **Not saving/publishing:** After pasting the code, it's easy to forget to save or publish the page. If you leave the editor without saving, your changes won't appear. **The fix:** Always click the "Update," "Publish," or "Save Changes" button in your blog's interface after adding the code.
- **Widget not loading:** Sometimes, a Space might be temporarily down or very busy, causing your widget to not load. The Hugging Face widget itself can't perform any actions beyond displaying information. **The fix:** Try refreshing the page, or go back to Hugging Face Spaces and try embedding a different, less busy weather model.

What this means for you

- **In everyday life:** You've learned a valuable skill in embedding interactive tools onto a website, which can be applied to many other types of widgets, not just weather. It's a simple way to make your personal site more dynamic and useful.
- **For your business or work:** Understanding how to embed third-party content is crucial for online presence. While this is a simple weather widget, the principle applies to embedding customer support chatbots, interactive data visualisations, or other AI-powered tools on a company website.
- **If you're just getting started:** You've taken your first step into integrating AI tools into your online life without needing to learn complex programming. This process demystifies how many online services share features across different websites.

Try it now

Visit the blog page where you embedded your new weather widget. Try typing in your current city, or a famous landmark city like "London" or "Tokyo," and see what weather information the AI provides!

How to Build a Simple Weekly Schedule in Google Sheets with Gemini

7 steps · Read online

> **In short:** You can ask Gemini, the AI built into Google Sheets, to design a weekly schedule from a plain description of your week — with time blocks, color categories, and a formula that totals hours automatically. You only need a Google account.

Most weeks you can list your time blocks in your head — work, gym, family dinner, the Saturday coffee — yet you never actually write them down. That's the gap this guide fills. In about ten minutes, you'll describe your week in plain English and end up with a color-coded schedule in Google Sheets that adds up your hours for you. The result is something you'll actually open on Sunday evening, not a spreadsheet that dies in your downloads folder.

Before you start

- A Google account (free — the same one you use for Gmail).
- A computer or phone that can open Google Sheets: go to sheets.google.com in any browser, or install the free Google Sheets mobile app.
- Roughly 10 to 15 minutes, and a rough mental picture of what your week actually looks like.
- Some Gemini features inside Sheets are only available on paid Google plans — if you don't see the Gemini option, check Google's official pricing page for the latest plan details. A free workaround: draft the schedule at gemini.google.com first, then paste it into Sheets.

Step 1 — Open Google Sheets and start fresh

Open sheets.google.com in your browser, or launch the Google Sheets app, and sign in with your Google account. Click the "+" button or "Blank spreadsheet" to create a new, empty sheet. Name it by clicking "Untitled spreadsheet" at the top-left and typing something you'll recognize later — "My weekly schedule" works fine.

💬 Example: name it "Spring weekly schedule" or "My week plan".

If it looks different: on a phone, the "+" usually sits in the bottom-right corner; if you only see templates, tap "Blank" inside that section. Menu wording changes often — just look for any option that creates a fresh, empty sheet.

You'll know it worked when you see a grid of empty cells with column letters (A, B, C…) across the top and row numbers (1, 2, 3…) down the side.

Step 2 — Open the Gemini side panel

Look at the top-right corner of the Sheets window for a small sparkle ✨ icon or a button labeled something like "Help me" or "Ask Gemini." Click it. A side panel slides in from the right edge — this is your chat window, and you'll stay inside your spreadsheet the whole time.

💬 Example: type "Hi" as a quick first message to confirm Gemini is responding.

If it looks different: Gemini's exact location in Sheets has moved over the years. If you don't see a sparkle, open the "Tools" menu (sometimes labeled "Extensions" on older versions) and look for anything mentioning Gemini or "Ask." On Workspace accounts, your admin may need to turn Gemini on — if the option is entirely missing, ask your IT admin or check Google's help pages.

You'll know it worked when you can type a question into the side panel and an answer appears in the same panel within a few seconds.

Step 3 — Describe your week and ask for a template

In the Gemini side panel, write a short paragraph describing a typical week. Be specific — even if it feels messy. Then ask for a weekly schedule layout with time blocks, categories you can color-code, and a formula that totals hours per category. Gemini will reply either with a written layout you can copy, or with a button to insert a table directly into your sheet.

💬 Example: "Here's my week. I work Monday to Friday, 9 a.m. to 5 p.m. I teach a yoga class Tuesday and Thursday evenings from 6 to 7:30 p.m. I hit the gym three mornings a week, around 6:30 to 7:30 a.m. on Monday, Wednesday, and Friday. Saturday morning is family breakfast from 9 to 11 a.m. Please build a weekly schedule in this sheet with rows for each day, columns for morning, midday, afternoon, and evening, and a SUM formula that totals the hours automatically."

If it looks different: some versions of Gemini in Sheets can only chat (no insert button). If that happens, copy the layout into the cells yourself — see Step 4.

You'll know it worked when you see either a table inserted into the sheet or a clear text description of where each activity should go.

Step 4 — Build the time-block grid yourself (the easy way)

If Gemini gave you a written description, copy it into your sheet by hand. In row 1, type the days of the week across the columns (Monday in A1, Tuesday in B1, and so on — or use one column per time slot if you prefer). In column A, starting from row 2, type your time blocks: "6:30–7:30

a.m.", "9:00 a.m.–5:00 p.m.", "6:00–7:30 p.m.", and so on. Then drop each activity into the cell where day and time meet.

Example: in cell B2 (Tuesday morning) type "Gym"; in cell C5 (Wednesday afternoon) type "Deep work block"; in cell A8 (Monday evening) type "Family dinner."

If it looks different: some people prefer one row per day and one column per time slot — both are normal layouts. Choose the one that's easier to read on your device.

You'll know it worked when you can read down a column to see one full day's plan, or across a row to see one time slot across the week, without guessing.

Step 5 — Add color categories so your eyes can scan it fast

Highlight the column or row headers by clicking the letter at the top of the column (or the row number on the left), then pick a fill color from the toolbar. Look for a paint-can icon or a "Fill color" button. Then color each activity by type: one color for work, one for fitness, one for family, one for rest. Pick colors that mean something to you.

Example: blue for work, green for fitness, yellow for family time, gray for rest or sleep.

If it looks different: the fill-color icon sometimes looks like a tipped-over paint bucket. On a phone, the same tool lives inside the formatting menu, often marked with an "A" and color bars underneath it.

You'll know it worked when you can glance at the sheet and instantly see which parts of your week are work, which are fitness, and which are family time — without reading a single word.

Step 6 — Add a formula that totals your hours

At the bottom of one day's column, click an empty cell and type a SUM formula. In Sheets, every formula starts with `=`. If you've put the duration of each activity in plain numbers (8 for an 8-hour workday, 1.5 for a 90-minute class), the formula adds them up for you.

Example: type `=SUM(B2:B10)` and press Enter. The cell instantly shows the total. If your durations are in seconds, multiply by 24 first: `=SUM(B2:B10)*24`.

If it looks different: if your time blocks are written as text like "9–5", Sheets can't add them up — they'd need to be plain numbers in a separate "Hours" column. Ask Gemini in the side panel, "How do I add up time blocks written as text in Sheets?" if you get stuck.

You'll know it worked when the total cell shows a number that updates the moment you edit any of the time blocks above it — that's the formula doing its job.

Step 7 — Save once and forget about it

Look at the top-left of the screen for your file name. As long as you're signed in, Google Sheets saves your work automatically every few seconds — there's no "Save" button to click. If you see a small gray "Saving…" message or a checkmark next to the title, you're done.

Example: just close the browser tab when you finish. Next time you open the file, your schedule will be right where you left it.

You'll know it worked when you can refresh the page or close and reopen the file and find your schedule exactly as it was.

Common mistakes

- **Mistake: asking Gemini something too vague.** "Make me a schedule" gives a generic template that won't fit your life. **Always include your real hours, real activities, and how many days — the more specific you are, the better the output.**
- **Mistake: writing time blocks as text like "9–5".** Sheets can't add text up, so your total stays at zero. **Fix: put the duration as a plain number in a separate "Hours" column (8 for an 8-hour workday, 1.5 for a 90-minute class), or ask Gemini how to convert time strings to numbers.**
- **Mistake: assuming Gemini inserted your schedule when it only chatted.** Some versions of Gemini in Sheets can only suggest, not insert. **Fix: read the side panel carefully — if you see an "Insert" button, click it; otherwise copy the suggestions into the cells yourself.**

What this means for you

- **In everyday life:** A weekly schedule you actually open turns vague plans into real commitments. The color categories mean you can spot at a glance if "rest" has quietly disappeared from your week — and fix it before burnout shows up.
- **For your business or work:** If you bill by the hour or track project time, the SUM formula and color categories double as a rough time-tracking sheet. Swap the activity names for your projects and you have a one-page weekly dashboard.

- **If you're just getting started:** Don't aim for the perfect schedule in one sitting. Type in the bones of your week today — work, sleep, family — and add the rest over a few days as you notice the gaps.

Try it now

Open a new Google Sheet, open the Gemini side panel, and paste this in: "I teach three morning classes Monday, Wednesday, and Friday from 8 to 11 a.m. I work part-time Tuesday and Thursday from 1 to 6 p.m. I have family dinner every weekday at 7 p.m. Build me a weekly schedule with time blocks, three color categories (teaching, work, family), and a SUM formula that totals hours per category." See what it gives you, change the colors to match your real life, and you've got a planner that took ten minutes to make and will save you hours of "wait, what day is it again?" by next Sunday.

Build a Simple Q&A Chatbot for Your Website with Cohere

6 steps · Read online

By the end of this guide, you'll have a basic chatbot on your website that can answer your most common customer questions. This approach uses Cohere's powerful artificial intelligence (AI) models and requires a simple copy-paste of some code onto your website, making it accessible even if you're new to coding. This chatbot can draft responses, but it won't process payments or manage complex user accounts without further development.

Before you start

- **Cohere Account:** A free or paid Cohere account. The steps in this guide work with a free plan, though paid plans offer higher usage limits.
- **Website Access:** You need administrative access to your website to add custom HTML (HyperText Markup Language) or JavaScript code. This is typically available through platforms like WordPress, Wix, Squarespace, or Shopify.
- **Common Questions:** A short list of the top 3-5 questions your customers frequently ask, along with your desired answers.
- **Time:** Allocate about 30-45 minutes to follow these steps and test your new chatbot.

Step 1 — Create or Log In to your Cohere Account

First, open your web browser and go to the Cohere website at `cohere.com`. Look for a **Sign Up** or **Log In** button, usually in the top-right corner. If you're new, follow the prompts to create your account, which will involve verifying your email address. If you already have an account, simply log in with your credentials. You'll know it worked when you see the Cohere console or dashboard after successfully logging in, which usually displays options like "Playground," "Models," or "API Keys."

Step 2 — Find Your API Keys and Playground

Once logged in, you'll need two main things: an **API key** and access to the **Playground**. An API key is like a secret password that lets your website talk to Cohere's AI services, and the Playground is where you can test how Cohere's AI models respond. Navigate your dashboard to find "API Keys" and "Playground" (or similar names). A new screen will open, showing your existing API keys or an option to create a new one. Click on **Create New Key** if you don't have one, and give it a memorable name like "MyChatbotKey." Then, switch to the "Playground" section. If you can't find these options directly, look for a left-hand navigation menu or a "Developers" section. You'll know it worked when you have a long string of letters and numbers (your API key) copied to your clipboard or saved safely, and you're on a page with a text input box labeled "Prompt" in the Cohere Playground.

Step 3 — Craft Your Chatbot's Responses in the Playground

In the Cohere Playground, you can teach the AI how to respond to specific questions. A **prompt** is the instruction or question you give to the AI, and the **response** is the answer it generates. For a simple Q&A chatbot, you'll provide specific instructions for common questions. Type your first question into the "Prompt" box. After typing your prompt, click the **Run** or **Generate** button, usually found below the prompt box. The AI will then generate a response in the output area. You can adjust your prompt or the AI's settings (like "Temperature" or "Max Tokens" which control creativity and length) until you get the perfect answer for your chatbot. 💬 Example: `As a friendly customer service bot for "My Awesome Shop", what is your return policy?` If the response isn't what you expect, try rephrasing your prompt with more detail. You'll know it worked when the AI generates a clear, helpful response that you'd like your chatbot to give customers for that specific question.

Step 4 — Prepare the Chatbot's Logic with a Simple Code Snippet

Now, you need a small piece of JavaScript code that will connect your website to Cohere's API and display your chatbot. This code will act as a digital messenger, sending customer questions from your site to Cohere and bringing back the AI's answers. You'll need your Cohere API key from Step 2. Open a plain text editor (like Notepad on Windows or TextEdit on Mac, or any code editor) and copy-paste the following basic JavaScript snippet into it. Replace `"YOUR_COHERE_API_KEY"` with the API key you saved earlier, and feel free to adjust the `welcomeMessage` or `botName`. This snippet creates a simple chat interface and makes a request to Cohere's Chat API. 💬 Example:

```
<div id="ai-world-chatbot"></div>
<script>
  const chatbotContainer = document.getElementById('ai-world-chatbot');
  chatbotContainer.innerHTML = `
    <div style="position: fixed; bottom: 20px; right: 20px; width: 300px; height: 400px; border: 1px solid #ccc; border-radius: 8px; box-sha
      <div style="background-color: #007bff; color: white; padding: 10px; border-top-left-radius: 8px; border-top-right-radius: 8px;">
        <h3 style="margin: 0;">My Shop Helper</h3>
      </div>
      <div id="chat-messages" style="flex-grow: 1; padding: 10px; overflow-y: auto; background-color: #f9f9f9;">
        <p><strong>Bot:</strong> Hello! How can I help you today?</p>
      </div>
```

```
      <input type="text" id="chat-input" placeholder="Type your message..." style="padding: 10px; border: none; border-top: 1px solid #eee;
    </div>
  `;

  const chatMessages = document.getElementById('chat-messages');
  const chatInput = document.getElementById('chat-input');
  const COHERE_API_KEY = "YOUR_COHERE_API_KEY"; // Replace with your actual Cohere API Key

  async function sendMessageToCohere(message) {
    const response = await fetch("https://api.cohere.ai/v1/chat", {
      method: "POST",
      headers: {
        "Content-Type": "application/json",
        "Authorization": `Bearer ${COHERE_API_KEY}`
      },
      body: JSON.stringify({
        message: message,
        model: "command-r-plus", // You can choose a different model
        // Add more context or conversation history here for a smarter bot
      })
    });
    const data = await response.json();
    return data.text;
  }

  chatInput.addEventListener('keypress', async function(e) {
    if (e.key === 'Enter' && chatInput.value.trim() !== '') {
      const userMessage = chatInput.value.trim();
      chatMessages.innerHTML += `<p><strong>You:</strong> ${userMessage}</p>`;
      chatInput.value = '';

      const botResponse = await sendMessageToCohere(userMessage);
      chatMessages.innerHTML += `<p><strong>Bot:</strong> ${botResponse}</p>`;
      chatMessages.scrollTop = chatMessages.scrollHeight; // Auto-scroll to latest message
    }
  });
</script>
```

If this snippet looks different from Cohere's official documentation in the future, always use the most up-to-date example from Cohere's developer resources. You'll know it worked when you have your customized code snippet, including your API key, saved in a plain text file.

Step 5 — Add the Code Snippet to Your Website

Now it's time to add your chatbot to your live website. You'll need to paste the code snippet you prepared in Step 4 into the HTML of your website. Most website platforms, or **CMS (Content Management Systems)** like WordPress, Shopify, or Wix, have a way to add custom code. Look for options such as "Custom HTML," "Theme Editor," "Code Injection," or "Embed Code." Paste the entire snippet (starting with `<div id="ai-world-chatbot">` and ending with `</script>`) just before the closing `</body>` tag of your website's HTML. This ensures the chatbot loads correctly after all your site's content. Save your changes and publish your website. If you're using a CMS like WordPress, you might use a plugin for custom code, or navigate to Appearance > Theme File Editor > footer.php. On Shopify, go to Online Store > Themes > Actions > Edit code > theme.liquid. Look for the `</body>` tag and paste your code above it. You'll know it worked when you've saved and published your website changes, and there are no immediate error messages from your CMS.

Step 6 — Test Your Chatbot

With the code now live on your website, it's time to test if your chatbot is working correctly. Open your website in a fresh browser window or tab. You should see a small chat bubble or a simple chat window appear, usually in the bottom-right corner of your screen, based on the basic styling in our snippet. Click on the chat window to open it and type in one of the common questions you prepared in Step 3. Press Enter. The chatbot should send your message to Cohere and display an AI-generated response shortly after. If the chat window doesn't appear, try clearing your browser's cache and reloading the page, or double-check your website's custom code area for any typos. You'll know it worked when you can type a question into the chat window on your website and receive a relevant AI-generated response back.

Common mistakes

- **Incorrect API Key:** Forgetting to replace `"YOUR_COHERE_API_KEY"` with your actual key, or using an expired/incorrect one. Always double-check your API key and ensure it's active in your Cohere dashboard.
- **Wrong Placement of Code:** Placing the JavaScript snippet inside the `<head>` tag or in a part of your website that isn't rendered properly can prevent the chatbot from appearing. Always aim for just before the `</body>` tag for client-side scripts.
- **Not Publishing Changes:** After adding code to your website's editor, you must save and publish the changes for them to go live. Many beginners forget this final step in their CMS.

What this means for you

- **In everyday life:** You've learned how to connect powerful AI models to a real-world application, automating simple tasks and understanding how AI services can be integrated into digital tools.
- **For your business or work:** This basic chatbot can handle routine customer questions 24/7, freeing up your time to focus on more complex tasks and improving customer satisfaction by providing instant answers.
- **If you're just getting started:** You've successfully deployed your first AI-powered tool! This hands-on experience builds a solid foundation for exploring more advanced AI features or customizing your chatbot further.

Try it now

Revisit your Cohere Playground and experiment with refining a prompt to give an even better answer to a common customer query, like "What are your business hours?". Make sure your custom JavaScript snippet on your website is still open in your text editor. Then, on your live website, open the chatbot and ask that question to see your AI in action.

Set up named medication reminders with Alexa+ for yourself or an older relative

6 steps · Read online

> **In short:** Alexa+ (Amazon's upgraded voice assistant, on Echo devices and the Alexa app) can speak reminders to a specific person by name, on a daily schedule, at a specific time of day. It's not a medical device, and it won't replace pillboxes or a caregiver — but it's a useful extra nudge, especially for older relatives who live alone.

Alexa+ is the newer version of Amazon's Alexa assistant that runs on Echo speakers and Echo Show screens. The "named reminder" feature is the part you care about here: you can tell it who the reminder is for, and Alexa+ speaks the right words when that person is nearby. This guide walks you through it on both the Alexa app and an Echo device, for yourself or for a parent or grandparent.

Before you start

- An **Amazon account**. Free to create at amazon.com if you don't already have one.
- The **Alexa app** on a phone or tablet (free on the App Store and Google Play), OR an **Echo device** like an Echo Dot, Echo Show, or Echo Spot. At least one of these two — the app, an Echo, or both — is enough.
- The person receiving the reminders needs to be set up in Alexa+ as a **named voice profile** (Alexa+ can learn to recognize a specific voice, then call that person by name). Setup takes about 5 minutes the first time.
- About **10 minutes total** for this guide. After that, reminders run themselves.
- One honest limit up front: Alexa+ speaks reminders when it hears the wake word ("Alexa"), and it works best when the Echo is in the same room as the person who needs to hear it. If your relative is hard of hearing, an **Echo Show with a screen** is the better pick — the reminder also pops up as text.

Step 1 — Open the Alexa app and sign in

Tap the **Alexa app** icon on your phone — it's a blue speech-bubble with a white ring around it. The home screen slides up, showing your devices and recent activity. If you're asked to sign in, type the email and password tied to your Amazon account.

If it looks different: if you don't see the app on your home screen, search "Alexa" in your App Store or Play Store and install it. If the home screen shows a different layout (Amazon updates this often), look for a **person icon or "More" tab** in the bottom-right corner — settings live there.

Example: at the sign-in screen, tap **"Already a customer? Sign in"**, then enter your Amazon email and password.

You'll know it worked when your first name appears in the top-left corner of the home screen, and you can see a list of your Echo devices under a heading like "Devices."

Step 2 — Create a Voice ID for the person who needs reminders

Tap the **More** tab (or three-line **☰ menu**), then choose **Settings → Your Profile → Voice ID**. Tap **"Create voice ID"** and follow the prompts. The app asks the other person to repeat five short phrases out loud, like "Alexa, play music" and "Alexa, what's the weather?" — they don't need to be perfect. When training finishes, Alexa+ knows that voice.

If it looks different: some screens say "Recognize your voice" instead of "Voice ID" — both lead to the same setup. On older Alexa (not Alexa+), the option may be under **Settings → Alexa Account → Recognized Voices**.

Example: the prompts look like this — *"Alexa, who am I?"* and *"Alexa, play my Flash Briefing."* Speak them naturally. If your relative has a soft voice, ask them to sit about a foot from the phone.

You'll know it worked when the app says "Voice ID created for [name]" and you can now set reminders addressed to that person.

Step 3 — Create a named medication reminder in the app

Still in the Alexa app, go to **More → Reminders & Alarms → Add Reminder** (or tap the **+** button on the Reminders screen). Type the reminder, choose the date and time, then look for a field labeled **"To"** or **"Announce to"**. Pick the person you just trained — by name, not "everyone." Save it.

If it looks different: the wording shifts often. On some screens you'll find **Routines** instead — that's fine, a Routine can do the same thing plus more. Also: many Echo Show devices have a touchscreen, and you can create reminders right there: swipe down from the top of the screen, tap **Settings → Reminders → New Reminder**. Same idea, different buttons.

Example: type something plain — *"Remind Mom to take her blood pressure pill"*, set it for **8:00 a.m.**, set it to **repeat every day**, and assign it **To: Mom**. Keep the wording short — Alexa+ reads it back aloud, so it should sound natural when spoken.

You'll know it worked when the reminder appears in your Reminders list with the person's name attached, and the app confirms "Reminder set."

Step 4 — Test it on the Echo, out loud

Walk over to the Echo device in the room where your relative spends the most time — usually the kitchen or living room. Say clearly: **"Alexa, set a reminder for Mom at 8 a.m. every day to take her blood pressure pill."** The light ring on top of the Echo spins blue while Alexa+ thinks, then it speaks back: *"Okay, I'll remind Mom every day at 8 a.m. to take her blood pressure pill."*

If it looks different: on an Echo Show (the one with the screen), the same words work, but you'll also see the reminder pop up as text on the display. If your Echo doesn't seem to hear you, try saying the wake word first ("Alexa...") before your sentence, and speak a little louder than normal — older microphones prefer a clear voice.

Example: *"Alexa, set a reminder for Dad every evening at 7 p.m. to take his evening pill."* Notice you don't need the app at all for this — the Echo itself can create, list, and delete named reminders.

You'll know it worked when Alexa+ repeats your exact request back to you out loud. That verbal confirmation is your success check.

Step 5 — Listen for it the next day, and tweak the time

The next morning at the set time, stand a few feet from the Echo and listen. Alexa+ should say something like: *"Mom, here's your reminder: take your blood pressure pill."* That name in front is the whole point — it's what tells you who's meant to hear it. If 8 a.m. turns out to be the wrong moment (maybe your relative is still asleep), delete the reminder and recreate it: **"Alexa, delete my reminder for Mom."** Then create a new one at a better time.

If it looks different: if you don't hear anything at the scheduled time, check that the Echo is **plugged in, not muted** (a red ring around the top means it's muted — press the mute button on top to unmute). Also, reminders won't fire if the Echo has been unplugged or lost Wi-Fi, so make sure the ring glows blue at the set time.

Example: *"Alexa, list my reminders"* — Alexa+ will read every reminder you have set, including who each one is for. Use this to double-check before the day starts.

You'll know it worked when you hear the person's name spoken out loud at the exact time you set, in the room where the Echo lives.

Step 6 — Add a backup: a repeat chime or visual cue on Echo Show

For an extra layer, especially with someone who doesn't always hear well, pair the reminder with a Routine that also plays a chime or shows a notification. In the app, go to **More → Routines → +** and set it to trigger at the same time as the reminder. Under "Add action," pick **Sounds → Play chime**, or for Echo Show, **Announcements → Display a message**. Save and turn it on.

If it looks different: if the Echo is an older model without a screen, skip this step — the spoken reminder on its own is enough. You can also buy a smart plug and have Alexa+ flash a lamp when the reminder goes off, but that's optional.

Example: under the Routine's "When this happens," pick the reminder you made in Step 3. Under "Add action," pick **Sounds → Bell chime** (or whatever chime the app offers). Save the Routine and you're done.

You'll know it worked when at 8 a.m. you hear both Alexa+'s spoken reminder and the chime — two ways to catch your relative's attention.

Common mistakes

- **Forgetting to assign the reminder to a person.** Without a name, Alexa+ says "you have a reminder" — which is useless if your relative isn't sure the message is for them. Always set the **"To"** field to the specific person.
- **Setting the time too early.** If your relative sleeps until 9 a.m., an 8 a.m. reminder just becomes background noise. Test for a few days and shift the time forward — voice reminders are easy to delete and recreate.
- **Relying on a single Echo device in a quiet corner.** The reminder only works in the room where the Echo sits, and only when it's plugged in and connected to Wi-Fi. For someone who's often in the kitchen or bathroom, place the Echo where the morning routine happens, not in a bedroom down the hall.
- **Treating Alexa+ as a medical device.** It's a friendly nudge, not a guarantee. Always keep using a labeled pillbox, and never skip human check-ins because "the Echo will remind them."

What this means for you

- **In everyday life:** if you take a daily vitamin or prescription yourself, you can set a quick reminder from the kitchen counter — "Alexa, remind me to take my vitamin at 7:30 a.m. every day" — without ever opening the app.

- **For your family or caregiving:** named reminders turn one Echo into a gentle, shared checklist for everyone in the household. Mom's pill, Dad's evening walk, your kid's homework time — each person hears only their own name.
- **If you're just getting started:** begin with one reminder, one person, one time of day. Get that working for a week, and only then add more — too many reminders at once is overwhelming, and people start ignoring them all.

Try it now

Set one reminder today. Walk to the nearest Echo and say out loud, plainly: **"Alexa, set a reminder for me every morning at 8 a.m. to take my vitamin."** That's it — one sentence, zero app menus. Tomorrow morning, Alexa+ will speak your name back to you, and you'll know the whole system works. From there, you can add a named reminder for your relative whenever you're ready.

Build a Basic AI Chatbot Brain Using OpenAI's API – Quick Start

6 steps · Read online

Ever wanted to give your website or application a smart assistant? This guide shows you how to set up the core intelligence, or "brain," of a simple AI chatbot using OpenAI's powerful **API (Application Programming Interface)**. By the end, you'll understand how to instruct an AI to answer questions in a specific way, ready for future integration into your projects.

Before you start

- **An active OpenAI account:** You can create one for free on the OpenAI website.
- **Access to the OpenAI API:** This requires adding billing details to your OpenAI account to activate API usage. While free credits may be available, charges apply once these are used.
- **A computer with a web browser:** Chrome, Firefox, Edge, or Safari will work.
- **Rough total time:** Allow about 15-20 minutes to follow all the steps.

Step 1 — Sign In to Your OpenAI Account

Your first step is to log into the OpenAI platform where you'll access the tools. Open your web browser and navigate to the OpenAI website. On the main page, look for a "Log in" or "Sign in" button, typically located in the top-right corner of the screen. Click this button and enter your email address and password.

You'll know it worked when a new screen loads, displaying your OpenAI dashboard or home page, where you might see options for "ChatGPT", "API", or "DALL-E".

Step 2 — Get Your Secret Key (API Key)

To let programs communicate with OpenAI's powerful AI models, you need an **API key**. Think of an API key like a secret password that authorises your requests. It tells OpenAI that it's you making the request and helps manage your account's usage.

Once logged in, look for your profile icon or name, usually in the top-right corner of the screen. Click it, and from the dropdown menu, select "API keys". On the API keys page, you'll see a button like "Create new secret key". Click this, then give your key a memorable name, such as "My Website Chatbot Brain", and click "Create secret key" again. **Important:** Immediately copy the string of characters that appears. You won't see it again after this window closes, so paste it somewhere safe (like a temporary text file) or have it ready for the next step.

You'll know it worked when you've securely copied a long string of letters and numbers (e.g., `sk-xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx`).

Step 3 — Open the OpenAI Playground

Now, let's head to the **OpenAI Playground**. This is a user-friendly web interface provided by OpenAI that allows you to experiment with their AI models and API features without writing any code. It's like a sandbox for AI experiments.

From your OpenAI platform page, look for a left-hand navigation menu or a central dashboard button that says "Playground". Click on "Playground" to load the interface. Once loaded, ensure you are in the "Chat" mode, which is usually indicated by a dropdown menu or a tab at the top of the Playground, sometimes labeled "Mode: Chat". If it looks different, you might be in "Completions" mode; simply find the "Chat" option to switch over.

You'll know it worked when you see a panel on the left with a section titled "System" and another below it for "User" messages.

Step 4 — Define Your Chatbot's "Personality" (System Message)

This is where you give your chatbot its core identity! The **system message** is a set of instructions you provide to the AI *before* any conversation begins. It tells the AI its role, desired tone, and what kind of information it should focus on. Imagine it as writing a detailed job description for your AI assistant.

In the OpenAI Playground, locate the large text box labeled "System" on the left side of the screen. Click inside this box and begin typing your instructions. Be clear, specific, and tell the AI exactly what you want it to be.

💬 Example: `You are a friendly and helpful customer service assistant for "AI World HQ" Your main goal is to answer questions about general AI topics and guide users to relevant guides on our website. Be concise, polite, and avoid giving medical, financial, or legal advice. If a question is outside your knowledge or scope, politely state that you cannot assist with that specific query.`

You'll know it worked when your chosen instructions are fully visible within the "System" message box.

Step 5 — Have Your First Conversation

With your chatbot's "personality" defined, it's time to test its brain by simulating a conversation! Below the "System" box in the OpenAI Playground, you'll find a section labeled "User". This is where you'll type a message, just like a website visitor would.

Type a question that fits within the role you defined in your system message. For instance, if you set its role as a customer service assistant, ask a relevant question. Once you've typed your message, look for a "Submit" button, usually green and found at the bottom of the chat interface. Click it to send your message to the AI. If it looks different, check for a "Send" icon or similar.

Example: **User:** `What is AI?` *(Click "Submit")*

You'll know it worked when the AI's response appears directly below your "User" message, often labeled "Assistant" or "AI's Reply".

Step 6 — Refine and Experiment

The power of the system message lies in its flexibility. You can tweak it endlessly to fine-tune your chatbot's behavior. After your first conversation, you might notice areas for improvement or new ideas for its role.

Go back to the "System" box and adjust your instructions. For example, you could tell it to be more empathetic, or to always ask a follow-up question. After making changes, type another question into the "User" box and click "Submit" again. The AI will now process your new message using the *updated* system instructions. Experiment by challenging it with queries outside its defined scope to see how it adheres to your boundaries.

You'll know it worked when the AI's subsequent responses clearly reflect the changes you made to its "personality" in the system message.

Common mistakes

- **Sharing your API key publicly:** Your API key is like a master password to your OpenAI account. Never embed it directly into front-end website code or share it anywhere public, as others could use it to make requests under your account, potentially incurring significant costs. Always keep it secure and private.
- **Unclear or contradictory system messages:** If your chatbot gives unexpected or unhelpful answers, the problem is often in its "job description." Ensure your instructions are specific about its role, desired tone, and what it *shouldn't* do. Ambiguity can lead the AI astray.
- **Expecting the AI to be infallible or magical:** While powerful, AI can **hallucinate** (confidently make up information) or struggle with complex nuances. This process defines its "brain," but the API alone doesn't integrate it into a user interface or guarantee perfect information. Always verify critical details, especially in a business context.

What this means for you

- **In everyday life:** Understanding how to craft a system message helps you get better, more tailored results from personal AI tools like ChatGPT. It teaches you to be a more effective communicator with AI, guiding it to focus on what matters to you.
- **For your business or work:** You've just laid the groundwork for a potential custom AI assistant. This core intelligence can power a customer service bot, a virtual sales assistant, or an internal knowledge base helper on your website or app. It can answer FAQs, provide product information, or assist with troubleshooting, freeing up your team's time.
- **If you're just getting started:** The easiest first move is to spend more time in the OpenAI Playground. Experiment with different system messages and user prompts to see how drastically you can change the AI's behavior and responses.

Try it now

Head back to the OpenAI Playground, update your system message to make your chatbot a "friendly, expert travel agent for the US," and then ask it, "What's the best way to plan a trip to the Grand Canyon?"

Create a personalized workout routine with DeepSeek's multilingual chat

6 steps · Read online

> **In short:** DeepSeek's AI chat can generate a personalized workout plan in many languages — just describe your goals, fitness level, and available equipment, then ask for a weekly routine.

Imagine typing a few sentences about your fitness goals and receiving a ready-to-use, day-by-day workout plan in your language. That's exactly what DeepSeek's multilingual chat can do for you. No gym membership, no personal trainer, just a simple chat and a plan you can start today.

Before you start

- A free DeepSeek account (sign up at deepseek.com).
- A phone, tablet or computer with internet.
- About 10 minutes of your time.

Step 1 — Open DeepSeek Chat

Go to deepseek.com and look for the chat option — it's usually a "Chat" button or "Try DeepSeek" link near the top of the page. If you can't find it, try navigating directly to the chat section of the site. You'll see a clean page with a large text box in the center. A blinking cursor will appear in the box, ready for your first message. You'll know it worked when you see the blinking cursor in a text box labeled "Message DeepSeek..." or similar.

Example: Open your browser and go to deepseek.com, then find the chat interface

Step 2 — Say hello and set the scene

Type a friendly greeting so DeepSeek knows you're ready. This helps the AI understand you're starting a conversation. After you press Enter, the chat will reply with something like "Hello! How can I help you today?" If the reply doesn't appear, check your internet connection or refresh the page. You'll know it worked when you see DeepSeek's welcome message in the chat.

Example: Hi there, I want to build a simple workout plan

Step 3 — Describe your goals and limits

Tell DeepSeek exactly what you want and any limits you have. Be clear about your fitness level, available days, equipment and any injuries. The more details you give, the better the plan will fit you. If you're unsure what to say, just write "I'm a beginner, I have no equipment, I can train 3 days a week, and I want to build strength." You'll know it worked when DeepSeek replies with a summary like "Got it! You want a beginner-friendly, no-equipment, 3-day strength plan."

Example: I'm a complete beginner, I have no weights or gym access, I can train Monday, Wednesday and Friday mornings, and I want to build strength and feel more energetic

Step 4 — Ask for a day-by-day plan

Now ask DeepSeek to create a weekly workout routine based on what you just told it. Use simple words like "create a weekly workout plan" or "give me a 3-day beginner routine." DeepSeek will reply with a clear list of exercises, sets and rest times. If the reply is too long or unclear, type "shorten it to 5 exercises per session" to get a tighter plan. You'll know it worked when you see a numbered list of exercises with instructions like "3 sets of 10 squats, rest 60 seconds."

Example: Create a simple 3-day weekly workout plan for me, no equipment, beginner level

Step 5 — Choose your language

DeepSeek can speak many languages. If you'd like the plan in your language, just ask. Type "please give the plan in [your language]" and DeepSeek will rewrite everything so it feels natural to you. If the language switch doesn't happen, try typing the same request again or refresh the page. You'll know it worked when the exercise names and instructions appear in your preferred language.

Example: Please give the plan in Spanish

Step 6 — Save or print your plan

Once you're happy with the plan, copy it to a notes app, print it, or save it as a PDF. DeepSeek can't save files for you, but it can format the plan nicely if you ask. Type "format this as a printable table" and DeepSeek will return a clean layout you can screenshot or print. If the formatting looks messy, just ask again or copy the text into any document. You'll know it worked when you have a neat table or list you can take to your training space.

Example: Format this as a printable table so I can print it

Common mistakes

- **Too vague**: If you only say "I want a workout," DeepSeek might give a generic plan that doesn't fit you. Fix: Add details like "beginner," "no equipment," "3 days," and "strength."
- **Ignoring rest**: If you skip rest times, your plan may feel impossible. Fix: Ask DeepSeek to include rest periods or look for "rest 60 seconds" in every exercise line.
- **Not checking language**: If the plan is in English and you prefer another language, the exercises might feel unfamiliar. Fix: Ask DeepSeek to rewrite the plan in your language before you start.

What this means for you

- **In everyday life**: You'll have a clear, personalized plan that fits your schedule and language, so you can start today without confusion.
- **For your business or work**: If you manage wellness programs, you can use DeepSeek to quickly draft beginner-friendly routines for colleagues, saving hours of planning.
- **If you're just getting started**: This is the easiest first move—open the chat, tell DeepSeek your goals, and receive a ready-to-use plan in minutes.

Try it now

Open your browser, go to deepseek.com, find the chat, type "Hi, I'm a beginner with no equipment and I want a simple 3-day strength plan," and copy the reply into your notes. You'll have your first workout plan in under two minutes.

Automate Git Commands with GitHub Copilot Suggestions

5 steps · Read online

If you've ever felt overwhelmed by the precise commands needed to save and organize your programming work, or simply track changes to any project, you are not alone. This guide will show you how to let GitHub Copilot do the heavy lifting, turning your natural language requests into ready-to-use Git commands directly inside your code editor. It's like having a helpful assistant for your project's version control.

Before you start

- **Visual Studio Code (VS Code)**: A popular, free code editor. Ensure you have it installed on your computer.
- **GitHub Copilot**: GitHub Copilot has a free tier for individual users. If you need advanced features or are part of an organization, you may need a paid subscription. Check the official pricing page to see what's right for you.
- **VS Code Extensions**: Install the `GitHub Copilot` and `GitHub Copilot Chat` extensions within VS Code. You can find these in the Extensions view (a square icon on the side panel).
- **Time**: This guide will take approximately 10-15 minutes to follow through.

First, let's clarify two essential terms. **Git** is a version control system—think of it as a digital time machine for your project files, keeping track of every change so you can undo mistakes or explore different versions. The **terminal** (or command line) is a text-only window where you type direct instructions to your computer, rather than clicking buttons, like giving specific commands to a digital assistant.

Step 1 — Open your project and the terminal

To begin, we need to open your project folder in your code editor and access the built-in terminal. This lets you ask for help and run commands exactly where you're working.

Open Visual Studio Code and use the "File" menu (top-left) to select "Open Folder...", then navigate to and select your project folder. Once open, you'll see your project files listed on the left. Next, open the integrated terminal: look for the "Terminal" menu at the top of VS Code and select "New Terminal." A new panel will appear at the bottom of your screen, showing a text-based input area. If your screen looks different, you might find a plus (+) icon within the bottom panel that creates a new terminal, or try the keyboard shortcut `Ctrl + \`` (backtick) on Windows or` Cmd + `` on a Mac.

You'll know it worked when: The bottom panel opens and shows a text prompt, usually ending with `$` or `>` and your current folder name.

Step 2 — Ask Copilot for a basic save command

In Git, a **commit** is like saving a checkpoint in a video game, marking a specific set of changes to your project. Before you commit, you usually **stage** your changes, which means preparing them to be included in that checkpoint—think of it like packing items into a box before sealing it. Instead of trying to remember the exact commands to stage and save your files, you can just ask Copilot.

With your terminal open, you'll see a prompt where you can type. Instead of typing a Git command, look for a Copilot icon, or a small text box within or above the terminal panel that says "Ask Copilot." Click on this or use the shortcut `Ctrl + I` (or `Cmd + I` on Mac). A small input field will pop up, ready for your question. What happens next is Copilot listens to your natural language. If it looks different, you might need to make sure the "GitHub Copilot Chat" extension is enabled, or simply start typing in the main terminal prompt, and Copilot might automatically offer suggestions after you type a few words or trigger it with `/git`.

Example: Type: *"How do I stage all changes and commit them with the message 'Added initial project setup'?"*

You'll know it worked when: The Copilot chat box or inline suggestion area shows a generated command, likely starting with `git add . && git commit -m "..."`.

Step 3 — Run Copilot's suggested command

Once Copilot has generated a command for you, the next step is to run it in your terminal. This is where your instructions come to life and Git performs the action.

After Copilot suggests the command, you'll usually see a button or an option to "Insert" or "Run" the command. For example, a small button labeled "Insert into Terminal" or a right-arrow icon might appear next to the suggestion. Click this button to place the command into your terminal's input line. What happens next is the command will appear in your terminal, ready to be executed. Press the `Enter` key on your keyboard. If it looks different, sometimes Copilot might directly offer a "Run Command" button, which will execute it immediately without you pressing Enter. Always review the command before running it.

You'll know it worked when: The terminal shows messages indicating success, such as "[main (root-commit) ...] Added initial project setup" or "2 files changed, 10 insertions(...)", and your files have been saved in Git's history.

Step 4 — Create a new "sandbox" (Git Branch)

A **branch** is like creating a parallel universe or a separate workspace for your project. It lets you safely test new ideas or build new features without messing up the main, working version of your website or application. You can switch between branches easily, keeping your experiments separate until they're ready.

Open the Copilot chat box in your terminal again (either by clicking its icon or using `Ctrl/Cmd + I`). Now, ask Copilot to help you set up a new isolated space for your work. What happens next is Copilot will suggest the Git command to create and switch to this new branch in one go. If it looks different, ensure your Copilot Chat panel is visible; you might need to expand the bottom panel further to see it clearly, or just type your request directly into the main terminal prompt and wait for the inline suggestion.

Example: Type: *"Create a new branch and switch to it, name it 'feature/contact-page-update'."*

You'll know it worked when: Copilot suggests a command similar to `git checkout -b feature/contact-page-update`, ready for you to insert and run.

Step 5 — Safely combine your work (Git Merge)

Once you're happy with your changes in your sandbox branch (e.g., `feature/contact-page-update`), you'll want to bring them back into your main project. This process is called **merging**. Copilot makes this safe by reminding you of the correct order of steps, ensuring you don't accidentally overwrite your hard work.

Trigger the Copilot chat in your terminal once more. In the prompt, explain what you want to achieve. What happens next is Copilot understands that merging often involves a sequence of commands (first switching back to your main branch, then performing the merge). It will provide you with clear, distinct commands to follow. If it looks different, or Copilot only gives you one command, you can refine your prompt by asking: "Can you give me the *full sequence* of commands to merge...?"

Example: Type: *"How do I merge my current branch (feature/contact-page-update) back into the main branch?"*

You'll know it worked when: Copilot suggests a sequence, typically starting with `git checkout main` followed by `git merge feature/contact-page-update`, ensuring a smooth process.

Common mistakes

- **Blindly running commands:** Always read the command Copilot suggests before pressing Enter or the "Run" button. While AI is incredibly helpful, it doesn't know your project's specific history or nuances as well as you do. **Fix:** Take a moment to read the command and understand what `git add`, `git commit`, `git checkout`, or `git merge` actually do. If unsure, ask Copilot: *"Can you explain what this command means?"*
- **Forgetting your location:** Before running a command to save or merge, make sure you are in the correct project folder in your terminal. Running Git commands in the wrong folder can lead to errors or applying changes to the wrong project. **Fix:** If you're unsure, ask Copilot: *"How do I check which folder I am currently in?"* or simply type `pwd` (print working directory) and press Enter to see your current path.
- **Not providing enough context:** If Copilot gives a generic or unhelpful suggestion, it might not have enough information about your goal. For instance, just saying "merge" isn't as helpful as "merge my 'feature/design' branch into the 'main' branch." **Fix:** Be specific with your requests, mentioning branch names, commit messages, or what files you've changed.

What this means for you

- **In everyday life:** You no longer need to keep a notebook of confusing technical commands next to your keyboard. You can focus on creating and building your project rather than memorising complex syntax or looking up commands online. Copilot acts as your memory aid and guide.
- **For your business or work:** New team members can get up to speed with your team's project workflow much faster, reducing the time spent troubleshooting simple file-saving mistakes. It lowers the barrier to entry for contributing to technical projects.
- **If you're just getting started:** Use Copilot as an interactive tutor. Instead of just running the commands, ask it: *"Can you explain what each part of this command does?"* or *"Why do I need to stage files before committing?"* to learn the underlying concepts as you go.

Try it now

Let's do a quick test. Open Visual Studio Code and your project's integrated terminal. Trigger the Copilot chat (click its icon or `Ctrl/Cmd + I`), and then type: **"How do I check the status of my current files?"**

Watch how quickly it suggests `git status`—the foundational command to see what changes are pending or have been made in your project. This simple command is your first step to understanding your project's current state with Git, and Copilot helps you find it instantly.

How to Add an AI Step to Your Zapier Workflow

5 steps · Read online

> **In short:** In the regular Zap editor, look for the **AI by Zapier** action, type the instructions you want it to follow, and optionally hand it a small list of tools (like "send a Slack message" or "create a Google Sheet row"). No separate "agent" product to learn — it's just one step in the workflow you're already editing.

If you already use **Zapier** — the service that connects your apps so something happening in one automatically triggers something in another — there's a new way to make a Zap smarter. According to Zapier, the AI capabilities that used to live in a separate product called Zapier Agents have been folded directly into the regular Zap editor as a single **AI by Zapier** action. You drop it in like any other step, write a plain-language instruction, and pick which tools (small actions like "send an email" or "add a row to a spreadsheet") the AI is allowed to call on while doing its job.

This guide walks through that exact process, from opening an existing Zap to testing the AI step end-to-end. It's written for someone who has built a basic Zap before but has never added an AI step.

Before you start

- A **Zapier account** — a free plan works for basic Zaps, but the AI step's usage may count toward an AI actions limit depending on your plan, so check Zapier's current pricing page if you plan to run it heavily.
- An existing **Zap** (or the willingness to build a short two-step one to practice on). The AI step is just a third piece in that chain.
- Roughly **10–15 minutes** for the first run-through, including testing.
- The apps you want the AI to reach — for example Gmail, Slack, Google Sheets, or a calendar — already connected to your Zapier account.

Step 1 — Open the Zap you want to enhance

In your Zapier dashboard, find the Zap you want to add AI to and click to open it in the editor. You'll land on the visual workflow view, where each step is a card stacked vertically with arrows between them. If you've never built a Zap before, create a simple one first (a trigger plus one action) so you have something to work with. The trigger is the event that starts the Zap ("a new row in Google Sheets"), and each action is something the Zap does in another app. If your screen shows a list of Zaps instead of a workflow, click the name of the Zap you want to edit. **You'll know it worked when you see the numbered step cards and the path between them.**

Step 2 — Add the AI by Zapier action between your trigger and your final step

Click the **"+" between two existing steps** (or after your trigger if your Zap has no actions yet). A small panel opens with a list of apps. Type **"AI by Zapier"** into the search bar. When the result appears, click it. Zapier then asks which AI action you want to use — for a general-purpose step that can also call other apps, choose **"Use AI"** (the option that lets you write a custom prompt). The action is added as a new card in your workflow. If you don't find a "+" between steps, look for an **"Add step"** button below the trigger instead — Zapier's editor layout shifts occasionally. **You'll know it worked when a new card labeled "AI by Zapier" appears in your workflow, right between the cards that come before and after it.**

Step 3 — Write the prompt the AI will read each time it runs

Click the new **AI by Zapier** card to open its settings. You'll see a field labeled something like **"What should this AI do?"** or **"Prompt"** — the exact label can shift as the product changes, but it's the one big text area at the top of the card. Type the instruction in plain English, exactly as you'd explain the task to a person. Be specific about the input the AI will receive (usually data from your trigger step) and the output you want back. You can reference data from earlier steps by clicking a **field reference** button or similar icon next to the prompt box — it lets you insert variables from the trigger (like the email subject or the spreadsheet row) so the AI sees real data, not placeholders. **You'll know it worked when you can see at least one piece of data from a previous step referenced inside the text — usually shown as a small colored chip or a label like `1. Subject`. Example:** `Read the email below and tell me in one sentence whether the customer is happy, neutral, or upset. Email subject: {{1. Subject}} Email body: {{1. Body}}`

Step 4 — Tell the AI which tools it's allowed to use

Still inside the AI step, scroll to a section labeled something like **"Tools"** or **"Available actions"**. This is where you decide what the AI is allowed to do besides read the prompt. Click **"Add a tool"** (or a similarly worded button) and choose from the apps already connected to your account — for example "Gmail: Send Email", "Slack: Send Channel Message," "Google Sheets: Create Row," or "Calendar: Create Event." Add only the tools the AI actually needs for this task; more tools means more chances for the AI to choose the wrong one. Each tool you add opens a small configuration panel where you fill in the recipient, the channel, the spreadsheet, and so on — you can either set those values yourself in advance (a static "send to this address") or leave them for the AI to fill in from the prompt. If you don't see a Tools area, you may be on a simpler version of the step that only returns a text answer — in that case, skip to Step 5 and come back to tools later. **You'll know it worked when each tool you added appears as a sub-card inside the AI step, with its fields either pre-filled or marked for the AI to fill in.**

Step 5 — Test the AI step before turning the Zap on

Click the **"Test"** or **"Test action"** button on the AI card. It runs your prompt once with a real (or sample) piece of data from your trigger and shows you exactly what the AI wrote back, plus which tools it chose to call and what it sent to them. Read the output carefully — this is your chance to catch a prompt that's too loose or a tool the AI is misusing. If the answer is off, click back into the prompt, tighten the wording (for example, add "Reply in exactly one sentence" or "Only use the Slack tool when the customer is upset"), and test again. When the output looks right, click **"Continue"** or **"Done editing"** to save, then turn the Zap on with the toggle at the top of the editor. **You'll know it worked when the test panel shows a non-empty AI response and, if you added tools, a short log of which tool ran and what it did — and your Zap's status switch shows "On."**

Common mistakes

- **Writing a vague prompt.** "Summarize this email" gives you uneven results; "Reply in one sentence saying whether the customer is happy, neutral, or upset, and explain your answer in plain English" gives you something you can rely on.
- **Giving the AI too many tools.** If you add Gmail, Slack, Sheets, and Calendar all at once, the AI has to guess which one fits. Stick to the one or two tools this step actually needs.
- **Skipping the test step.** The test is the only place you'll see what the AI actually does before it runs on real data for real people. If you don't test, you're publishing blind — and AI mistakes (wrong recipients, wrong tone) are easy to miss without that preview.

What this means for you

- **In everyday life:** A Zap with an AI step can quietly sort your incoming email — flagging the urgent ones, drafting short replies, or pushing the important ones into a priority folder — so your inbox doesn't eat your morning.
- **For your business or work:** You can build Zaps that read new support tickets, summarize each one, and post a short digest into a team Slack channel, without anyone manually copying text from one app to another.
- **If you're just getting started:** The easiest first Zap with AI is one trigger you trust (like "new email in Gmail") plus the AI step with just a prompt and no extra tools — that lets you see the AI's plain-text answer in the Zap history and decide from there whether to give it tools.

Try it now

Open any Zap you already have, add an **AI by Zapier** action right after it, and write this prompt in the prompt box: `Read the file under the file upload here and let me know what it's about in 2 sentences.` Then run the test. You'll see exactly how the AI step talks, what its output looks like, and whether you want to give it tools — all in under two minutes.

How to Plan a Simple Weekly Workout Routine That Fits Your Real Life

5 steps · Read online

> **In short:** A working weekly plan is built around your real week, not an ideal one — pick 3-4 sessions, match them to the time and gear you actually have, and keep one rule: when life gets busy, shorten the session, never skip it. This guide walks you through that in five concrete steps.

A workout routine only works if you can actually do it on a normal week — the one with the late meeting, the kid who woke up sick, the grocery run that ran long. The plan below takes about 20 minutes to set up once, then lives in your notes app. It's intentionally boring, because boring plans get repeated.

Before you start

- A notebook, a notes app, or a piece of paper — anything you can carry or pin to a fridge
- A rough idea of which days of your week are usually free (most people pick 3-4)
- A list of what you have access to: bodyweight only, dumbbells, a gym, a park bench, a yoga mat
- About 20 minutes for the first setup, then 2 minutes a week to check in on Sunday
- No apps, no accounts, no paid plans — this is a thinking exercise, not a subscription
- This guide does not replace a doctor, physiotherapist, or certified trainer. If you have a health condition, an old injury, or you're returning after a long break, get a quick medical check first.

Step 1 — Pick your weekly "training windows"

Sit down and look at the next seven days as they actually are — not as you wish they were. Mark every half-hour block that is already spoken for: work, commute, school run, dinner, sleep. What's left is your real availability. Now pick 3 or 4 blocks of 30 to 45 minutes and write them down by day and time. Be honest: if Saturday morning is your only quiet window, that's your anchor session — work the rest around it.

Example: "Tuesday 7:00am, Thursday 7:00am, Saturday 8:30am, Sunday 5:00pm (backup)."

If it looks different: if your week is unpredictable (shift work, freelancing), pick a weekly *quota* instead — "4 sessions this week, in any order" — and place them on the day they become possible. You'll know it worked when you can point at the calendar and say "this is when I train" without re-thinking it.

Step 2 — Match each window to a session type

A balanced week usually mixes three things: something that makes your heart work (cardio, like brisk walking, cycling, or jumping rope), something that challenges your muscles (strength — push-ups, squats, dumbbells, resistance bands), and something gentle that keeps you moving without strain (mobility — stretching, yoga, a slow walk). Slot them so the harder days are spaced out. A common beginner rhythm is strength on day 1, cardio on day 2, mobility on day 3, strength again on day 4. Write the session type next to each window from Step 1.

Example: "Tuesday — strength 25 min • Thursday — cardio 30 min • Saturday — strength 30 min • Sunday — mobility 20 min."

If it looks different: if you only have three windows, drop one strength session and keep one cardio, one strength, one mobility. If you only have two, alternate weeks between cardio-focus and strength-focus. You'll know it worked when each window has a clear label and you're not staring at the calendar on Tuesday wondering "what was I supposed to do?"

Step 3 — Build a starter set of exercises (5 to 7 moves per session)

Pick exercises you can do with the gear you marked down in *Before you start*. For most beginners, a full-body strength session needs only a squat (or sit-to-stand from a chair), a push-up (against a wall, a table, or the floor — any version counts), a plank, a glute bridge, and a row (with a bag of books if you have no weights). For cardio, the move that works best is the one you'll actually repeat — walking, cycling, dancing in your kitchen, all count. For mobility, pick five stretches for the areas that feel tight after your day. Write the moves in a simple list, in the order you'll do them.

Example: "Strength — 1) squats 10×3, 2) push-ups 8×3, 3) plank 20 sec ×3, 4) glute bridge 12×3, 5) backpack row 10 each arm ×3."

If it looks different: if a move hurts or feels wrong in your body, swap it — there is almost always a gentler version. You'll know it worked when you can read your own list and picture yourself doing each move without needing a video.

Step 4 — Set a "minimum useful session" (the busy-week rule)

This is the most important step, because it decides whether your routine survives a hard week. For each session type, write down the smallest version that still counts as a workout: 10 minutes of brisk walking is a cardio session. 3 sets of 5 squats and 3 sets of 5 push-ups is a strength

session. 5 minutes of stretching is a mobility session. Write these minimums next to your main plan. The rule is simple: when life gets busy, you do the minimum, never skip the session entirely. A 10-minute walk is infinitely better than a 0-minute plan you feel guilty about.

Example: "Minimum strength = 3×5 squats + 3×5 push-ups. Minimum cardio = 10 min walk around the block. Minimum mobility = 5 min stretch before bed."

If it looks different: if even the minimum feels impossible that day, drop to 2 minutes of any movement and count it. You'll know it worked when you have, on paper, a "bad day" version of every session that takes less than 15 minutes total.

Step 5 — Do a 2-minute Sunday check-in

Every Sunday, look at the coming week. Move sessions if a meeting shifted, swap days if your kid has a recital, decide now which days are realistic. Write the new version in your notes. This is the difference between people who train for years and people who stop after three weeks — they plan the week before the week, not the morning of. The check-in also lets you notice patterns: if you keep missing Thursday mornings, Thursday morning is not your training window. Change the day, don't fight the day.

Example: "Sun check — moving Thursday to Friday because of the team offsite. Sat 8:30am still good."

If it looks different: if you forget the Sunday check-in, do it whenever you remember — Monday morning, Tuesday night. A check-in that happens late beats a perfect one that doesn't happen. You'll know it worked when you can say "here's my week" in one sentence, without scrolling through your phone.

Common mistakes

- **Mistake: copying a "perfect" 6-day plan from a fitness influencer.** Fix: cut it in half. Three sessions you actually do will beat six sessions you quit in three weeks. Add a day only after four consistent weeks.
- **Mistake: making every session long and intense.** Fix: shorter, easier sessions done often beat occasional heroic ones. If you finish a session thinking "I could have done more," the dose is right for building the habit.
- **Mistake: skipping a session because you missed the "scheduled" day.** Fix: do the minimum version that day, or shift the day by 24 hours. A routine that bends survives; a routine that breaks from one missed Tuesday is too brittle.
- **Mistake: changing exercises every week because something new looked fun.** Fix: keep the same 5-7 moves for at least 4 weeks so you can actually measure progress. Boredom, here, is a feature.

What this means for you

- **In everyday life:** you'll spend about 90-150 minutes a week training — less than two episodes of a TV show — and most of the rest of the week you'll feel slightly more capable walking up stairs, carrying groceries, and sleeping well.
- **For your business or work:** regular movement is one of the cheapest ways to protect focus and energy during a heavy week. A 20-minute walk between meetings is a known reset for the brain, not a luxury.
- **If you're just getting started:** ignore everything else in this guide except Steps 1 and 4. Pick three days. Write down a minimum version of each session. Start next week. Everything else is fine-tuning.

Try it now

Open your notes app right now, write the days and times of your next three training windows, and put the smallest possible session next to each one — even if it's just "10-minute walk." That's your week 1. Everything else can wait.

Organize Your Daily Schedule by Voice with Alexa+

6 steps · Read online

Ever wished you had a personal assistant to keep track of your daily tasks and appointments? This guide will show you how to use your Alexa+ enabled device to set reminders and manage your calendar, all with the power of your voice. By the end, you'll be able to add and adjust your schedule effortlessly, making sure you never miss a beat.

Before you start

- An **Alexa-enabled device** (like an Echo speaker or a smart display) that's already set up and connected to Wi-Fi.
- The **Alexa app** installed on your smartphone or tablet, logged in with your Amazon account.
- A **free Amazon account**.
- A **digital calendar service** (like Google Calendar, Apple Calendar, or Outlook Calendar) if you plan to manage calendar events.
- Rough total time: 10-15 minutes.

Step 1 — Check Your Alexa+ Device is Ready

To begin, make sure your Alexa+ device is connected and ready to respond. This foundational step ensures Alexa+ can hear you and interact with your personal schedule. Your Alexa+ device, whether it's a speaker or a smart display, acts as your voice-controlled assistant, and it needs to be online and linked to your Amazon account to work.

You'll usually see your Alexa+ device with a blue light or a display screen showing a clock, indicating it's powered on and connected. Once you speak the "wake word" (the phrase that gets Alexa+'s attention, like saying "Alexa"), the light will illuminate or the screen will animate, showing it's listening. If you don't see any lights or the device doesn't respond, try unplugging it for 30 seconds and plugging it back in, or check its Wi-Fi connection in the Alexa app under "Devices."

You'll know it worked when **you can say "Alexa, what time is it?" and your device responds with the current time.**

Step 2 — Link Your Digital Calendar

To manage calendar events using your voice, Alexa+ needs permission to connect to your existing digital calendar. This is like giving your personal assistant access to your diary so they can add or check appointments for you. By linking your calendar, any events you create or modify through Alexa+ will sync automatically with your chosen calendar service.

Open the **Alexa app** on your smartphone or tablet. Look for the "More" button, often represented by three horizontal lines or dots, usually located in the bottom right corner of the screen. Tap it to open a menu. From the menu, find and tap "Settings," then scroll down and select "Calendar." You'll see options to link various popular calendar services. Tap "Add Account," choose your preferred calendar (e.g., Google, Apple, Outlook), and follow the on-screen prompts to sign in and grant permission.

You'll know it worked when **your chosen calendar service is listed under "Linked Calendars" in the Alexa app, and you see a confirmation message in the app that your account is connected.**

Step 3 — Create a Voice Reminder

Now that Alexa+ is ready, let's set your first voice reminder. A reminder is a simple alert that Alexa+ will speak at a specific time, perfect for remembering everyday tasks like taking out the rubbish or calling a friend. It's like having a sticky note that talks to you when it's time to act.

Approach your Alexa+ device and clearly say your "wake word" followed by your instruction. Keep your voice natural and clear. Alexa+ will process your request, using its artificial intelligence to understand what you're asking. Once it understands, it will confirm the reminder details with you. If it sounds different, try rephrasing your request.

Example: "Alexa, remind me to water the plants at 3 PM today."

You'll know it worked when **Alexa+ says, "Okay, I'll remind you to water the plants at 3 PM today," or a similar confirmation.**

Step 4 — Add a Calendar Event by Voice

Moving beyond simple reminders, you can also add more detailed events directly to your digital calendar using just your voice. This is great for scheduling appointments, meetings, or social gatherings. Since you've already linked your calendar, Alexa+ can place these events precisely where they belong in your schedule.

Start by saying your "wake word," then clearly state the event you want to add, including the name, date, and time. Alexa+ will listen and interpret your request, then use the linked calendar from Step 2 to add the event. Sometimes it might ask for more details, like which calendar to use if you have multiple linked, or to confirm the event name. If it seems confused, try giving more specific details in your instruction.

Example: "Alexa, add 'dentist appointment' to my calendar for next Tuesday at 10:30 AM."

You'll know it worked when **Alexa+ says, "Okay, I've added 'dentist appointment' to your calendar for next Tuesday at 10:30 AM," and the event appears in your digital calendar app on your phone.**

Step 5 — Check Your Upcoming Schedule

With reminders and events set, it's helpful to quickly review what's coming up. Alexa+ can easily read out your schedule, acting like your helpful secretary giving you a quick overview of your day or week. This saves you from having to look at your phone or computer every time you want to know what's next.

To check your schedule, simply use your "wake word" followed by a question about your reminders or calendar. Alexa+ will access the information it has stored (for reminders) and your linked calendar (for events) and read the details aloud. If it lists things you weren't expecting, it might be accessing a different linked calendar than you intended, or picking up old reminders.

Example: "Alexa, what's on my calendar for tomorrow?" or "Alexa, what are my reminders?"

You'll know it worked when **Alexa+ lists your upcoming events or reminders, giving you a clear audio summary of your schedule.**

Step 6 — Manage (Edit or Delete) Entries by Voice

Life changes, and so does your schedule. Thankfully, you can easily modify or remove reminders and calendar events using your voice, giving you flexible control without needing to type anything. This ability to make quick adjustments ensures your digital assistant remains helpful and your schedule accurate.

To manage an entry, use your "wake word" and clearly state what you want to change or remove. For instance, to change a reminder's time, you might say, "Alexa, change my 'water the plants' reminder to 4 PM." For calendar events, you'll need to be specific about the event you're referring to. If Alexa+ has trouble finding the right entry, try being more exact with the event's name, date, or time. For very complex changes, it might be easier to use your calendar app directly.

Example: "Alexa, cancel my 'dentist appointment' for next Tuesday." or "Alexa, delete the 'take out the bins' reminder."

You'll know it worked when **Alexa+ confirms the change or deletion, such as "Okay, I've canceled 'dentist appointment' for next Tuesday," and the entry is updated in your calendar or reminder list.**

Common mistakes

- **Not speaking clearly or accurately:** Alexa+, like any digital assistant, needs clear instructions. If you mumble or use unusual phrasing, it might misunderstand.
  - **The fix:** Speak clearly and at a normal pace. If it misunderstands, rephrase your request using slightly different words.
- **Forgetting the wake word:** You must start every command with "Alexa" (or your chosen wake word) to activate the device. Without it, your device won't listen.
  - **The fix:** Always begin your command with your designated wake word. Think of it like calling someone's name before you speak to them.
- **Calendar events not syncing (or not being created):** This usually happens if your digital calendar account wasn't linked properly, or if Alexa+ isn't set to use the correct calendar.
  - **The fix:** Go back to the Alexa app settings (as in Step 2) and verify that your preferred calendar account is linked and selected as the default.

What this means for you

- **In everyday life:** Imagine cooking dinner and suddenly remembering you need to call your mom. A quick "Alexa, remind me to call Mom at 7 PM" saves you from messy hands and forgotten tasks.
- **For your business or work:** You can quickly add a meeting to your calendar or set a deadline reminder without interrupting your workflow, keeping your focus on the task at hand.
- **If you're just getting started:** Begin with a simple reminder for something you do daily, like "Alexa, remind me to stretch at 8 AM every morning." This helps build familiarity.

Try it now

Grab your Alexa+ device and set a reminder for yourself in the next hour. For example, say: "Alexa, remind me to take a break in 30 minutes."

Set Up an Alexa+ Routine to Automate Your Everyday Tasks

6 steps · Read online

By the end of this guide you'll have a working Alexa+ Routine that starts your morning (or any other moment) with a spoken phrase, turning on lights, playing music and delivering a custom greeting. It's perfect for anyone who wants a little extra convenience without needing any coding skills.

Before you start

- An **Echo speaker**, **Echo Show**, or any **Alexa-enabled smart display** plugged in and connected to your home Wi-Fi.
- The **Alexa app** installed on a smartphone or tablet and signed in with your Amazon account.
- Any **smart-home devices** (e.g., smart bulbs or plugs) already added to the Alexa app.
- About **5–10 minutes** of quiet time.

Step 1 — Open the Alexa app

Tap the Alexa app icon on your phone or tablet. When the app opens you'll see a home screen that shows a picture of your device at the top and a row of icons across the bottom (Home, Devices, Routines, etc.). Tap the **Routines** icon – it often looks like a clock or a series of stacked squares and may be labeled "Routines". The app will slide to a new page listing any existing routines. If you don't see a Routines icon, look for a three-line "Menu" button in a corner and choose "Routines" from the list. You'll know it worked when a screen appears titled "Your Routines" with either a list of routines or a note that none exist yet.

Step 2 — Start creating a new routine

On the "Your Routines" screen, locate a **plus (+) button** in the top-right corner (or a button that says "Create Routine"). Tap it. You'll now see a blank page that asks you to "Add a name" and "When this happens". If the plus sign isn't visible, look for a colorful "New Routine" button near the center of the screen – the wording can vary between devices. You'll know it worked when the top of the page shows a text field for the routine name and a section titled "When this happens".

Step 3 — Name your routine

Tap the text field that says "Add a name". A keyboard slides up and you can type any label you'll recognize later – for example **Morning Wake-Up**. (If you're on a tablet, the keyboard may appear at the bottom of the screen.) After typing, look for a **Next** or **Continue** button – usually at the bottom right – and tap it. You'll know it worked when the view switches to a new screen that says "When this happens" and shows a button labeled "Add trigger".

Example: Type `Morning Wake-Up` then tap **Next**.

Step 4 — Choose a voice trigger

Press the **Add trigger** button. A menu of trigger types appears: *Voice*, *Schedule*, *Smart Home* and *Device*. Select **Voice** – this tells Alexa to listen for a spoken phrase. A field labeled "Enter phrase" opens; type the exact words you want to say, such as **Alexa, good morning**. Tap **Next** or **Add** (the button may be at the bottom of the screen). If you don't see a Voice option, look for any item that mentions "spoken phrase" or a microphone icon. You'll know it worked when the top of the routine page now reads "When you say 'Alexa, good morning'".

Example: Type `Alexa, good morning` then tap **Add**.

Step 5 — Add an "Alexa Says" action

Below the trigger section, tap **Add action**. A list of action categories appears – "Alexa Says", "Music & Podcasts", "Smart Home", etc. Choose **Alexa Says**; you may need to scroll a little. Now pick **Customized** (or "Speak text") – this opens a text box where you can type what Alexa should say. Enter a friendly greeting like **Good morning! Time to start the day.** Tap **Next** or **Add** to save the action. If the category is named "Speech" or "Announcement" instead of "Alexa Says", select the one that lets you type a spoken line. You'll know it worked when the routine summary shows a line reading "Alexa will say 'Good morning! Time to start the day.'".

Example: Type `Good morning! Time to start the day.` then tap **Add**.

Step 6 — Add a music action (optional but fun)

Press **Add action** again, which now appears under the previous action. Choose **Music & Podcasts** from the list. If you have a music service linked (Spotify, Amazon Music, etc.) you'll see its logo – tap the service you prefer. A new field lets you type an artist, genre or playlist name; for a breezy start type **Acoustic Sunrise** (or any favorite playlist). Tap **Next** or **Add**. If no music service shows, you'll need to link one in the Alexa app's Settings first – the guide assumes a service is already connected. You'll know it worked when the routine page lists a second action: "Play 'Acoustic Sunrise' on Spotify".

Example: Type `Acoustic Sunrise` then tap **Add**.

Common mistakes

- **Skipping the "Add trigger" step:** Without a trigger Alexa won't know when to run the routine. If you miss it, go back to the top of the routine page and tap "Add trigger" again.
- **Typing the wrong phrase:** Alexa matches the spoken phrase exactly, so a typo will prevent activation. Double-check the text in the "Enter phrase" field and re-type if needed.
- **Forgetting to link a music service:** The "Music & Podcasts" action will appear blank if no service is connected. Open the Alexa app's Settings → Music, add your favorite service, then return to the routine.

What this means for you

- **In everyday life:** You can wake up to a personalized greeting, have the lights turn on and your favorite playlist start without lifting a finger – a smooth start to a busy day.
- **For your business or work:** If you run a home-based office, a routine can flip on a desk lamp, start a focus-music station and announce the day's agenda, helping you shift into work mode quicker.
- **If you're just getting started:** The easiest first move is to create a voice-triggered routine that only says a greeting. Once you're comfortable, add more actions like smart-home controls or music.

Try it now

Open the Alexa app, tap **Routines**, create a new routine called **Quick Test**, set the voice trigger to **Alexa, test routine**, and add a single "Alexa Says" action that says **Hello, I'm ready!**. You'll have a working routine in under two minutes.

How to Add an Alibaba Cloud AI Chatbot to Your WordPress Site

5 steps · Read online

> **In short:** An Alibaba Cloud chatbot can handle visitor queries on your site 24/7, answering common questions like opening hours while you focus on other work.

By the end of this guide you'll have a chat bubble that greets visitors, answers basic queries, and works on every page of your WordPress site – no programming experience required.

Before you start

- **Alibaba Cloud account:** sign-up at alibabacloud.com (the free tier is enough for a simple chatbot).
- **WordPress site:** you need administrator access to the dashboard.
- **A web browser** on a computer or tablet (Chrome, Edge, or Firefox work fine).
- **A few minutes:** the whole process typically takes about 20-30 minutes.
- **Optional:** a free plugin such as "Insert Headers and Footers" makes adding code easier.

Step 1 — Create a new chatbot instance

Open your web browser, go to console.alibabacloud.com and sign in. Look for the AI chatbot or conversation AI service in the console – you can usually find it by searching for "Chatbot" or "Conversation AI" in the top-right search bar. Click the button to create a new bot or robot (it may be labeled **Create Robot**, **Create Instance**, **New Bot**, or similar). A form appears asking for a name and a brief description. After you type a name such as "FrontDeskHelper" and confirm, the platform creates a new bot and shows it in your list, typically with a status of "Running" or an indicator showing it's active. 💬 Example: `FrontDeskHelper`

You'll know it worked when the new entry appears in your bot list with an active status indicator.

Step 2 — Define the chatbot's welcome greeting

In the list of bots, click the name of the bot you just created to open its settings page. Look for a section where you can set the first message visitors see – this might be labeled **Greetings**, **Welcome Message**, **Chat Prompt**, or something similar. There's usually a text box where you can type what the bot says when someone first opens the chat. After you type a sentence like "Hi! I'm your digital assistant – ask me about opening hours or services," a preview of the message often appears below the box. Click **Save** to store the change. 💬 Example: `Hi! I'm your digital assistant – ask me about opening hours or services.`

You'll know it worked when you see a confirmation that the settings were saved successfully.

Step 3 — Choose a visual style that matches your site

While still in the bot's settings, look for options to customize how the chat bubble looks – this might be in a tab labeled **Appearance**, **Style**, **Customization**, **Theme**, or similar. You'll typically find options to pick a color (you can copy your site's primary color as a hex code, e.g., #0066CC) and possibly upload a small logo or icon for the chat bubble. The preview on the screen should update to show what the chat bubble will look like on your site. 💬 Example: set color to `#0066CC` and upload `logo.png`

You'll know it worked when the on-screen preview shows the color and icon you selected.

Step 4 — Copy the widget embed code

Return to the main bot page and look for a button to publish, deploy, or make the bot publicly available – it may be labeled **Publish**, **Deploy**, **Enable**, or similar. Once activated, a panel should appear (it might be labeled **Embed Code**, **Web Widget**, **Integration Code**, or something similar) containing a snippet of code that starts with `<script src="..."></script>`. Look for a **Copy** button next to the code box; if there isn't one, click inside the box and press Ctrl + C (or Cmd + C on a Mac) to copy it manually. 💬 Example: `<script src="https://chatbot.alibabacloud.com/widget.js"></script>` — *your actual code may differ*

You'll know it worked when you can paste the snippet into a plain-text document and see a full `<script>` tag appear.

Step 5 — Insert the code into WordPress footer

Log into your WordPress dashboard (your-site.example.com/wp-admin). You can add the chatbot code either through **Appearance → Widgets** (look for a widget area like **Footer** or **Footer 1**, add a **Custom HTML** block, and paste the code) OR by installing a free plugin like "Insert Headers and Footers" and going to **Settings → Insert Headers and Footers**, then pasting the code into the "Scripts in Footer" box. Either method makes the script load on every page of your site. Press **Update** or **Save** to confirm. 💬 Example: paste the `<script>` snippet exactly as copied

You'll know it worked when you open your website in a new browser tab (or an incognito window) and see a friendly chat bubble sitting at the bottom right of the page.

Common mistakes

- **Skipping the Save step in the greeting settings:** the message may revert to a default. Fix it by always clicking the **Save** button and watching for a confirmation.
- **Placing the widget code in the header instead of the footer:** the chat may load too early and cause page-load delays. Fix it by using the Footer widget area or the "Scripts in Footer" field of a plugin.
- **Not testing on a mobile device:** the bubble can cover navigation icons on small screens. Fix it by opening your site on a phone, checking the bubble's position, and adjusting the placement or size options in the Appearance or Style settings if available.

What this means for you

- **In everyday life:** Your chatbot can answer simple questions like "What are your opening hours?" so you can relax while it handles routine enquiries.
- **For your business or work:** The bot captures leads overnight, giving you contact details even when your staff are offline, which can increase bookings without extra staffing.
- **If you're just getting started:** You don't need any code – just copy, paste, and a few clicks are enough to make a professional-looking assistant appear on every page.

Try it now

Open your WordPress dashboard, go to **Appearance → Widgets**, add a **Custom HTML** block to the Footer, paste any short piece of text (like "Hello world") and click **Update**. In the next two minutes you'll see that the widget saves correctly – a quick win that shows you're ready for the chatbot code.

Every prompt and example in this pack

Every line marked 💬 in the guides above, collected in one place: the prompts you can copy and type (marked ➤), and the example labels and screen text that go with them. Each guide title links back to the guide it came from.

731 examples from 147 guides — 280 of them prompts you can type

How to check if a suspicious message is a scam by pasting it into an AI

> open ChatGPT, Gemini, Claude, Microsoft Copilot, or Le Chat — whichever one you already have an account on.

> turn "my SSN is 123-45-6789, please call 555-0123" into "my SSN is [REDACTED], please call [REDACTED]".

> ➤ "Here's a message I just received. Can you act as a security expert and tell me whether it looks like a scam? List the specific red flags you see, and tell me whether I should click any links in it. Message: [paste here]"

> if the AI says "the sender domain doesn't match the company", check that yourself — does the "from" address really end in the company's real website, or in something like "support-service.example.com"?

> a "USPS" tracking link points to `usps-tracking.example.com/pay-now` — that is NOT usps.com, so it's a fake.

> text a family member — "Hey, I just got a weird message about my Amazon account, does this sound legit to you? [paste]"

How to Spot an AI Scam Call, Fake Email, or Deepfake Voice Today

> ➤ "Grandma, it's me, I was in an accident, please send money" → hang up → dial your grandchild's saved number.

> ➤ "Hi, this is Maria from your bank's fraud department" → hang up → look up the bank's number on your statement → dial it yourself.

> ➤ "Your package is held at customs, pay a small fee to release it" → hover the link → see it points to a domain you've never heard of → delete.

> ➤ "Pay this fee now or your account closes in five minutes" → pause → ask the ten-minute question → almost always the answer is no.

> Text your sister: "If anyone calls saying they're me and needs money, ask them for our safe word. The word is platypus."

> Realized mid-call that the "bank agent" was fake → hung up → called the number on my debit card → told them what happened → account flagged for monitoring.

What You Should Never Share With an AI Assistant: A Privacy Guide

> ➤ “What types of personal information should I never share with you?” You’ll know it worked when the assistant lists three groups such as (1) identifiers (name, address, ID numbers), (2) financial details (bank accounts, credit-card numbers), and (3) health or biometric data.

> type “Show me my last five messages” (many assistants support a simple request like this). You’ll know it worked when the assistant displays the recent prompts you entered, giving you a chance to review them.

> after selecting the message, type “Delete this message”. You’ll know it worked when the message disappears from the chat view and the history no longer shows that content.

> ➤ “Turn off chat history”. You’ll know it worked when the toggle changes color or moves to the “off” position, and the app confirms “Your chats will not be saved”.

> ➤ “Show my recent chats”. You’ll know it worked when the list only contains the recent, innocuous question you just asked.

> ➤ “Dear Support, please remove any personal identifiers from my account.” You’ll know it worked when you receive a confirmation email or a system message saying “Your request has been received”.

How to use AI safely with your personal information

> Instead of typing: *"My name is Jane Smith, I live at 24 Smith Street and need help budgeting my $1,500 weekly salary. My email is jane.smith@example.com."*

> In ChatGPT, open Settings (tap your profile picture or the gear icon), then look for "Privacy" or "Data Controls" and toggle off any option related to using your conversations for model improvement.

> If you want the AI to polish a professional email, change the real names and figures to generic labels like this:
> Original: *"Hi Sarah, can you please transfer the $450 refund for the project at 12 High St to my Commonwealth account by Friday?"*

> In Claude, look for a "Temporary Chat" button next to "New Chat" in the sidebar. Click it to start a new, unsaved conversation.

How to Spot a Phishing Email or Scam Text Before You Click

> if it's "Your package is delayed, confirm here: bit.ly/x9zq", just read the words. Don't tap the link.

> ➤ "Act as a fraud detector. Read the message I'm pasting next and list every red flag you can see, then tell me clearly whether it's safe to reply, click any link, or open any attachment. If you're not sure, say so."

> of what to paste: "From: Amazon Support <> — Dear customer, your account has been suspended due to unusual activity. Click here to verify: hhtp://amazon-secure.example.com/login — please respond within 24 hours."

> of an AI answer: "Red flags: sender domain is amazon-support.example.com, not amazon.com; the link is 'hhtp' (typo) and uses a zero in 'amazon'; urgency tactic ('within 24 hours'); generic greeting. Verdict: do NOT click, do NOT reply. Report and delete."

> of a good follow-up question to the AI: "Can you explain in one sentence what an 'urgency tactic' means, so I know what to look for next time?"

How to Spot AI Scams, Deepfakes, and Fake Content

> You see a photo of a well-known person in an unusual situation. Zoom in on their hands and face. If you notice anything unusual — extra fingers, distorted features, or inconsistent lighting — investigate further before sharing or believing it.

> You receive an urgent video message from your company's director asking for an immediate fund transfer. Pause the video and look closely at their jawline, neck, and hairline. If you notice anything unusual — a fuzzy outline, unusual shadows, or lip movements that don't quite match the audio — treat it as suspicious and verify through another channel before acting.

> You receive an urgent phone call from a friend asking for financial help. The voice sounds exactly like them, but the request is unusual or the tone feels strange in context. This is a red flag. Hang up immediately and call them back on their usual, trusted number to verify.

> A supplier sends an email with "updated bank details" for a payment, urging you to use the new account. Do not reply to that email or call the phone number listed in its signature. Instead, find the phone number from your original paper contract or the supplier's official website and call them directly to confirm the change.

> You receive an email from a supposed friend, but it uses overly formal language like "Furthermore, I require immediate assistance" rather than their usual casual style, and every sentence is perfectly structured without any contractions or slang. This unnatural perfection suggests AI involvement.

How to Spot AI-Generated Scam Messages Before You Click or Reply

> a message you might want to check — *"Hi Mom, I dropped my phone and I'm using a friend's number. Can you urgently transfer $450 to this account for a new one? I'll pay you back tonight."*

> open the assistant you already use, and start a brand new chat titled something like *"Check this message for scam signs"*.

> prompt to copy: *"I got this message and something feels off. Can you list any scam warning signs you notice? Please look at the tone, the request, any pressure to act fast, and anything that doesn't quite add up. Don't tell me what to do — I just want a list of red flags."*

> ➤ *"The sender claims to be from USPS. Their email is 'usps-delivery-support@gmail.com' and the link is 'usps-tracking.example.com/pay-now'. Does that look like a real USPS address? Just point out anything suspicious — don't visit the link for me."*

> ask a follow-up like *"Based on the red flags you found, is this more likely a scam or a real message? Give me your best guess and a one-sentence reason."*

How to Spot a Deepfake Video or Voice Clone Before You Share It

> ➤ "Hey, I got a voice note from someone claiming to be you asking about gift cards. Did you actually send it? Just checking."

Build a Simple Personal Budget Spreadsheet Using Any AI Assistant

> ➤ "Design a simple monthly budget spreadsheet for a single person earning a regular salary. Include columns for Category (Rent, Groceries, Transport, Utilities, Entertainment, Savings, Other), Planned amount, Actual amount, and Difference. Add 2 example rows per category. Keep it under 20 rows total."

Decode Any Confusing Phone, Internet, or Medical Bill in 10 Minutes

> a small note on your paper — "Today's goal: figure out what each line means and which ones look off."

> keep visible — "$84.50 — Base Plan — Sept 3 to Oct 2". Black out everything else.

> prompt to copy: "Here are the line items from my phone bill. For each charge, tell me: (1) what it usually means, (2) whether it's normal or worth questioning, (3) one question I should ask the provider."

> mark "monthly plan $40" green, "regulatory fee $3.20" yellow (you don't know if it's accurate), and "international roaming $27" red (you didn't travel).

> search: type `"administrative fee"` plus the company's name into your browser to see how they explain the charge in their own words.

> message you can adapt: "Hi, on bill dated 3 October, line 'Equipment Fee — $12.00' appeared for the first time. I don't recall being notified about a new charge. Can you explain what this covers and the start date? Thank you."

> closing reply — "Thanks. Could you confirm the credit amount in writing and the date it will appear on my next statement?"

How to read and summarize a long lease or insurance policy with Gemini

> *If you don't see a chat box, look for a button labeled **"New chat"** or a **pencil icon** somewhere on the page.*

> *If the icon names don't match, try the **"+" or plus button** — most chat interfaces tuck uploads there.*

> ➤ *"Please summarize this lease in plain English. I want to know what I have to pay each month, what I'm not allowed to do, what happens if I break the rules, and how I can end the lease."*

> ➤ *"Which clauses in this document could cost me the most money or cause the biggest problems if something goes wrong? List them and quote the key sentence from each one if you can."*

> ➤ *"What does 'indemnify and hold harmless' mean in everyday language? And does this clause apply to me as the renter, or only to the landlord?"*

> *If you'd rather skip the copy-paste, take a screenshot of the summary and save it to your phone. It's not searchable later, but it's better than nothing.*

How to Use ChatGPT to Explain a Confusing Bill or Invoice

> Take one photo straight down, then pinch to zoom and check the smallest line — if you can read it, ChatGPT probably can too.

> ➤ "This is my home internet bill. Please list every charge on it in plain English, and flag anything that looks unusual or higher than I'd expect for a typical plan."

> ➤ "What does 'regulatory cost recovery fee' usually cover on these bills, and is it something I can ask the company to remove?"

How to Decode a Confusing Medical Bill with AI

> ➤ "Scan the front and back of the bill into one PDF named 'bill-jan2026.pdf'."

> ➤ "Black out my SSN, member ID, and the last four digits of my card before pasting."

> ➤ "Here's my itemized bill from a clinic visit. In plain English, please list every charge, what each one was for, what my insurance paid, and what I still owe. Flag anything that looks unusual."

> ➤ "On line 4 it says 'CPT 99213, allowed amount $84.20, patient owes $25'. What is a 99213, what does 'allowed amount' mean, and why is the patient share $25?"

> ➤ "Please re-add every line on this bill and confirm the total. I'm getting $342.10 when I add it by hand but your last answer said $298.50."

How to Read Through Insurance Policies and Bank Contracts with Claude

> there is nothing to type yet — just open the chat.

> a PDF named `home-insurance-2026.pdf` shows up as a thumbnail before you send.

> ➤ "I'm a homeowner, not a lawyer. Please summarize this home insurance policy in plain English. Cover: what am I covered for, what am I NOT covered for, my deductible, and anything that could void my coverage. Keep it under 400 words."

> ➤ "Now read it again like a skeptical friend. List every clause that could limit my payout or surprise me. For each one, quote the original wording, then explain what it actually means in plain English."

> ➤ "I've pasted two home loan offers below. Compare them on: total interest paid over 5 years, early-repayment penalties, fixed vs variable rate, and any hidden fees. Show your work in a simple table."

> paste into a note titled `home-insurance-summary-2026` so future-you can find it.

Build a Smart Packing List for Any Trip in Just a Few Minutes

> ➤ "Lisbon, Portugal, 6 days in mid-October. Mild days around 20°C, cooler nights. Hiking one day, two nice dinners out, sightseeing the rest."

> (paste this, then edit the city and dates to match your trip): "I'm going to Lisbon, Portugal for 6 days in mid-October. Days will be mild around 20°C and nights cooler. I'll hike one day, have two nice dinners out, and sightsee the rest."

> ➤ "Please turn that into a packing list organized by category (clothes, toiletries, electronics, documents, activity-specific). Focus on layers for mild days and cooler nights, and call out anything I'm likely to forget."

> ➤ "Add a compact rain jacket, my prescription medications, and a USB-C power adapter. Remove the gym clothes and the formal shoes — I won't need them."

> ➤ "Now add a short 'often forgotten' list tuned to a 6-day city-and-hike trip in mild autumn weather."

> ➤ "Please rewrite the full list as a checklist I can copy and paste — one item per line with empty checkboxes."

> to save: "I'm going to [city] for [X days] in [month]. Weather: [warm/cold/rainy/unknown]. Activities: [list]. Please make a packing list by category for mild days and cooler nights, then add an 'often forgotten' section, then format as a checklist."

Plan a Week's Meals from What's Already in Your Fridge

> pull out the cheese drawer, the veg crisper, and the little door shelves too. A 5-minute job.

> a half-empty milk carton with a 2-day use-by date goes in "use first". A sealed jar of pickles goes in "use later".

> ➤ "We have basmati rice, two tins of tomatoes, a bag of frozen peas, and half a loaf of bread in the freezer."

> week: Monday stir-fry (uses the soft capsicum and half the chicken), Wednesday frittata (uses the eggs and any leftover veg), Friday curry (uses the sad carrots and a tin of lentils), with toast, yoghurt, and leftovers filling the gaps.

> ➤ "Tues lunch: last night's chicken + salad leaves. Sat dinner: pantry fried rice (rice, egg, frozen peas, soy sauce)."

> list: "1 loaf bread, 1 L milk, 1 lettuce, 6 tortillas." That's it. Estimated cost: about $12.

> while the curry is simmering, boil an extra cup of rice. Tomorrow's lunch is sorted in 30 seconds.

Plan a Family Weekend Getaway with DeepSeek: A Step-by-Step Guide

> just type the word **"Hi"** for now to make sure the box works.

> ➤ *"We're a family of four — two adults, a 7-year-old and a 10-year-old — based in Chicago. We want to drive somewhere within 3 hours for the weekend of 14–15 June. We love outdoorsy stuff and kid-friendly museums. Budget is medium."*

> ➤ *"Please make a day-by-day itinerary for Saturday and Sunday. For each day, give me morning, midday and afternoon activities that work for kids aged 7 and 10. Add rough drive times between stops."*

> ➤ *"Add lunch and dinner suggestions for each day — nothing too fancy, kids need to see familiar food on the menu. Estimate the total daily cost for our family of four, and give me a backup indoor plan in case it rains."*

> ➤ *"Swap the Sunday morning museum for an outdoor activity — the kids get restless indoors. Also, can you suggest a town within 2 hours of Chicago instead of 3?"*

> copy the final plan, then paste it into a new note titled **"Family weekend — 14/15 June"**.

> type *"Quick check — can you remind me of the backup plan for Sunday if it rains? I want it in one short list."*

Create a soothing bedtime story routine with Alexa

> Tap `More` → `Routines`.

> Type `Bedtime Story` into the name field.

> Type `start bedtime story`.

> Tap `Skills` → search for `bedtime stories` → select a skill like `Kids Bedtime Stories` and tap `Add`.

> Tap `Alexa's Voice` and choose `Matthew`.

> Tap `Music & Audio` → search for `rain sounds` → select a track and tap `Add`.

> Say: `Alexa, start bedtime story`.

How to Plan a Kids' Birthday Party Without Losing Your Mind

> ➤ "Saturday, June 14, 2:00–4:00pm. Guests: 8 kids + 2 siblings."

> ➤ "Theme: dinosaur dig. Budget total $200. Food $80, decor + activities $90, surprises $30."

> ➤ "4 weeks: send invites. 2 weeks: buy decor + cake ingredients. 1 week: confirm RSVPs. Day before: set up."

> ➤ "Hi! Maya's turning 6, dinosaur-dig party, Saturday June 14, 2–4pm at our place. Please RSVP by June 7."

> ➤ "Food: mini sandwiches + grapes + juice + cake. Anchor: backyard dinosaur dig with toy fossils."

> ➤ "1:00 setup. 2:00 arrive. 2:15 snack. 2:45 dig. 3:30 cake. 3:50 bags. 4:00 done."

How to use Qwen Chat to plan a weekend road trip in the US without getting lost

> `Hi`

> `I want a weekend road trip from Portland to the Pacific Coast Highway with my partner, leaving Friday after work and returning Sunday evening. Can you plan every stop, time, and backup idea so we don’t get lost?`

> `Here’s what I’d like to keep: the Twelve Apostles at sunset, a short rainforest walk at Maits Rest, and a seafood dinner in Apollo Bay. Please add those and any other must-sees, plus backup ideas if a stop is closed or crowded.`

> `Please give me turn-by-turn directions for every driving leg, including street names and kilometres, so my partner can navigate while I drive.`

> `Add a backup plan for every main stop in case it’s raining or too crowded. Include indoor alternatives and shorter walks.`

> `Please give me one clean summary I can save or print, with every stop, time, address, and backup idea in one place.`

> `Please review the entire plan for anything I might have missed—toll roads, fuel stops, or time buffers—and suggest any tweaks.`

Turn messy phone photos into a family album with Google Photos AI

> Open the Google Photos app → tap **Sign in** → type `your.name@gmail.com` and your password → tap **Next**.

> Tap **Search** → scroll to **People & pets** → tap the heading → wait for the list to appear.

> Leave the app open or close it — either way, come back in 2–3 minutes.

> Tap **Unknown 1** → tap the pencil icon → type `Dad` → tap **Save**.

> Tap **Library** → **Utilities** → **Album** → name it `Family Summer 2025` → tap **Add photos** → choose **From your groups** → select “Mom”, “Dad”, “Liam” → tap **Done**.

> Tap **Share** → choose **Messages** → pick “Mom” from your contacts → tap the send arrow.

> Open Google Photos → tap **Search** → tap **People** → name any new “Unknown” groups → close the app.

Organize your favorite recipes with Perplexity Q&A

> Type: `My Recipe Book`

> Pumpkin Soup
> Ingredients: 1 kg pumpkin, 1 onion, 2 carrots, 400ml coconut milk, 1 tsp cumin
> Instructions: Peel and chop veggies. Sauté onion. Add pumpkin, spices, and water. Simmer 20 mins. Blend. Stir in coconut milk.

> Save this as a recipe in My Recipe Book.

> Show me all recipes with coconut milk.

> What can I cook tonight with pumpkin and coconut milk?

> https://perplexity.ai/space/12345-my-recipe-book

> What’s a quick vegan dinner I can make?

Create a personalized weekly meal plan based on your diet with Gemini

> `I need a 7-day vegetarian meal plan for one person. I prefer easy recipes that take less than 30 minutes to prepare each night. I have lots of lentils, spinach, and sweet potatoes that I want to use.`

> `This looks great, but could you swap out the dinner on Tuesday for something with chickpeas instead of lentils? Also, can you add a simple snack idea for each day, like a piece of fruit or some nuts?`

> `Can you create a shopping list for this entire week's meal plan? Please group the items by type, such as 'Produce,' 'Pantry,' 'Dairy,' and 'Frozen.'`

> `For the 'Sweet Potato and Lentil Curry' on Thursday, could you give me a very quick overview of the main steps? Also, what's a good alternative if I don't have fresh ginger on hand?`

How to write heartfelt birthday cards, wedding toasts and thank-you notes with Copilot

> type `Help me write a birthday card for my sister.`

> `Write a warm birthday card for my dad, who is turning 70. He's a quiet, kind man who taught me to fish and always fixes things around the house. I want it to feel grateful and a little funny, not too long. About 80 words.`

> `That is good, but can you make it sound less formal? Add a line about the time he fell asleep in his fishing chair and the fish got away. Keep it under 100 words.`

> `Can you add a line about how he still hums the same tune from the radio while he works in the garage? It's something I will always remember.`

> change *"I am eternally grateful for your unwavering guidance"* to *"Thanks for always being there, Dad. I notice it more every year."*

> paste the draft into a blank card, print it on nice paper, or copy it into the notes app on your phone for the big day.

Plan a 3-day weekend road trip with Meta AI in 15 minutes

> ➤ *"I'm starting in Asheville, North Carolina, and I have three days. I want scenic mountain drives, good local food — nothing chain restaurants — and mostly quiet, not too touristy spots. Can you give me a day-by-day itinerary with approximate driving times between stops?"*

> ➤ *"Day 2 looks good but I don't want to drive more than 2 hours between stops — can you tighten it up?"*

> search *"Blue Ridge Bakery Asheville hours"* to see if it opens early enough for your planned breakfast stop.

Translate a foreign-language recipe or instruction manual with Hugging Face's free translator

> type `huggingface.co/tasks/translation` into the address bar.

> set the source dropdown to "Spanish" and the target dropdown to "English."

> if your dropdown only had ten languages and you need Thai or Vietnamese, switch the model to "facebook/nllb-200-distilled-600M" and the language list grows.

> paste the Spanish line "Mezcle el arroz con el caldo y deje cocer a fuego lento durante 20 minutos."

> after clicking Translate, the right box should show something like "Mix the rice with the broth and let it simmer for 20 minutes."

> paste the English result into a note titled "Paella — translated" alongside the original Spanish so you can double-check the measurements.

Build a Simple Weekly Meal Plan With Grok Based on What's in Your Fridge

> type "Hi" to test the connection, or jump straight to Step 2 with your list ready.

> ➤ "Fridge: 2 chicken breasts, half a head of broccoli, 4 eggs, a wedge of cheddar. Pantry: rice, canned chickpeas, soy sauce, olive oil, garlic. Freezer: frozen peas. Serving 2 adults."

> ➤ "Here's what I have: [paste your list]. Build me a 7-day dinner plan for 2 adults. No seafood. Keep weeknight meals under 30 minutes. Use what I already have first, then suggest what to buy."

> ➤ "Swap Tuesday's curry for something quicker — under 20 minutes. Add a slow-cooker recipe for Sunday. And please skip mushrooms entirely; nobody here eats them."

> ➤ "Great. Now give me a shopping list of only what I'm missing, grouped by aisle (produce, dairy, pantry, frozen, meat). Quantities for 2 adults. Keep it under 25 items if you can."

> ➤ "OK, I'm ready to save this. Can you format the final 7-day plan and the shopping list as plain text I can copy into my notes app?"

Plan a Budget-Friendly City Vacation in One Afternoon

> ➤ *"I want to spend 4 days in Lisbon in mid-October. Just me, no car. My total budget including flights from the US East Coast is $1,800. I like walking, local food, neighborhoods, and one or two paid museums. I don't love big crowds. Give me a first-pass plan."*

> ➤ *"Thanks. Now break the $1,800 into lodging, food, transport within the city, attractions, and a buffer. Which two categories usually eat most of the budget for a 4-day solo trip, and which is the easiest to cut without ruining the trip?"*

> ➤ *"Rebuild the 4 days using mostly free or cheap anchors — parks, walking streets, viewpoints, free museum hours, markets. Put one paid anchor per day, and group nearby things together so I'm not crossing the city for every stop. Flag anything that's usually very crowded."*

> ➤ *"Here's the Day 2 plan you gave me: Alfama walking tour → Sé cathedral → lunch in Graça → National Tile Museum → sunset at Miradouro da Senhora do Monte → dinner in Bairro Alto. Push back on this: how much walking is that, do I need to book the museum, and what's a good rainy-day swap?"*

> ➤ *"Give me a short list of Lisbon dishes with a 'cheap-eats version' and a 'sit-down version' for each. Then give me 5 'don't do this' warnings — the kind of thing a first-timer always falls for, like taxi scams, tourist-menu restaurants, or transport passes that don't pay off for 4 days."*

> ➤ *"Now compress everything we discussed into one tight page: budget shares, 4-day plan, food list, top 5 warnings, and a 5-line packing note for October weather. Keep it tight — I want to fit this on one screen on my phone."*

Plan Your Weekly Meals and Shopping List with Any AI Assistant

> don't type anything yet — just make sure you can see the empty text box.

> ➤ *"Family of four in Chicago: two adults, a hungry 14-year-old, and a 7-year-old who won't eat spicy food. No nuts or shellfish. One of us is vegetarian most nights. Budget around $120 for five dinners. We have an oven and a stovetop, no microwave."*

> ➤ *"Great, thanks. Please give me a five-night dinner plan based on that. Keep each dinner under 30 minutes of active cooking, repeat ingredients across nights where it makes sense (saves money and waste), and call out any night that needs overnight prep like defrosting or marinating."*

> ➤ *"Nice plan. Can you turn that into one shopping list grouped by aisle (produce, meat/fridge, pantry, dairy, frozen) with rough amounts? And please mark staples I likely already have at home — oil, salt, pepper, basic spices — so I don't have to buy them."*

> ➤ *"Looks good but please swap the salmon for chicken thighs (cheaper this week), change the couscous to jasmine rice, and double-check I'm not buying the same herb or spice twice across the five nights. Then re-print the shopping list."*

> a quick message you can send to yourself or a family member: *"Here are this week's dinners: Mon – chicken stir-fry, Tue – veggie pasta, Wed – sheet-pan sausages and veg, Thu – rice bowl with leftover chicken, Fri – omelette night. Shopping list attached."*

Plan a Weekend Trip From a One-Line Idea With Cohere

> ➤ *"You are a practical trip planner. I travel with one other adult, no kids. We like walking, local food, and one museum or landmark per trip max. Avoid luxury recommendations. Keep each day to three time blocks: morning, afternoon, evening. Never invent specific restaurant names — describe the cuisine and neighborhood instead."*

> ➤ *"Weekend trip idea: Portland, Oregon, second weekend of October, two adults, arriving Friday evening, leaving Sunday afternoon. Build a day-by-day plan in three time blocks per day (morning, afternoon, evening). For each block, give: one main activity, one place to eat (describe the cuisine and neighborhood, don't invent names), and a backup if it rains."*

> ➤ *"Looks good overall, but tighten this: replace any block that requires a car with a public-transit option, and move any suggestion that starts before 9 a.m. to 10 a.m. or later. Keep everything else the same."*

> ➤ *"For Sunday afternoon only, give me a backup plan that works if it rains: an indoor activity, a nearby warm place for lunch, and a coffee stop I can walk to between them. Keep Saturday exactly as it is."*

> ➤ *"Now compress the whole plan — both days plus the Sunday rain backup — into a one-screen version I can read on my phone. Each item is one short line. No descriptions, no flowery language. Order: Friday evening, Saturday morning/afternoon/evening, Sunday morning/afternoon, Sunday rain backup."*

> ➤ *"Now write me a copy-paste version of the one-screen summary that I can paste into Apple Notes or Google Keep. Use plain text, no markdown bold, one item per line, with a blank line between days."*

Use ChatGPT to Plan a Road Trip with Stops, Fuel Breaks, and Drive Time

> ➤ "My partner and I are driving from Portland, Oregon to San Francisco, California on June 12, leaving around 8am. We like small towns, good coffee, and short hikes. We want to drive no more than 6 hours a day."

> ➤ "Great. Now please give me a day-by-day plan. For each day, tell me the start city, the end city, total driving hours, and one sentence on why that stop is worth it."

> ➤ "Add a fuel stop about every 3 hours of driving, and one proper meal stop per day. Just give me the town or exit name and roughly when in the day I'd hit it."

> ➤ "For each overnight stop, suggest one outdoor thing, one local food place, and one quick museum or viewpoint. Keep it short — we're not trying to cram the whole trip in."

> ➤ "Now read the plan back to me like a skeptical friend. What looks too rushed? Which day has the longest stretch between fuel stops? Anything I should double-check before I leave?"

> ➤ "Now condense the whole trip into a clean numbered list I can paste into my notes app. Keep fuel stops inline with each day. Don't add anything new."

Creating a Travel Packing List for Any Trip with AI Guidance

> open ChatGPT, Gemini, Claude, Copilot, Le Chat or DeepSeek — whichever you already have an account with, or download a free one from your app store.

> ➤ "I'm going to Tokyo for 5 nights in early December with my partner. We'll stay in a hotel in Shinjuku, do a mix of sightseeing, two day trips, and 3–4 nicer dinners. One of us has dodgy knees so we don't want to walk more than about 10 km a day."

> ➤ "Weather forecast says 8–14°C with a chance of rain on two days. We're planning one full-day Disneyland trip, one day in Hakone with onsen, and lots of walking in Tokyo. I take a daily medication and one of us wears glasses. We each have a phone, one small camera, and one universal power adapter."

> ➤ "Turn this into a printable checklist grouped by category (clothes, toiletries, electronics, documents, medications, miscellaneous), and add a short 'easy to forget' section at the end. Use markdown checkboxes so I can copy it into my notes app."

> ➤ "Now make a separate 'to buy' list with just the items I probably don't already have, like travel-size toiletries or any specific gear for the onsen."

> ➤ "Save this whole conversation as 'Packing prompt template'. Next time: paste it, change the destination and dates, delete anything that no longer applies, and re-send."

How to scale a family recipe up or down for any number of guests

> original serves 6, you need 14 → factor = **2.33**; original serves 8, you need 4 → factor = **0.5**.

> 1 cup flour × 2.33 = about 2¼ cups; 2 tbsp butter × 1.5 = 3 tbsp.

> 1 tsp salt in the original × factor 2.33 (full math) would be ~2.33 tsp; gentler multiplier 2 → about 2 tsp is the right target.

> original 9×13 pan for 8 servings → 16 servings → use **two** 9×13 pans, or one large roasting pan.

> original bake time = 45 min → start checking at **30 min**, then every 5 min after, using a toothpick or thermometer.

> ➤ "Scaled 2× for Thanksgiving. Next time: bigger pan, 10 min less in oven, +½ tsp salt."

Plan a Relaxing Weekend Getaway From Scratch in One Sitting

> ➤ "I live in Austin, Texas, and I want a relaxing weekend trip for the second weekend of March. I have a car, around $400 total to spend, and I do not want to fly. I am traveling with one other adult and we both like quiet nature and good food, not nightlife."

> ➤ "Please redo it as an hour-by-hour itinerary for Saturday and Sunday. Keep mornings slow, plan one main activity per day, and put dinner before 8pm. We do not drink, so please skip bars and breweries."

> ➤ "Now give me a rough budget breakdown in USD for the whole weekend. Split it into lodging, food, transport, activities, and a small extras bucket. Keep the total around $400 for two people."

> ➤ "Now give me a Plan B for the same weekend if it rains nonstop or if we are too tired to do the outdoor stuff. Keep it calm, indoors, and cheap."

> ➤ "Based on this plan, what should I book in advance? Please list the items in the order I should handle them this week."

> ➤ "Thanks, please summarize everything into one short message I can paste into my notes app — Plan A, Plan B, the budget, and the booking list."

Brainstorm birthday gift ideas with Le Chat for any budget and personality

> ➤ "My sister is turning 32. She loves cooking, has a tiny apartment kitchen, and already owns a stand mixer. She hates clutter and reads a lot of mystery novels."

> ➤ "Based on that, suggest birthday gift ideas for her grouped into three price tiers: a small thoughtful gift, a mid-range main gift, and a bigger splurge option. For each idea, tell me in one sentence why it suits her."

> ➤ "I love the cookbook idea but she already has a shelf full of cookbooks. Can you replace it with something equally thoughtful in the same price range? Also, swap one of the splurge ideas for something experience-based instead of a physical object."

> Ask Le Chat: "Give me a clean final list of the top 4 gift ideas, each with a one-sentence reason and the price range, so I can paste it into my notes."

How to create a custom bedtime story starring your child with any AI

> nothing to type yet — just open the app and tap the new-chat button.

> `You are a kind bedtime story writer for young children. Please write in simple, gentle language with a happy ending.`

> `Please write a story starring my 6-year-old daughter Lily, her hamster named Biscuit, and their backyard with the big oak tree. Lily is shy but very curious.`

> `Please make it about 500 words long — the length of a 5-minute bedtime read — and gently show how being brave helps, without being preachy.`

> `This is lovely, but please change the ending so Lily and Biscuit fall asleep peacefully under the oak tree instead of running home.`

> tap the **share icon**, choose **"Copy link"** or **"Send to email"**, and save the story in a folder called "Bedtime stories."

How to organize a chaotic shared family calendar in one afternoon

> ➤ *"We have three iPhones and a Mac — let's use Apple iCloud Calendar so we don't have to install anything new."*

> ➤ *"Calendar name: Family. Color: orange. Description: shared with mom, dad, Mia, and Leo."*

> ➤ *"Tuesday 16:00–17:30, Soccer practice, Mia. Repeats every Tuesday. Reminder: 90 min before (popup) and 15 min before (notification)."*

> ➤ *"Saturday 14:00–16:00, Leo's birthday party at the community center. Notes: RSVP link — bring a swimsuit."*

> ➤ *"Sunday 19:00–19:10, Family calendar review, repeats weekly. Reminder: at start time."*

Drafting a Basic Family-Tree Outline with AI: A Beginner's Guide

> ➤ "Me — Sara. Parents: Mom = Linda, Dad = Tom. Mom's parents: Gran = Joan, Pop = Bill. Dad's parents: Nan = Margaret, Grandpa = Alan. Siblings: brother Jake. Aunts/uncles: Aunt Mary (Mom's sister), Uncle Rob (Dad's brother). Cousins: Lily (Aunt Mary's daughter)."

> ➤ "Here's my family as I know it: me (Sara), Mom Linda, Dad Tom, brother Jake, Mom's parents Joan and Bill, Dad's parents Margaret and Alan, Aunt Mary (Mom's sister) and her daughter Lily, Uncle Rob (Dad's brother). Please draw a simple text-based family tree outline with three generations — me, parents, grandparents — using indents or a basic tree shape. Note any branches I haven't filled in. This is just a starting draft."

> ➤ "Thanks. Please move Aunt Mary to sit under Mom's parents as their daughter, and add a cousin called Lily as her child. Also use dashes and greater-than symbols (├── └──) instead of dots so it copies cleanly into my document."

> ➤ "One more version please — same tree, but plain text only, no markdown or code blocks, so I can paste it straight into my notes app without fixing the formatting."

> ➤ "Add these new people to the tree: my partner Sam, our daughter Mia, and Sam's parents Pat and Chris. Please show the in-law branch as a separate section so it doesn't get mixed up with my birth family."

How to Troubleshoot Your Home Wi-Fi Step by Step with an AI Assistant

> ➤ "My home Wi-Fi has been cutting out for the past two days. My laptop loses connection every 20–30 minutes, but my phone stays connected fine. I haven't changed any settings. Where should I start?"

Generate unique profile pictures for social media with Midjourney

> `/imagine prompt: a joyful wombat wearing a tiny chef hat, cartoon style, vibrant colours`

> `/imagine prompt: a happy person with short curly hair, wearing glasses, smiling gently, digital art portrait, soft pastel colours, simple background, --ar 1:1`

How to generate AI images in Picsart: a beginner's step-by-step guide

> this step has nothing to type — just finding the right entry point.

> this step is a tap, not a typing step — the key choice is **Text to Image** for a first run.

> for a square post, choose the square option; for a phone wallpaper, choose the tall rectangle; if you see a quality choice, pick the lowest setting for a quick test.

> ➤ "A friendly poster for a Saturday morning yoga class, soft pastel colors, with the words 'Saturday Morning Yoga' clearly readable in the center, hand-drawn style."

> after the yoga poster above, send: "Make the background soft pink, change the words to 'Bring a friend!', and add a small sun in the top-right corner."

> upload a group photo, then type: "Remove the person in the red shirt on the right." Or paint over the background with the brush tool and write: "Replace this background with a sunset beach."

> not a typing step — the key choice is **JPG** for everyday sharing, **PNG** only if you need a transparent background.

Auto-write captions for your photos using Hugging Face's free image captioning tool

> the verification email is usually titled something like "Confirm your Hugging Face account".

> searching for "BLIP image captioning" may return a Space built by Salesforce, which has been one of the popular captioning demos on the site.

> a typical blurb is "A short image captioning model based on BLIP, fine-tuned on COCO" — COCO is just the name of a popular photo dataset the model learned from.

> a beach sunset photo might return something like "a sunset over the ocean with orange clouds", and a family group shot might return "a group of people smiling and standing together".

> a clear photo of a coffee cup on a wooden table might come back as "a cup of coffee on a table" — accurate, but plain.

> a simple spreadsheet column might read `beach_sunset.jpg → "a sunset over the ocean with orange clouds"`.

> a LinkedIn post becomes "Wrapped up a long week with a walk on the beach 🌅 — a sunset over the ocean with orange clouds. Felt good to breathe."

How to remove a photo background cleanly with Meta AI on your phone

> choose a clear, well-lit photo of a friend standing in front of a plain wall, or a product on a tabletop.

> `Remove the background from this photo and give me the subject on a transparent background.`

> `Zoom in around the hair — I can still see a faint gray halo. Please clean up the edges.`

> a beach sunset from your camera roll, so your friend looks like they're standing on vacation.

> `Remove the background from this photo and give me the subject on a transparent background. Keep the subject exactly the same — only the background should change.`

> `Remove the background from this photo and give me the subject on a transparent background.`

Create a multi-room music playlist with Alexa+

> Look for the house-shaped icon at the bottom of the Alexa app and tap it.

> Open the Kitchen room and look for the ··· menu in the top-right corner.

> Tap **Add to Routine** (or the equivalent option in your app version).

> Tap the voice trigger option and then tap inside the text box.

> Confirm the phrase and move to the next step. Alexa may show a summary of the phrase you've entered.

> Type "Alexa, play dinner playlist" and confirm it.

> Tap the music action and then tap the option to choose a playlist.

> Select "Dinner Vibes" playlist and tick Kitchen, Lounge and Bedroom.

Sharpen Your Photos with AI Using Upscayl

> Select a file named `old_family_photo.jpg` from your `Pictures` folder.

How to create and edit images in ChatGPT

> ➤ "You click 'New chat' and immediately type: 'Generate an image of a peaceful cottage garden.'"

> `Generate an image of a peaceful, sun-drenched cottage garden in the American countryside, filled with native wildflowers, a small birdbath, and a maple tree in the background, painted in an impressionistic art style with soft, warm colours.`

> ➤ "You notice the maple tree looks more like an oak. You also see a small download icon, usually a downward arrow or three dots, when you hover over the image."

> `Edit the image to make the maple tree in the background taller and add a small, rusty tin shed near the cottage.`

> ➤ "You click the downward-pointing arrow in the top-right corner of the image, and a file named 'image.png' appears in your browser's download bar."

> ➤ "You avoid asking for images of specific copyrighted characters or content that promotes hate speech."

How to Sort and Clean Up a Messy Phone Photo Library with Gemini

> tap **New chat**, then wait for the empty text box to appear.

> open your gallery, select the 8 beach photos from last weekend, tap **Share** → **Gemini**.

> ➤ "Here are 8 photos from the same sunset. Can you group any that look like near-duplicates, and tell me which one of each group is the sharpest or best composed?"

> ➤ "Please transcribe the text in this screenshot. Format it as: store name, date, total amount, and any reference number."

> ➤ "Based on what you saw, give me a short checklist of which photos I can safely delete and which I should keep. Keep it under 5 bullet points."

> open the 8 beach photos, delete 5 based on Gemini's notes, star the best one, and paste the transcribed receipt text into your notes app.

Safely Rename a Batch of Photos with a GitHub Copilot Script

> `Write a Python script that renames every .jpg and .png file in the current folder to "vacation-2024-001.jpg", "vacation-2024-002.jpg", and so on. Skip files that already match this pattern.`

> `Write a Python script that lists every .jpg file in the current folder and renames them to "sample-001.jpg", "sample-002.jpg", in the order they appear. Print the new names before doing it.`

How to Turn a Long Podcast Transcript Into Shareable Notes With Claude

> a clean opening line — "Host: Today we're talking about how small teams ship faster than big ones."

> no typing yet — you're just clearing the deck.

> after pasting, add this on a new line — "The text above is a podcast transcript. Please read it carefully before I ask my next questions."

> ➤ "Now please summarize this podcast in exactly 5 bullets for a busy teammate who hasn't listened. Lead with the single most useful idea."

> ➤ "Now give me two more sections: (1) the 3 most quotable lines, with who said them; (2) a 4-item action list for someone who wants to apply this tomorrow. Label each section clearly."

> ➤ "Summary: [paste 5 bullets]. Best quotes: [paste 3 lines]. Action list: [paste 4 items]. Source: [episode name and link]."

How to Summarize a Long YouTube Video in 60 Seconds with Any AI

> Your copied link will look something like youtube.com/watch?v=AbCdEfGh123 — the letters and numbers at the end are the video's ID.

> ➤ "Here's a YouTube link: [paste your link]. Summarize this video in 5 bullet points and keep it under 150 words."

> ➤ "Now pull out 3 key quotes from that video, with the speaker's name for each one. Keep each quote to one sentence."

> ➤ "In one sentence, what's the single most important takeaway from that video, written so I can send it to my manager?"

Sharpen Old Videos with NVIDIAAI Super-Resolution Online

> Open your web browser and search for "NVIDIA AI Video Super Resolution online demo". Look for a link from an official NVIDIA website or a reputable AI demonstration platform like Hugging Face or Replicate, often with "NVIDIA" or "Super Resolution" in the title.

> Look for a button or box labeled "Upload your own video," "Choose File," or "Drag and drop video here." Click it, select your short video clip (e.g., an MP4 file under 10 seconds), and click "Open" or "Upload."

> On the demonstration page, find the settings panel or enhancement options. Select "4x Upscaling" or "Ultra Quality" to get the maximum possible detail out of your old video file.

> Click the "Generate," "Process," "Start Upscaling," or "Run" button to begin the transformation of your video.

> Review the side-by-side comparison of your original and enhanced videos. Once you are satisfied, click the "Download," "Save," or "Export" button to save your new, high-quality video to your computer.

Create Royalty-Free AI Music Loops for Podcasts and Videos

> `A calm, ambient electronic loop for a meditation podcast, featuring soft synth pads, a subtle bassline, and a gentle, pulsing rhythm. Focus on a peaceful and reflective mood.`

How to use Kling AI video templates in Picsart Flow

> search terms to try inside the app — `Kling`, `AI video`, `photo to video`.

> `The woman slowly turns her head to the right and smiles, with soft sunlight on her face.`

> `A young woman with short black hair, wearing a yellow raincoat, walks through a misty forest at dawn.`

> `Change the setting to a neon-lit Tokyo street at night, keep the same walking motion.`

> trim suggestion — `keep only seconds 2 through 6` (drag the handles until those seconds are highlighted).

How to create images with AI: a beginner's first try

> ➤ "Create an image of"

> ➤ "A fluffy ginger cat wearing a blue knitted jumper, sitting on a wooden park bench, bright sunny day, soft watercolour painting style."

> ➤ "Change the jumper color to yellow, and add a small friendly magpie sitting on the bench next to the cat."

How to Edit Your Phone Photos with Simple Text Instructions in Meta AI

> Open WhatsApp → tap the new-chat icon → tap "Meta AI" in the contacts list.

> One landscape photo where the sky looks flat and gray — perfect for your first edit.

> ➤ "Make the sky look like a bright sunset with orange and pink clouds" — or simpler, "Remove the person on the right" or "Brighten the photo so it doesn't look so dark."

> Ask yourself — is the person really gone, or is there a faint shadow where they stood? Is the sky the color you wanted, or does it look painted on?

> First try didn't remove the person cleanly? Send: "Try again — remove the person on the right side, but keep everything else exactly the same."

> Save your favorite version, then close the chat. Your original photo is untouched on your phone, and the edited version is now saved next to it.

How to Generate Images with Google Gemini: A Beginner's Guide

> ➤ "Create an image of..."

> ➤ "Create an image of a fluffy golden retriever puppy wearing a bright red collar, sitting on a sun-drenched green lawn next to a yellow tennis ball."

> ➤ "Generate a watercolour painting of a quiet beach at sunset, with soft purple and orange colors in the sky."

> ➤ "Can you make the sky in that last picture a bit more blue, and add a single sailboat in the background?"

Create Your First AI Image with Midjourney: A Beginner Walkthrough

> when asked for your email, use the one you check most often — that's where receipts and account notices will go.

> `a cozy reading nook by a window, warm lamp light, watercolor illustration, soft colors`

How to Create AI Logos with Recraft V4.1 Vector

> type a one-line description like "warm coffee shop, hand-lettered feel, soft cream and forest green" into the prompt box.

> prompt: "minimalist logo for a yoga studio, lotus flower, calm pastel palette, no text in the image."

> if you loved the lotus shape but the colors felt too cold, change "calm pastel palette" to "warm peach and sand palette" before you continue.

> paste your tightened prompt from Step 3 — for instance, "minimalist lotus logo, warm peach and sand palette, simple shapes only, no text."

> open the file in a browser tab, then zoom out until it looks tiny. If the lotus shape is still recognizable, you're good.

> `lotusyoga-logo-minimal-peach.svg`, `lotusyoga-logo-simple-mono.svg`, `lotusyoga-logo-with-text.svg`.

How to Create Custom Images with Copilot for Invites and Social Posts

> don't type anything yet — just make sure you can see the image-creation prompt box.

> ➤ "A watercolor birthday invitation background with soft pink peonies, gold confetti, and the words 'Emma turns 30' written in elegant script at the top."

> ➤ "Same image but make the peonies larger, and put the text 'Emma turns 30 — Saturday June 14' in the center."

> for the invite, ask for "A watercolor birthday invitation background with soft pink peonies and gold confetti, no text." Then you add "Emma turns 30 — Saturday June 14" yourself.

> click the **Download** button below your favorite peony image and wait for the file to appear in your Downloads folder.

> in Canva, click **"Uploads"**, drag in your downloaded peony image, drop it onto a square Instagram template, and add the text "Emma turns 30 — details in bio" on top.

Use Gemini on Your Phone to Pull Editable Text from Any Photo or Screenshot

> nothing to type yet — this step is just opening the app.

> no text needed yet, just pick the image.

> ➤ "Pull out all the text from this photo, line by line, exactly as written." — or — "Read the recipe in this image and list just the ingredients."

> paste it into Notes and you'll have a searchable copy of that faded parking sign — no more squinting at the photo.

> ➤ "Now translate that to Spanish." — or — "Shorten this list to the five most important points." — or — "Turn these ingredients into a grocery checklist with quantities."

How to Turn Your Saved Articles and Notes into a Personal Podcast with NotebookLM

> type a short name like "Weekend reading" or "Q3 research" into the notebook title field so you can find it later.

> paste the URL of one long article you've been putting off.

> add 4 items — a long-form article, a PDF report you downloaded last week, a Google Doc with your own notes, and one more link.

> type an optional instruction in the custom prompt box if you see one, such as "Focus on the parts about pricing strategy."

> while listening, jot down a question in the Notes section on the right so you can ask NotebookLM's chat about it after.

How to Design a Simple Logo for Your Side Project with Midjourney

> in Discord, type a forward slash in the message box and you'll see a menu. Pick /imagine from the list.

> ➤ "minimalist line drawing of a coffee cup with steam, single dark color on white background, modern, clean, simple"

> if you liked the third image in the first row, click V3.

> if your favorite is in the bottom-right slot of the latest grid, click U4.

> in Inkscape, select the image, choose Path → Trace Bitmap, click OK, then save as SVG.

Practice a new language or presentation with Claude for instant feedback

> `Act as a friendly French tutor for a beginner. Ask me simple questions about my daily life in French, and gently correct any grammar or vocabulary mistakes I make, explaining why.`

> `You are my presentation coach. I need honest, constructive feedback on my script's clarity, conciseness, and how engaging it sounds. Imagine I'm presenting to a general audience. I'll share parts of my script one by one.`

Build a personal study-buddy chatbot in 15 minutes using Hugging Face Spaces (no code)

> You can copy this into your browser:
> https://huggingface.co/spaces

> Space name: `my-study-buddy`

> You'll see a single line:
> `def main():`

> You can copy the code above and paste it directly into the editor.

> Type: "Define study buddy"

> Replace the entire `respond` function with the code above.

Build a simple weekly study plan with Meta AI in 10 minutes

> *Just open the app or website.*

> *Math exam 15 June, History essay due 20 June, Chemistry quiz 12 June, English speech 25 June*

> *Create a 7-day study plan that spreads the work evenly and includes short breaks.*

> *Move the Math block to 4 pm each day instead.*

> *For Math, list the exact chapters and topics I should study each day.*

> *Tap the share icon and choose “Copy.”*

> *Paste the plan into tomorrow’s calendar events.*

Getting Started with Mistral's Le Chat: Your Private AI Guide

> Click the "Sign up" button. On the next screen, enter your email address in the designated field, then create a strong, secure password. Click the "Submit" or "Continue" button. You will then receive a verification email from Le Chat; open this email and click the link inside to activate your account.

> From the menu that appears, select "Settings" or "Data Privacy." You will then likely see an option labeled something like "Store Chat History" or "Use my data for training." Toggle this option **off** (so it shows as gray or inactive, rather than blue or green). This tells Le Chat not to use your specific conversations to further train its AI.

> Copy and paste the following text into the box, then press the "Send" button (often a paper airplane icon or similar) located on the right side of the text box:
> ➤ "Please organize this messy shopping list into neat categories like 'Fresh Produce', 'Dairy', and 'Pantry': apples, milk, carrots, cheese, pasta, bananas, butter, eggs, bread, coffee, chicken."

> Type this follow-up message into the chat box and send it:
> ➤ "Thank you! Now, can you rewrite that list but put a little checkbox symbol [ ] next to each item so I can print it out and tick them off when I go shopping?"

> Type this prompt into the box and send it using the paper airplane or "Send" button:
> ➤ "I'm planning a small outdoor birthday party for a friend who loves vintage sci-fi movies and board games. Give me five unique theme ideas for the party, including a brief description for each."

How to Translate and Understand Other Languages with AI

> Click the "New chat" button.

> Please translate the following Italian text into English and keep a relaxed, conversational tone:
>
> Ciao! Come stai? Spero che la giornata vada bene. Ci vediamo stasera?

> If the AI translated "Ciao! Come stai?" as "Hello! How are you?"

> Following on from the previous step, you might type:
> Could you make the greeting even more casual, like "Hey there! How's it going?" for "Ciao! Come stai?" And ensure the entire message keeps that very relaxed feel.

> Translate the following section of a legal document into plain English, maintaining a formal and precise tone, without using any legal jargon where a simpler term will suffice:
>
> [Paste the first few paragraphs of your legal document here]

> Highlight the AI's translation, press **Ctrl + C** (or **Cmd + C** on Mac), open a new Google Doc, and press **Ctrl + V** (or **Cmd + V** on Mac) to paste.

Use ChatGPT to Practice a Foreign Language Without Feeling Awkward

> name the chat *"Spanish practice — beginner"*.

> ➤ *"Please reply only in simple Spanish sentences. I'm a complete beginner. After each of my messages, correct any mistakes in a kind bullet list, then ask me a follow-up question to keep the conversation going. Keep your replies short."*

> if ChatGPT asked *"¿De dónde eres?"*, type *"Yo soy de México"* — or even the imperfect *"Yo es de México"* — so ChatGPT has something to correct, which is the whole point.

> say out loud *"Buenos días, ¿cómo estás?"* and wait for ChatGPT to answer in voice.

> ➤ *"I keep mixing up 'ser' and 'estar.' Can you give me 5 more practice sentences and tell me which one to use, with a short explanation?"*

> ➤ *"Please summarize today's lesson in 5 bullet points I can review before our next session, and list 3 things we should practice again."*

How to Build a 30-Day Plan to Learn Almost Any New Skill from Scratch

> ➤ "I want to learn beginner conversational Spanish so I can order food and have small talk on a trip to Mexico."

> ➤ "20 minutes a day on weekdays, weekends flexible, complete beginner, no paid courses."

> ➤ "I want to learn beginner conversational Spanish for a Mexico trip in 30 days. I have 20 minutes a day on weekdays, weekends are flexible, and I'm a complete beginner. Please build me a day-by-day plan that: 1) fits in 20 minutes on weekdays and up to 60 minutes on weekends, 2) focuses on speaking and listening first, 3) uses free resources only, and 4) includes a small checkpoint at the end of each week. Format as a numbered list, one entry per day."

> ➤ "Days 18–22 ask for 90 minutes each. I only have 60 minutes on weekends. Please rewrite that week with shorter sessions and remove the paid app recommendation."

> ➤ "I'll paste the plan into my calendar as 20-minute events on weekdays, and a 60-minute block on Saturday morning."

> ➤ "I'm on day 8. I skipped day 5 and day 6 because of work. Days 1–4 went well, but the 30-minute listening exercise felt too long. Can you adjust the rest of the plan so I still finish in 30 days total, and shorten the listening parts to 15 minutes?"

How to turn your messy study notes into practice quiz questions

> no typing in this step — you're just opening a clean chat.

> (a short chunk you might paste):
> ➤ "Cell biology chapter 7. Mitochondria = powerhouse of the cell, makes ATP. Double membrane: outer is smooth, inner folded into cristae. ATP made via the Krebs cycle in the matrix and the electron transport chain on the cristae. Chloroplasts are found only in plant cells and do photosynthesis. Both mitochondria and chloroplasts have their own DNA — endosymbiosis theory."

> (a prompt you can copy and adapt):
> ➤ "Turn the notes above into: (1) 5 multiple-choice questions with 4 options each, (2) 5 short-answer questions, (3) 5 flashcard-style Q&A pairs. Put the correct answer on a separate line for each one."

> (a quick fix prompt):
> ➤ "Question 3 says the answer is 1972, but my notes say 1927. Please correct it and double-check the rest for similar mistakes."

> (a paste title): "History Unit 4 practice questions — 12 March 2026"

> (a self-test routine): "10 questions, cover the answers, answer first, then check — repeat for any I missed."

Boost Your Learning: Get Quick Explanations and Study Guides with Copilot

> type `Hi Copilot, can you help me understand a tricky topic?` and press Enter — you'll see a short, friendly reply appear in seconds.

> `Explain how photosynthesis works in simple terms, like I'm 12 years old.`

> `Can you give me a real-world example of photosynthesis that I see in everyday life?`

> `Great, now turn this into a short study guide with 5 key points and 5 practice questions I can test myself with.`

> paste it into a note titled `Photosynthesis – quick study guide` so you can find it easily next time.

> `Can you give me one short sentence I can remember to recall how photosynthesis works?`

How to use Qwen Chat to create personalized study flashcards for any subject

> ➤ *"Hi Qwen — I'm studying for a Year 9 biology exam on cell structure. Here's my notes: [paste your notes here, or describe: 'The main topics are animal cells vs plant cells, the function of the nucleus, mitochondria, ribosomes, cell membrane, and how osmosis works.' Please make 12 flashcards covering these."*

> ➤ *"Great, please rewrite those as a numbered list, with each one formatted exactly like this: 'Card 1 — Question: [question] | Answer: [answer]'. Use a vertical bar to separate Q and A so I can copy them straight into a spreadsheet."*

> ➤ *"Please add 6 more cards, but make them harder — focus on application and 'why' questions rather than simple definitions. Then quiz me on the first 5 cards one at a time."*

> ➤ *"I got these wrong: [list 2-3 topics]. Please make 4 more cards on those specific points, and include a tiny hint on the back of each one."*

Turn a Reading into a Study Quiz with GitHub Copilot

> don't type anything yet — just confirm the chat is open and you can click into the box.

> paste this short paragraph — *"The Amazon rainforest produces about six percent of the world's oxygen and stores billions of tons of carbon. It also hosts ten percent of all known species. Fires there release stored carbon and reduce the forest's ability to absorb more."*

> ➤ *"Read the passage I pasted. Write 5 multiple-choice questions about it, with 4 answer choices each. Make 2 easy recall and 3 about meaning or application. Don't give me the answer key yet."*

> ➤ *"Now give me the answer key. For each question, write the correct letter (A/B/C/D) and quote one short phrase from the original passage that supports it. Don't paraphrase — quote directly."*

> if the answer says "B — 'ten percent of all species'" but the original paragraph says the rainforest *hosts* (not produces) ten percent of species, you've caught a subtle mistake — correct the question before using it.

> paste your verified list into a note titled "Quiz — Amazon rainforest (WIP)" so you can come back and revise it later.

How to Create a Simple Weekly Study Plan with Any AI Assistant

> no typing here yet — just confirm the empty chat is open and ready.

> Here's everything I have to do this week (Monday to Sunday). Some classes are fixed times, the rest is flexible. Build me a realistic day-by-day plan with blocks for studying, breaks, meals, sleep by 11pm, and one hour of exercise. Flag anything you couldn't fit. Mon 9am Biology 101 lecture (fixed); Tue 6pm cafe shift until 10pm (fixed); Wed Psychology essay due 11:59pm; Thu 7pm soccer practice; Fri nothing fixed; Sat family lunch 1-3pm; Sun Statistics problem set due.

> Now read that plan back like a strict friend. Tell me which days are unrealistic, where I forgot breaks, and what you would cut first if I only had four hours of real focus time per day. Give me a v2.

> prompt, paste in your real list, and see what comes back.

Turn Your Messy Class Notes into Study Flashcards with Cohere

> a one-line email like "yourname@example.com" works fine to create the account.

> copy only the bullet points about "phases of mitosis," not the whole biology chapter.

> ➤ "Turn these notes into 10 study flashcards. Format each one as: Q: [question] / A: [answer]. Cover the main definitions, dates, and examples."

> ➤ "The answer to card 3 is too long. Make every answer under 15 words, and remove any dates I didn't include in the notes."

> paste into a new Google Doc titled "Bio 101 — Mitosis flashcards" so you can find it again later.

> spend 5 minutes quizzing tonight, then 5 more tomorrow morning on the cards you missed.

AI Made Simple: A Gentle First Guide for Beginners

> Type **chatgpt.com** or **gemini.google.com** into your browser's address bar. Then, click "Sign up" and follow the prompts to create your free account.

> Type **Hello, how are you today?** into the chat box and press Enter.

> Instead of typing short, blunt terms like "scone recipe easy", type:
> **Could you please share a simple recipe for traditional scones that uses basic ingredients I likely have in my pantry?**

> If the recipe it gave you looks too detailed, you can type:
> **That recipe looks good, but could you please simplify it into just five clear, numbered steps?**

> **- Safe prompt: What are some gentle exercises to help improve balance at home?**

How to pick the right free AI app when you're just getting started

> ➤ "I want to draft a short email to my manager explaining why a project deadline needs to move."

> prompt to try in two tools: "Write a 150-word product description for a small handmade candle business."

> prompt: "Explain what a budget spreadsheet is to someone who has never made one."

> prompt: "Write a simple Python script that asks for your name and prints a greeting."

> After comparing results from Steps 2 through 4, note which tool you would open first if you had a real task due tomorrow.

How to study and learn faster with AI – a beginner's step-by-step guide

> Visit `https://chat.openai.com`. On the first page, click the "Sign up" button. Follow the prompts to enter your email, create a password, and verify your email address. You may also need to provide your name and date of birth. After completing the short setup, you'll land on the main chat interface. If it looks different, look for "Log In" or "Sign Up" usually in the top-right corner or center of the screen.

> Type into the message box:
> Explain photosynthesis in simple terms, like I'm a high-school student with no science background.

> Type into the message box:
> Summarize the main steps of photosynthesis in 5 bullet points.

> Type into the message box:
> Make three multiple-choice questions about the bullet-point summary of photosynthesis. Include one correct answer and three distractors for each question.

> Type into the message box:
> Here are my answers to the photosynthesis quiz: 1-B, 2-A, 3-C. Which are correct, and can you explain any mistakes?

Using Perplexity AI for Quick Study Summaries and Exam Prep

> Click on the **Focus** button and select **Academic** if you are preparing an essay for a history or science topic.

> ➤ "I am a biology student preparing for an exam. Summarize the key differences between mitosis and meiosis. Present this as a simple comparison table followed by bullet points explaining the main phases of each."

> ➤ "Based on the summary of cell division you just created, write a 5-question multiple-choice quiz. Hide the answers at the very bottom of your response so I cannot see them while I test myself."

Use Alexa+ to Pick Up Common Phrases in Another Language

> say *"Alexa, how do I say 'where is the bathroom?' in Spanish?"*

> say *"Can you use that in a short example sentence? And could you say it more slowly?"*

> say *"Alexa, can you spell that phrase for me? I'd like to write it down."*

> say *"Could we walk through ordering a coffee in Spanish? Show me what I would say, then what the barista would say back."*

> say *"Thanks! Please set a reminder for tomorrow at 6pm to practice my Spanish phrases again."*

The main AI assistants explained simply (and what each is best at)

> ➤ "Give me a step-by-step plan to organize my garage this weekend, including how to categorise items."

> ➤ "Read this draft email to my landlord and help me make the tone more polite but firm."

> ➤ "What are the latest updates on the local public transport extensions, and where can I read the official announcements?"

> ➤ "Translate this business greeting into polite conversational Japanese and explain any cultural etiquette I should keep in mind."

How to use Meta AI inside WhatsApp – step-by-step for beginners

> Type "Meta AI" into the search bar.

> Type "Hi Meta AI, what can you do?" and press the send button (usually a paper plane icon).

> Type "What's a simple, healthy breakfast idea I can make in under 5 minutes?" and send it.

> Tap and hold Meta AI's breakfast recipe, then choose "Copy" from the menu that appears.

> After getting the breakfast idea, type "Can you suggest some healthy lunch ideas too, keeping the 5-minute prep time in mind?" and send.

> Go back to your main WhatsApp "Chats" screen, then tap the green sparkle icon or the "new chat" icon to begin a fresh conversation.

ABeginner's Guide to Using ChatGPT Voice for Hands-Free Tasks

> ➤ "Hi! I'm planning a weekend trip to the Blue Ridge Mountains. Can you suggest some scenic bushwalks that are good for beginners?"

> ➤ "Those sound great! What's the best way to get to the Wentworth Falls track if I'm coming from Chicago?"

> (While ChatGPT is explaining) "ChatGPT, hold on a moment please."

> ➤ "Okay ChatGPT, tell me a short mystery story, but pause every two minutes so I can guess the suspect."

Getting started with DeepSeek for everyday help

> Copy and paste this prompt into the chat box to try it out:
> ➤ "I have a busy Saturday ahead. I need to take the dog to the vet at 9:00am (takes 1 hour), buy groceries (takes 45 minutes), and drop off a birthday present at my sister's house, which is a 20-minute drive from the vet. I also need to be home by 1:00pm for a delivery. Can you organize the perfect, stress-free schedule for me, including realistic travel times?"

Using Claude to Understand Complex Topics for Study or Work

> `Explain the concept of 'quantum entanglement' in physics.`

> `Can you explain that in simpler terms, like you're talking to a high school student?`

> `Please give me two everyday examples where this concept might apply.`

> `You mentioned 'decoherence' in the first explanation. Can you elaborate on what that means specifically?`

> `Can you summarise our conversation about quantum entanglement in three short paragraphs?`

How to Use ChatGPT Voice Mode to Practice Conversations in a New Language

> nothing to type yet — just tap the icon.

> pick the voice whose short sample feels the clearest in your earphones, then move on.

> ➤ "I'd like to practice ordering at a café in Italian. I'm a beginner, please speak slowly, and correct my biggest mistakes gently."

> order your coffee — "Buongiorno. Vorrei un caffè e un cornetto, per favore."

> ➤ "Was that sentence correct? And how would a native speaker actually say 'I'd like a coffee'?"

> ➤ "Thanks! Can you list three new words I used today with a quick English meaning for each?"

Organize your work with Claude Projects – a step-by-step beginner's guide

> `email@example.com` + `YourPassword` → press **Enter**.

> Name: Q3 Marketing Campaign
> Description: All copy, ad ideas and feedback for the upcoming quarter

> Tick "Ad copy brainstorming" and "Email subject line ideas", then click **Add**.

> Drag `CampaignBrief.pdf` from your desktop onto the upload box.

> Summarize the key ideas from all attached chats and the brief, and suggest three headline options.

> Click **Share**, then press **Copy link**.

> `Give me a quick summary of the chat and the uploaded file.`

Automate Email Drafts with AI for Faster Inbox Management

> (for OpenAI): Go to `platform.openai.com/api-keys`. If you're signed in, you'll see your API Keys page. Click the "+ Create new secret key" button. A new key will pop up. Click the green 'copy' icon next to it and save it.

> (for Gmail): Click ⚙ > "See all settings" > "Filters and Blocked Addresses" > "Create a new filter." In the "Subject" field, type `AI Draft Request`. Click "Create filter" and confirm any warnings.

> prompt:
> You are a professional customer service assistant. Write a polite, concise draft email reply in US English. Acknowledge the sender's message and explain that we've received their query. Clearly state that someone will respond within 2 business days. Keep the reply under 60 words.
>
> Email content to reply to:
> {{email_body}}

> (for Zapier):
> 1. On the "Trigger" screen, search for and select "Gmail."

> (for Zapier):
> 1. On the "Action" screen, search for and select "OpenAI."

> (for Zapier):
> 1. On the "Action" screen, search for and select "Gmail."

> After a test run, you find the draft is a bit too formal. Go back to Zapier, edit your OpenAI action, and change the prompt line "Write a polite, concise draft email reply" to "Write a friendly, informal, and very brief draft email reply." Then, send another test email to yourself.

How to Get Quick Summaries from Long Documents with Claude

> ➤ "Please summarize this document for me, highlighting the main points and key takeaways relevant to small businesses."

> ➤ "Can you list the top three challenges mentioned in the document and suggest potential solutions?"

Turn Messy Notes into a Polite Meeting Agenda in Five Minutes

> ➤ "Below is the email thread about our Q3 launch. Goal: 30-minute sync to unblock the design freeze. Attendees: Maria, John, Sarah."

> ➤ "Now produce the draft. Use a polite opener, list each topic with a duration in minutes and a single owner, leave a 2-minute buffer between topics, and end with a one-line confirmation request."

> ➤ "The opener feels stiff. Loosen it. Our team uses first names and avoids phrases like 'circling back'. Keep the same structure and timings."

> ➤ "Confirm each topic has exactly one owner. Add up the minutes and show the total at the end. If anything overflows, suggest what to cut."

> ➤ "Add a 'Next steps' block at the end: Maria sends the calendar invite by tomorrow morning; John brings the latest design mockup; follow-up sync the following Monday."

> a doc titled "Agenda from messy notes" with three labeled blocks: [PASTE INPUT], [STRUCTURE REQUEST], [TONE NOTES].

Build a simple review form for your small business with GitHub Copilot

> name your repository `review-form` so it's easy to find later.

> ➤ *"Write a clean HTML form for a small bakery called 'Sunrise Sourdough'. Ask for the customer's name, a star rating from 1 to 5, and a comment box. Include a nice 'Submit review' button."*

> if Copilot's reply was a chat message instead of a code block, copy it from the chat, then paste it into the file the same way.

> ➤ *"Now add CSS that makes the form look warm and friendly — soft cream background, rounded corners, and a brown 'Submit' button. Keep it simple."*

> ➤ *"Update the form so it submits to Formspree. Use this endpoint: https://formspree.io/f/your-code-here, and change the button text to 'Send my review'."*

How to Clean Up a Flooded Email Inbox Using AI in One Afternoon

> open chatgpt.com, claude.ai, gemini.google.com, or copilot.microsoft.com.

> in Gmail, tick the checkbox at the top of the inbox. If a "Select all conversations" link appears, click it. Open the ⋯ menu and choose "Export."

> prompt: *"Here is a list of my unread emails. Group them into 5–7 folder names that make sense, and tell me which emails go in which folder. Use plain folder names a normal person would understand. Don't add explanations yet — just the list."*

> in Gmail, look in the left sidebar, scroll to "Labels," click the "+" next to it, and type "Bills & Receipts." In Outlook web, right-click "Folders" in the sidebar and choose "New folder."

> type `label:unread from:newsletter` in the search box, select all results, move them to your "Newsletters" folder in one click.

> prompt: *"Here are the remaining unread emails. Write me one short paragraph (under 100 words) telling me what I actually need to act on today, what's due this week, and what I can ignore. Start with the most urgent."*

> copy the paragraph into a new email titled "Today" addressed to yourself, and pin it at the top of your inbox.

Check your resume for ATS-friendly keywords using Hugging Face's Inference API

> Email: yourname@example.com
> Username: resumehelper
> Password: MyPassword123!

> I need help finding ATS-friendly keywords for a job application.
> Here is the job ad:
> We are looking for a digital marketing specialist with experience in SEO, Google Analytics, and content creation.
>
> Here is my resume:
> I have 3 years of experience in digital marketing, including social media management and email campaigns.

> Rewrite this sentence from my resume to include the keyword SEO:
> "I have 3 years of experience in digital marketing, including social media management and email campaigns."

Write a Polite Complaint Email to a Company with DeepSeek

> type nothing yet — just confirm the box is there.

> ➤ "I bought a coffee maker from BrewCo on March 12, order #44291. It arrived broken — the carafe was cracked and water leaks everywhere. I emailed them a week ago and nobody replied. I want a replacement or a full refund."

> ➤ "Please turn this into a polite complaint email I can send to BrewCo customer service. Keep it firm but respectful, around 150 words, and give me a subject line too."

> ➤ "Please rewrite it so I clearly ask for a full refund of $89 within 7 business days, and mention that if I don't hear back I'll file a chargeback with my bank. Keep the same polite tone."

> ➤ "Make it a bit warmer — I'm not angry, just disappointed. Keep the firmness but add one sentence saying I have liked their products in the past."

> replace "[Your Name]" with your real name, "[Order #]" with the real order number, and double-check the dollar amounts.

> write a short follow-up — "Hi, I'm following up on my email from March 20 about order #44291. I haven't received a reply yet. Please let me know how you'd like to resolve this."

Drafting Professional Presentations from Scratch Using Microsoft Copilot in PowerPoint

> Type or paste:
> Help me create a presentation about the benefits of remote work for a team of 20 software developers.

> Create an outline for a 10-minute presentation on 'How Remote Work Boosts Developer Productivity'. Include slides on flexibility, tools, and team culture.

> For a slide titled “Flexibility,” Copilot might add:
> Speaker notes: Remote work lets developers choose their most productive hours. Studies show 60% of developers report higher focus at home. Mention our recent productivity report from Q2.

> Add a quote from Satya Nadella about remote work culture to the last slide.

Digitize Handwritten Notes into Editable Text with AI

> Place your handwritten notes on a plain, uncluttered surface, like a clear desk. Position them under natural light from a window, if possible, or use a well-lit lamp to avoid casting shadows over your writing. Make sure the paper is flat and smooth, with no creases.

> Tap the "Google Keep" app icon on your home screen or app drawer (it looks like a yellow and white notepad). Alternatively, tap the "Microsoft OneNote" app icon (a blue notebook icon).

> In Google Keep, tap the plus (+) icon at the bottom of the screen, then choose the camera icon (it might be labeled 'Take photo' or 'Image'). In Microsoft OneNote, select 'Camera' from the options at the bottom, then hold your phone steady over your notes and tap the shutter button.

> In Google Keep, after your photo is captured, tap on the image to view it full screen. Then, look for a three-dot menu (⋯) in the top-right corner of the screen and tap the option labeled 'Grab image text'. In Microsoft OneNote, after taking a photo, tap the 'Text' icon that appears on the captured image to extract the text.

> Carefully read through the digitized text. If you spot a word like "flow" incorrectly appearing as "flew," tap on "flew" and type "flow" to correct it. Also, remove any extra spaces, line breaks, or unwanted characters the AI might have picked up from smudges or background textures on the paper.

> Tap the 'Save' or 'Done' button within your chosen app to store the note. Alternatively, look for a 'Share' icon (often a three-dot connection or an upward-pointing arrow) to copy the text to your phone's clipboard or send it directly to an email, a document app like Google Docs or Microsoft Word, or cloud storage like Google Drive or OneDrive.

How to turn a long PDF contract into a plain-English summary in 5 minutes

> ➤ "I'm uploading a contract. Please don't analyze it yet — I'll ask my questions one at a time."

> ➤ "In plain English, what is this document, who are the parties, and what is the main thing being agreed to?"

> ➤ "From this contract, list: 1) every deadline or time limit I need to know, 2) every fee, charge, or price that might apply, and 3) exactly how I cancel, including any notice period. Use bullet points."

> ➤ "What happens if I miss a payment, who covers legal fees in a dispute, can the company change the terms mid-contract, and which state's laws apply? Point me to the section for each."

> ➤ "Summarize everything we just found into a one-page checklist I can save: dates I must remember, amounts I'll be charged, and the steps to cancel. Keep it under 15 lines."

> ➤ "Cancellation clause says 30 days, fees say a charge after year one, and auto-renewal is 60 days before expiration — can you re-read sections 7, 9, and 12 and tell me if your summary matches?"

Unlock Document Insights with Google NotebookLM

> `My Research Notes on Renewable Energy`

> `What are the key challenges mentioned in the "Renewable Energy Report.pdf"?`

> `Can you summarise the solutions proposed for reducing these challenges?`

Summarize Any Long Document or Article with AI

> Open a long email thread. Press `Ctrl+A` (Windows) or `Cmd+A` (Mac) to select the entire email body, then press `Ctrl+C` (Windows) or `Cmd+C` (Mac).

> Open your web browser and navigate to `gemini.google.com`. Log in to your account if prompted, then ensure you have a blank conversation ready.

> Type this into the chat box (do not press Enter or Send yet):
> ➤ "Please summarize the following document in three concise bullet points. Keep the language neutral and factual, avoiding any personal opinions or interpretations."

> ➤ "Please summarize the following document in three concise bullet points. Keep the language neutral and factual, avoiding any personal opinions or interpretations.

> ➤ "The third bullet point mentions 'Section 5'. Please confirm the exact wording of Section 5 from the original text and update that bullet point to be more precise. Also, can you ensure the overall tone is strictly professional?"

> Paste the finalised bullet points into a new email, a team chat application like Microsoft Teams, or a note-taking app on your phone to share or save.

How to Prepare for a Job Interview Using a Pasted Job Post

> name it "Acme – Marketing Lead interview prep."

> JOB POSTING:
> [Paste the full job description here — responsibilities, requirements, and any "nice to haves"]
>
> ABOUT ME:
> [Paste 4–6 sentences: current role, years of experience, two recent achievements with numbers if you have them]

> ➤ "Based on the job posting and my background above, give me ten interview questions a hiring manager is most likely to ask me. Mix behavioral, technical, and situational questions. For each, add one sentence on what they're really trying to find out."

> ➤ "Now for each of those ten questions, give me two concrete talking points I could use, drawing from my background. Keep each point to one or two sentences. Flag any question where my pasted background seems thin."

> ➤ "Give me five strong questions I, the candidate, can ask back at the end of the interview. Make them specific to this role and company. Skip anything generic like 'what's the culture like'."

> pick the talking points that felt truest to your real work — maybe one about leadership, one about a measurable result, one about a hard problem you solved — and re-read them once before you walk into the interview.

Write a Resume From Scratch (Even With No Work Experience)

> section headers to copy — `Summary`, `Experience`, `Education`, `Skills`, `Projects`.

> ➤ "High school senior finishing in May 2026. Strong with customers, cash handling, and staying calm under pressure. Looking for part-time retail or front-desk work."

> (babysitting): `Cared for three children (ages 4–9) after school each day. Prepared snacks, helped with homework, and organized weekly craft activities.`

> (school project): `Led a 4-person team on a science fair project. Researched topics, delegated tasks, and presented findings to a panel of 5 judges.`

> education entry: `Lincoln High School — Expected graduation May 2027. Coursework: Business Math, Intro to Computer Science, Spanish II.`

> skills line: `Skills: Google Docs, Sheets, Canva, social media scheduling (Meta Business Suite), conversational Spanish, POS cash handling.`

> project entry: `Project: Neighborhood plant-swap app — designed a simple scheduling tool in Google Sheets, promoted via local Facebook group, coordinated 3 successful swaps with 40+ attendees.`

> cut this — "Responsible for various tasks including helping customers." Replace with — "Greeted 80+ customers daily, restocked shelves, and handled returns at a family-run grocery store."

How to Organize a Messy Email Inbox Into Simple Folders in One Sitting

> tell yourself out loud, "I'm spending the next 35 minutes on this. No answering anything that needs more than 60 seconds."

> search "Amazon" and count how many emails you find. If the number is in the double digits, that's a sender you'll handle in Step 3.

> for the Amazon thread - open one message, click "unsubscribe" near the bottom, wait for the confirmation page, then select all Amazon emails (look for a checkbox at the top of the list) and click "Archive" or "Delete."

> name them "To act on," "To read later," "To keep" - exactly those three, no more. Resist the urge to add a fourth on the same day.

> a newsletter you want to read on the weekend → drag to "To read later." A package tracking update → drag to "To keep." A friend's question that needs a real answer → drag to "To act on."

> in Gmail settings, choose "Create new filter" → From: newsletter@example.com → "Apply the label: To read later." Every future issue lands in the right folder automatically, no work from you.

Practice job interview questions with Qwen before the real interview

> just type "Hi" first to confirm the chat is responsive, then move to Step 2 to set up the real practice.

> ➤ "I have an interview next week for a marketing coordinator role at a mid-size SaaS company. I have 2 years of experience in social media and email campaigns. Please act as the hiring manager and interview me one question at a time. After each answer, give me feedback on clarity, completeness, and confidence, then move to the next question."

> if Qwen asks "Tell me about a time you handled a difficult customer," you might type something like: "At my old retail job, a customer wanted to return a laptop past the 30-day window. I listened to her frustration, explained the policy, then asked my manager if we could do a store credit as a one-time exception. She agreed, and the customer left happy. I learned that sometimes bending the rules a little keeps a customer for life."

> ➤ "Before the next question, can you give me feedback on my last answer? Specifically: (1) Was it clear and easy to follow? (2) Did I actually answer what was asked? (3) Did I sound confident or unsure? (4) What is the one thing I should improve? Please be honest — I'd rather hear it from you than from a real interviewer."

> ➤ "That feedback helps. Now please ask me a follow-up question about the same story — something a real interviewer would press me on. I'll answer, you give me feedback, and then we move on."

> ➤ "Can you give me a summary of the three strongest answers from this session, plus one thing I should remember to improve next time?"

How to write a polite but firm complaint email that actually gets a reply

> ➤ "Bought a blue wool coat from Example Outfitters on March 4 (order #1234). It arrived March 10 with a torn lining. Emailed support on March 11, no reply. I'd like a replacement shipped by March 25, or a full refund."

> type "I want to draft a complaint email" as your first line, just to confirm the chat is responding.

> ➤ "Please draft a polite but firm complaint email. Facts: [paste your five facts here]. Tone: calm, professional, no sarcasm, no insults. Length: under 150 words. End with one specific ask and a reasonable deadline of 5 business days. Do not invent any order numbers, dates, or names I didn't give you."

> ➤ "Make this sound more confident. I'm being reasonable, but I expect a real answer."

> ➤ "Shorten this to one ask: I want a replacement coat shipped by March 25. Remove the part about the discount."

> follow-up: "Following up on my email below — I haven't received a response. I'd still like a replacement coat shipped by March 25. Please let me know either way."

Translate and Explain a Confusing Email or Letter with Perplexity

> don't type yet — just leave the cursor blinking in the box.

> paste a paragraph like "Sehr geehrte Mieter, wir möchten Sie daran erinnern, dass die Miete für November bis spätestens zum dritten Werktag des Monats zu überweisen ist..." — the exact wording matters, so paste the real thing.

> ➤ "Please translate this email into plain English, list the three things the sender wants me to do, and tell me how urgent it sounds."

> follow-up: "Can you explain 'bis spätestens zum dritten Werktag' like I'm 10 years old?"

> ➤ "Now write a short, polite reply in English confirming I'll pay by Friday and asking for a receipt."

> edit: change "I will remit payment by Friday" to "I'll send the payment by Friday — thanks for the reminder, Maria."

How to Use Gemini Deep Research for Thorough Reports

> ➤ "Act as an expert researcher for a homeowner. Please conduct a deep, multi-step search of the web to investigate current trends in home solar battery storage for residential properties in the US. Focus on average cost ranges, typical lifespan, available government rebates or incentives, and environmentally friendly recycling options. Organize the report with clear headings, bullet points, and provide specific examples where possible. Include links to your sources."

> If Gemini asks, "Are you interested in global statistics or a specific country for the solar battery research?" you might reply:
> ➤ "Focus primarily on American residential properties, especially those in suburban and regional areas, not just cities."

> ➤ "Can you expand on the environmental impact of lithium-ion battery production? Also, please summarize the key points of this report in three bullet points, suitable for a quick presentation slide."

How to Draft a Polite Follow-up Email After a Job Interview

> *Asked about handling an angry customer → I told the café story. Mentioned the new Nashville office opens in March.*

> `Subject: Thank you — Marketing Coordinator interview`

> *Hi Sarah, thank you so much for taking the time to chat with me this morning about the Marketing Coordinator role. I really enjoyed hearing how the team is planning the March launch.*

> *Our chat about managing the café's social media during the weekend rush reminded me how much I enjoy fast-paced, customer-facing work — that's exactly the energy I'd love to bring to your team.*

> *I'm excited about the role and would love to hear how the next steps look. Thanks again, Sarah — have a lovely rest of your week.*

> (final email): *Hi Sarah, thank you so much for taking the time to chat with me this morning about the Marketing Coordinator role. I especially enjoyed hearing how the team is planning the March launch. Our conversation about handling busy weekend rushes reminded me how much I love fast-paced, customer-facing work, and I'd be thrilled to bring that energy to Wynnum Co. I'm excited about the role and would love to hear how the next steps look. Warm regards, Jordan.*

How to practice for a job interview with an AI mock interviewer tonight

> just open the chat — no typing yet.

> ➤ "I'm preparing for a Senior Customer Success interview at a B2B SaaS company. Here is the job description: [paste]. My background: 4 years in support, promoted once, led a small team. Please ask me one question at a time, like a real interviewer would."

> ➤ "Ask me one question at a time. Start with: 'Tell me about a time you had to win back an unhappy customer.'"

> ➤ "Rate my answer from 1 to 5 on clarity, specificity, and impact. Give me one strength and one specific thing to improve. Be honest — I'd rather hear it from you than from an interviewer."

> ➤ "Now run a 5-question mock interview, one question at a time. Wait for my answer before moving on. Mix behavioral and role-specific questions."

> ➤ "Based on this whole session, list the three patterns that hurt my answers most, and the one question I should drill tomorrow."

Polish Your Resume Bullet Points with Le Chat in 10 Minutes

> ➤ *"I helped customers with their bills and account questions at the front desk of a phone store for about two years."*

> full message to send — *"I helped customers with their bills and account questions at the front desk of a phone store for about two years. Please rewrite this as 3 strong resume bullet points a recruiter would notice. Start each one with a strong action verb."*

> ➤ *"Now give me a stronger version with measurable impact. Put any numbers you don't know in brackets like [X%] or [# of customers] — I will fill in the real numbers myself."*

> ➤ *"I like bullet #2 the best. Please make it slightly shorter and swap 'helped' for a stronger verb. Keep it to one line under 25 words."*

> (next round): *"Now do the same for this bullet: 'answered phones and took messages for a real estate office'. Rewrite as one strong resume line with placeholders for any numbers."*

Pull Names, Emails, and Phone Numbers Out of Messy Text with Cohere

> Below the pasted text, add: `Extract every person's full name, email address, and phone number from the text above. Return the result as a plain list with one contact per line, in this format: Name – Email – Phone.`

Use Alexa+ as a Hands-Free Meeting Note-Taker and Task Builder at Work

> say *"Alexa, create a note called Q2 planning sync."*

> ➤ *"Alexa, write this down: Maria will finalize the budget by Friday. John is drafting the comms. Sarah is booking the follow-up."*

> ➤ *"Alexa, pull the action items from my Q2 planning sync note and read them back to me."*

> ➤ *"Alexa, add all of those action items to my Work to-do list."*

> ➤ *"Alexa, draft a follow-up email from my Q2 planning sync note. Keep it under five sentences."*

> ➤ *"Alexa, show my note called Q2 planning sync so I can review it."*

How to build or improve your CV with AI

> 2022–2024 Café Manager, Brew House
> - Led a team of 8 during peak hours
> - Created weekly rosters and handled payroll
> - Reduced waste by 15 % through better ordering
> 2020-2022 Barista, The Daily Grind
> - Prepared coffee and served customers
> - Managed cash register and end-of-day reports
> - Maintained cleanliness and stock levels

> List:
> 2022–2024 Café Manager, Brew House
> - Led a team of 8 during peak hours
> - Created weekly rosters and handled payroll
> - Reduced waste by 15 % through better ordering
>
> Turn the list into a clean, professional CV in American English.
> Use a simple reverse-chronological layout.
> Keep it to one page.
> Use bullet points, not paragraphs.
> Do not add any made-up details.

> Rewrite every bullet point to start with a strong action verb and include a quantifiable result where possible.

> Please rewrite my CV bullet points to specifically match the keywords and skills mentioned in this job description:
>
> [Paste the full job description here, e.g., "Seeking a café manager with proven leadership skills, experience in staff rostering, and a track record of reducing operational costs and food waste."]

> Give me the final CV in plain text so I can easily copy it into Word or Google Docs for final formatting.

How to Ask GitHub Copilot to Explain a Confusing Error Message

> (no typing needed in this step.)

> (no typing needed in this step.)

> `TypeError: Cannot read properties of undefined (reading 'map') at App.render (App.jsx:14)`

> `Explain this error to me like I'm new to coding. What is it actually telling me to fix?`

> `List the 3-5 most likely causes of this error. For each one, tell me what to look at in my code.`

> `Here's the function that produces the error. Which of your suggestions fits this code? function render() { return items.map(...) }`

> `You mentioned "undefined". I don't fully get what that means in this context — can you explain just that word?`

> `Explain this error to me like I'm new to coding, then list the 3 most likely things to check.`

Compare two job offers side-by-side with Le Chat in plain English

> ➤ *"Hello Le Chat, I need help comparing two job offers."* You'll know it worked when the chat area shows a greeting like "How can I help you today?" or "New conversation started."

> ➤ *"Offer A: $95,000 base salary; 15 days annual leave; 1 month notice; $5,000 signing bonus. Offer B: $92,000 base salary; 20 days annual leave; 2 weeks notice; 3% profit share."* You'll know it worked when the copied text appears in the box after you paste (Ctrl+V).

> ➤ *"Compare these two offers side-by-side and explain the pros and cons of each in simple language."* You'll know it worked when you see a two-column table with headings like "Offer A" and "Offer B" and simple sentences underneath each row.

> ➤ *"Can you explain in everyday terms what the different notice periods mean for my life?"* You'll know it worked when the assistant adds the requested row or provides a short, jargon-free explanation beneath the table.

> *No extra prompt needed – just use the copy command.* You'll know it worked when the entire table appears in your notes or email, ready for you to review side-by-side with a coffee.

Quickly Research Any Topic and Summarize Web Pages with Microsoft Copilot in Edge

> Just click the sparkle icon in the top-right corner of Edge.

> In the Copilot sidebar, click the dropdown and choose “Summarize”.

> You’ll see something like:
> - “This page explains how to compost at home.”

> Type: “What are the top 3 benefits of composting mentioned here?”

> Open a new article or research page, and Copilot will prompt: “Want me to summarize this page?”

> Type: “Save this summary to my notes” to store it in your Copilot chat history.

> In the Edge mobile app, tap the Copilot icon, then choose “Summarize”.

Fact-Check Online Claims with Perplexity's Sources

> `Is it true that taking a cold shower every morning can cure the common cold?`

> `Can you find government health organisation statements about cold showers and colds?`

How to Summarize a Long Article or PDF into Key Bullet Points

> type nothing yet — just confirm the chat is open and ready.

> highlight the first few paragraphs of a long article, copy them, and have your clipboard ready.

> paste a 3,000-word article into the chat, or upload a 12-page PDF called *"Q3-strategy-review.pdf"*.

> ➤ *"Please summarize this into bullet points. I want: 1) the main argument or thesis in one sentence, 2) the 5–7 key takeaways as short bullets, 3) any important numbers, dates or names mentioned, 4) one sentence on what the author is asking the reader to do or believe. Keep it under 250 words."*

> ➤ *"That's good. Now give me just the 3 bullets most relevant to a small business owner."* or *"Cut this down to 5 bullets maximum."*

> paste the bullets into a note called *"Reading — June 2026"*, with a one-line note above them saying where they came from.

Translate a product description into Spanish and French with Cohere

> use your real work email so future receipts, API keys, and account-recovery messages all reach you.

> Cohere's "Command" family of models handles everyday translation tasks well; no need to change it for a product description.

> prompt: "Translate the following product description into natural, shopper-friendly Spanish and natural, shopper-friendly French. Label each version clearly. Match the warm, persuasive tone of the original. Here is the description: [your English copy here]"

> of what a good output looks like: "Spanish: ... [paragraph] ... French: ... [paragraph] ..."

> follow-up: "Make the Spanish version about 20% shorter and a bit warmer — like a friend recommending the product. Keep the French version unchanged."

> spreadsheet columns: English | Español | Français — one row per product, so any future edit stays in sync across all three languages.

How to compare two rental listings side by side with weighted pros and cons

> ➤ "Listing A — 2-bed in Eastside, $X/month, 12-month lease. — — Listing B — 1-bed + den in Northgate, $Y/month, month-to-month. Compare these two rentals."

> ➤ "Here are my priorities, weighted 1–5: 1) Commute under 30 min by transit (5). 2) Quiet street, no main road (4). 3) Utilities included in rent (3). 4) Allows my cat (5). 5) In-unit laundry (2). Score both listings against these."

> ➤ "Make a table with these rows: rent, lease length, square feet, utilities included, parking, laundry, pet policy, commute time to downtown, nearest noise source. One column for Listing A, one for Listing B, one for 'verified or missing.' If a fact isn't in the listing, write 'not stated.'"

> ➤ "Now score both listings 1–10 on each priority from Step 2, multiply by the weight, and show the totals. Add a two-sentence recommendation naming which one wins on my priorities and why."

> ➤ "List every fact in the table you marked 'not stated' and anything you had to infer. For each one, write a one-line question I can email the landlord."

> ➤ "PROMPT 1: [paste listings + intro line]. PROMPT 2: [paste weighted priorities]. PROMPT 3: [paste score-and-recommend request]."

How to use Gemini Deep Research to plan a big purchase decision

> when the picker appears, click the dropdown that currently shows "Fast" or "Pro" and switch it to **Deep Research**.

> ➤ "I need a quiet, energy-efficient dishwasher for a small kitchen, budget roughly $700–$1,000 USD, must fit a 24-inch opening and run under 45 dB. I cook most nights for two, sometimes four. Compare the top 5 models available in my country, with pros, cons, and where each one tends to be cheapest."

> hit the "Start Research" button, then make a coffee; the first draft of the report appears inline as it finishes each section.

> open the report, jump to the "Sources" list at the bottom, and click the first three links — if any of them are dead, irrelevant, or from a brand's own press release, treat the matching claim with caution.

> ➤ "Now compare models A and B from your report on noise, energy use, and reliability over five years — which one wins for a household of two who cooks daily?"

> ➤ "Summarize this whole conversation into a checklist: top 3 dishwashers I should see in a store, what to test in person, and what to double-check on the manufacturer's site before I buy."

Catch Up Fast: Summarize Long Chats in WhatsApp and Messenger with Meta AI

> ➤ *"I copied messages from a 40-person family group covering about a day's worth of chat — roughly 30 messages with photos, links and a few voice notes."*

> ➤ *"Here are the messages from my family group from yesterday. Please summarize them in 5 short bullet points, focus on any decisions or plans people made, and ignore the photos. Reply in plain English."*

> ➤ *"Great, but please make it 3 bullet points instead of 5, and add a one-line 'action items' section at the end showing what I might need to reply to."*

> ➤ *"I forward the summary to a private chat with myself called 'Catch-up notes' so I can search it later without scrolling through the original group."*

Summarize a long PDF into 5 bullet points with Claude

> `Summarize this document into 5 key bullet points. Please use simple, plain English.`

Summarize Long Articles into Key Bullet Points with Grok on X

> `Please summarise this news article into 5 clear bullet points, focusing on the main findings: [paste your article link here]`

> `Can you elaborate on the economic implications mentioned in bullet point 2?`

Translate and Explain Any Restaurant Menu with Qwen on Your Phone

> prompt to use after pasting a text menu: "Here's a menu from a small restaurant in Lyon, France."

> ➤ "Translate this menu into English. For each dish, give a one-sentence description, list the main ingredients, and flag any common allergens like dairy, gluten, eggs, nuts, shellfish, soy, or pork."

> ➤ "What does 'confit de canard' usually taste like, and is it ever made with something other than duck?"

Quickly gauge public mood with Hugging Face Sentiment Analysis

> `I absolutely love this new coffee shop! The staff are so friendly and the flat whites are perfect every time.`

Unpack Complex Ideas with DeepSeek's Simple Explanations

> ➤ "Explain quantum physics in simple terms for someone who knows nothing about science."

> ➤ "That's helpful, but can you explain 'wave-particle duality' from the previous explanation even more simply? Use an everyday analogy."

> ➤ "Can you summarize the main points of quantum physics in three bullet points?"

How to Ask Gemini to Explain Your Bank Statement in Plain English

> open the Gemini app, tap "New chat," and you should see an empty text box with the prompt "Ask Gemini" at the bottom.

> ➤ "I'm pasting my bank statement below — it's about 3 pages of transactions."

> ➤ "Please read the statement I just uploaded and give me a plain-English summary of what happened in November. List the top 3 things I spent money on and any fees I was charged. Keep it under 200 words."

> ➤ "Now look only at the charges that happened more than once this month, and anything that looks like a fee or service charge. For each one, tell me the name, the amount, and how many times it appeared."

> ➤ "Group all my spending from this statement into simple categories like food, transport, subscriptions, bills and other. Show me the total for each category and the percentage of my total spending. Then tell me which category grew the most compared to a normal month."

> ➤ "Turn the summary from Step 3 and the fee list from Step 4 into a short email I can send to my bank asking about the $12 monthly fee. Keep the tone polite and include the date and amount for each charge."

How to Research a Medical Procedure or New Medication with Perplexity

> ➤ "Explain laparoscopic gallbladder removal in plain English. I'm not a doctor."

> ➤ "What is metformin? Explain it like I'm smart but not a doctor. What does it actually do in the body, and what are the most common side effects people notice in the first month?"

> ➤ "Give me 6 smart questions I should ask my doctor before agreeing to this surgery. Make them practical — about recovery time, what could go wrong, and what I'd need to arrange at work and at home."

> Pick the Academic or All focus, then type: "What does recent research say about metformin and weight loss? Are there any big studies from 2024 or later?"

> ➤ "You said recovery from this surgery is usually 1–2 weeks. Does that match what major hospital systems say, or is it more commonly 3–4 weeks?"

Research and Draft Short Social Media Captions with Perplexity

> ➤ *"I'm writing a 3-sentence Instagram caption for a small American coffee roastery. What are 5 trending topics in sustainable coffee right now in the US, with sources from 2025 or 2026?"*

> ➤ *"Tell me more about the rise of compostable coffee pods in the US. Give me 3 specific statistics or quotes from American sources I could mention in a post."*

> ➤ *"Write a 3-sentence Instagram caption for our compostable coffee pods. Tone: warm and friendly, not salesy. Mention that 1 in 3 people now look for compostable packaging, according to a 2025 survey. End with a question to invite comments."*

> ➤ *"Make sentence 2 shorter and a bit punchier. Keep the question at the end. Suggest 5 relevant hashtags."*

Compare Products with Perplexity Before You Buy

> `Compare the Google Pixel 8 vs Samsung Galaxy S24 on camera quality and battery life.`

> `Which one is generally considered better for low-light photography according to user reviews?`

> `What are the pros and cons of an OLED TV screen for watching sports, compared to a QLED screen?`

How to Compare Phone Plans and Pick the Right One Without the Jargon

> ➤ "I used 12 GB last month — mostly maps, Spotify, and one Zoom call a day."

> list: "Must-haves: at least 15 GB, hotspot, 5G. Nice-to-haves: international roaming."

> ➤ "I'll compare a Verizon postpaid plan, an AT&T-based MVNO, and a T-Mobile prepaid plan."

> ➤ "$40 plan + $6 taxes + $25 phone = $71/month total. After promo ends: $95/month."

> ➤ "I asked my neighbor: 'How's your T-Mobile signal in our building?' He said: 'Two bars, calls drop.'"

> row: "Plan A: $55/mo, 25 GB, hotspot yes, support fair. Plan B: $45/mo, 15 GB, no hotspot, support good. Pick A — hotspot is a must-have."

Use Copilot to explain code you found online in plain English

> (no typing yet — just opening the chat)

> a real-looking snippet you might paste: `import requests; r = requests.get("https://api.example.com"); print(r.json())`

> ➤ "Explain this code line by line, like I'm 12 years old and have never programmed before."

> ➤ "Can you explain `requests.get(...)` like I'm ordering food at a restaurant?"

> ➤ "Are there any security risks, missing dependencies, or things that could break if I run this on my computer?"

> (no typing — just clicking the copy icon and pasting into a note)

Summarize a long chat thread into 3 bullet points with Copilot

> `Summarize the following chat thread into exactly three short bullet points, focusing on key decisions and action items:`

> `Please make the bullet points even shorter, just focusing on who needs to do what.`

Read the Room on X: Summarize Public Opinion with Grok

> toggle behavior: you turn on **Use X posts**, and the icon turns from gray to blue.

> ➤ *"Summarize what people on X are saying about the new AI safety rules announced this week. Give me the main viewpoints, the most common arguments for and against, and the overall mood."*

> ➤ *"Good. Now add the most common arguments people are making on each side, and tell me the overall mood — positive, negative, or mixed — and roughly what share of posts fall into each."*

> ➤ *"Show me 3 real X posts that support the negative mood you described, with the usernames and links."*

Ask Copilot to compare two product specs and list the differences

> Type or paste:
> ➤ "Hello"

> Paste:
> ➤ "Here are the specs for the first laptop: Processor: Intel Core i7-13700H, RAM: 16GB LPDDR5, Storage: 1TB SSD, Screen: 14" 2.8K OLED 120Hz, Weight: 1.4kg, Price: $2,199"

> Type:
> ➤ "Format these specs as a clean bullet-point list."

> Paste:
> ➤ "Now for the second laptop: Processor: AMD Ryzen 9 7940HS, RAM: 32GB LPDDR5, Storage: 512GB SSD, Screen: 16" 2.5K IPS 165Hz, Weight: 1.8kg, Price: $1,799"

> Type:
> ➤ "Compare these two products and list the differences in a clear table or bullet list."

> Type:
> ➤ "Which of these two products is better for video editing on a budget?"

Summarize Long Research PDFs with Cohere's Clever AI

> Your key will look like a long string of random letters and numbers, such as `a1B2c3D4e5...`

> Paste the text you copied from your PDF into the large input box on the left side of the Playground.

> output:
> - The study tested 500 participants over twelve months.

Create a tiny AI weather-helper widget for your blog with Hugging Face widgets

> `[Enter your Email Address]`, `[Choose a Password]`, `[Click Sign Up]`

> `weather` (typed into the search bar)

> `Chicago` (typed into the city input box)

> `[Click on the </>]`

> `[Click the "Copy" button next to the HTML code]`

> `[Switch to "Text" editor, paste HTML code]`

> `[Click "Update" or "Publish" in your blog editor, then open your live blog page.]`

How to Build a Simple Weekly Schedule in Google Sheets with Gemini

> name it "Spring weekly schedule" or "My week plan".

> type "Hi" as a quick first message to confirm Gemini is responding.

> ➤ "Here's my week. I work Monday to Friday, 9 a.m. to 5 p.m. I teach a yoga class Tuesday and Thursday evenings from 6 to 7:30 p.m. I hit the gym three mornings a week, around 6:30 to 7:30 a.m. on Monday, Wednesday, and Friday. Saturday morning is family breakfast from 9 to 11 a.m. Please build a weekly schedule in this sheet with rows for each day, columns for morning, midday, afternoon, and evening, and a SUM formula that totals the hours automatically."

> in cell B2 (Tuesday morning) type "Gym"; in cell C5 (Wednesday afternoon) type "Deep work block"; in cell A8 (Monday evening) type "Family dinner."

> blue for work, green for fitness, yellow for family time, gray for rest or sleep.

> type `=SUM(B2:B10)` and press Enter. The cell instantly shows the total. If your durations are in seconds, multiply by 24 first: `=SUM(B2:B10)*24`.

> just close the browser tab when you finish. Next time you open the file, your schedule will be right where you left it.

Build a Simple Q&A Chatbot for Your Website with Cohere

```
As a friendly customer service bot for "My Awesome Shop", what is your return policy?
```

```
<div id="ai-world-chatbot"></div>
<script>
const chatbotContainer = document.getElementById('ai-world-chatbot');
chatbotContainer.innerHTML = `
<div style="position: fixed; bottom: 20px; right: 20px; width: 300px; height: 400px; border: 1px solid #ccc; border-radius: 8px; box-shadow: 0 2px 10px rgba(0,0,0,0.1); display: flex; flex-direction: column; background-color: white; z-index: 1000;">
<div style="background-color: #007bff; color: white; padding: 10px; border-top-left-radius: 8px; border-top-right-radius: 8px;">
<h3 style="margin: 0;">My Shop Helper</h3>
</div>
<div id="chat-messages" style="flex-grow: 1; padding: 10px; overflow-y: auto; background-color: #f9f9f9;">
<p><strong>Bot:</strong> Hello! How can I help you today?</p>
</div>
<input type="text" id="chat-input" placeholder="Type your message..." style="padding: 10px; border: none; border-top: 1px solid #eee; border-bottom-left-radius: 8px; border-bottom-right-radius: 8px;">
</div>
`;

const chatMessages = document.getElementById('chat-messages');
const chatInput = document.getElementById('chat-input');
const COHERE_API_KEY = "YOUR_COHERE_API_KEY"; // Replace with your actual Cohere API Key

async function sendMessageToCohere(message) {
const response = await fetch("https://api.cohere.ai/v1/chat", {
method: "POST",
headers: {
"Content-Type": "application/json",
"Authorization": `Bearer ${COHERE_API_KEY}`
},
body: JSON.stringify({
message: message,
model: "command-r-plus", // You can choose a different model
// Add more context or conversation history here for a smarter bot
})
});
const data = await response.json();
return data.text;
}

chatInput.addEventListener('keypress', async function(e) {
if (e.key === 'Enter' && chatInput.value.trim() !== '') {
const userMessage = chatInput.value.trim();
chatMessages.innerHTML += `<p><strong>You:</strong> ${userMessage}</p>`;
chatInput.value = '';

const botResponse = await sendMessageToCohere(userMessage);
chatMessages.innerHTML += `<p><strong>Bot:</strong> ${botResponse}</p>`;
chatMessages.scrollTop = chatMessages.scrollHeight; // Auto-scroll to latest message
}
});
</script>
```

Set up named medication reminders with Alexa+ for yourself or an older relative

> at the sign-in screen, tap **"Already a customer? Sign in"**, then enter your Amazon email and password.

> the prompts look like this — *"Alexa, who am I?"* and *"Alexa, play my Flash Briefing."* Speak them naturally. If your relative has a soft voice, ask them to sit about a foot from the phone.

> type something plain — *"Remind Mom to take her blood pressure pill"*, set it for **8:00 a.m.**, set it to **repeat every day**, and assign it **To: Mom**. Keep the wording short — Alexa+ reads it back aloud, so it should sound natural when spoken.

> ➤ *"Alexa, set a reminder for Dad every evening at 7 p.m. to take his evening pill."* Notice you don't need the app at all for this — the Echo itself can create, list, and delete named reminders.

> ➤ *"Alexa, list my reminders"* — Alexa+ will read every reminder you have set, including who each one is for. Use this to double-check before the day starts.

> under the Routine's "When this happens," pick the reminder you made in Step 3. Under "Add action," pick **Sounds** → **Bell chime** (or whatever chime the app offers). Save the Routine and you're done.

Build a Basic AI Chatbot Brain Using OpenAI's API – Quick Start

> `You are a friendly and helpful customer service assistant for "AI World HQ" Your main goal is to answer questions about general AI topics and guide users to relevant guides on our website. Be concise, polite, and avoid giving medical, financial, or legal advice. If a question is outside your knowledge or scope, politely state that you cannot assist with that specific query.`

> `What is AI?`

Create a personalized workout routine with DeepSeek's multilingual chat

> Open your browser and go to deepseek.com, then find the chat interface

> Hi there, I want to build a simple workout plan

> I'm a complete beginner, I have no weights or gym access, I can train Monday, Wednesday and Friday mornings, and I want to build strength and feel more energetic

> Create a simple 3-day weekly workout plan for me, no equipment, beginner level

> Please give the plan in Spanish

> Format this as a printable table so I can print it

Automate Git Commands with GitHub Copilot Suggestions

> Type: *"How do I stage all changes and commit them with the message 'Added initial project setup'?"*

> Type: *"Create a new branch and switch to it, name it 'feature/contact-page-update'."*

> Type: *"How do I merge my current branch (feature/contact-page-update) back into the main branch?"*

How to Plan a Simple Weekly Workout Routine That Fits Your Real Life

> ➤ "Tuesday 7:00am, Thursday 7:00am, Saturday 8:30am, Sunday 5:00pm (backup)."

> ➤ "Tuesday — strength 25 min • Thursday — cardio 30 min • Saturday — strength 30 min • Sunday — mobility 20 min."

> ➤ "Strength — 1) squats 10×3, 2) push-ups 8×3, 3) plank 20 sec ×3, 4) glute bridge 12×3, 5) backpack row 10 each arm ×3."

> ➤ "Minimum strength = 3×5 squats + 3×5 push-ups. Minimum cardio = 10 min walk around the block. Minimum mobility = 5 min stretch before bed."

> ➤ "Sun check — moving Thursday to Friday because of the team offsite. Sat 8:30am still good."

Organize Your Daily Schedule by Voice with Alexa+

> ➤ "Alexa, remind me to water the plants at 3 PM today."

> ➤ "Alexa, add 'dentist appointment' to my calendar for next Tuesday at 10:30 AM."

> ➤ "Alexa, what's on my calendar for tomorrow?" or "Alexa, what are my reminders?"

> ➤ "Alexa, cancel my 'dentist appointment' for next Tuesday." or "Alexa, delete the 'take out the bins' reminder."

Set Up an Alexa+ Routine to Automate Your Everyday Tasks

> Type `Morning Wake-Up` then tap **Next**.

> Type `Alexa, good morning` then tap **Add**.

> Type `Good morning! Time to start the day.` then tap **Add**.

> Type `Acoustic Sunrise` then tap **Add**.

How to Add an Alibaba Cloud AI Chatbot to Your WordPress Site

> `FrontDeskHelper`

> `Hi! I'm your digital assistant – ask me about opening hours or services.`

> set color to `#0066CC` and upload `logo.png`

> `<script src="https://chatbot.alibabacloud.com/widget.js"></script>` — *your actual code may differ*

> paste the `<script>` snippet exactly as copied