AI Assistants, Explained

A Practical Guide for the Rest of Us

Plain answers about Claude, ChatGPT, and Gemini, for people who have better things to do than debug an app.

Prepared by RAW Insight LLC

Version 1.1 · Last updated November 26, 2026

Contents

PART ONE: IS THIS ACTUALLY SAFE?

Before Anything Else: Is This Safe?

Spotting Scams and Fakes

PART TWO: THE BASICS

What These Things Actually Are

Getting Started: The Universal Steps

Try This Today: Five Low-Stakes Things to Try First

A Few Things Worth Knowing

Talking Instead of Typing

What It Actually Costs

Know Your Assistant

PART THREE: THE BIGGER PICTURE

How Smart Is It, Really?

What About the Headlines Saying AI Threatens Humanity?

Straight from Anthropic: Facts About Claude

PART FOUR: PUTTING IT TO WORK

For Small-Business Owners: Running the Shop Without a Staff

Demo Project: Cleaning Up a Messy Gmail Inbox

Demo Project: Organizing a Google Drive

PART FIVE: WHEN YOU NEED A HAND

Common Worries, Answered Plainly

When Something's Not Working

A Few Words You'll See

Quick-Start Cheat Sheet

How to Use This Guide

This isn't meant to be read start to finish in one sitting. It's organized into five parts so you can jump to whichever one answers today's question and ignore the rest.

> **The short version.** Open the Contents, click the section that matches what you're trying to do right now, read that part, and stop. Come back for the rest whenever the next question comes up. Nobody is grading you on the other pages.

If you only ever read one section, make it "Spotting Scams and Fakes" in Part One. If you only ever do one thing, do one of the five prompts in "Try This Today" in Part Two. It takes about thirty seconds and costs nothing.

Part One: Is This Actually Safe?

The two sections worth reading closely, even if you read nothing else in this guide.

Before Anything Else: Is This Safe?

Short answer: yes, if you use it the way this guide shows you. Here is what actually happens behind the scenes, in plain terms.

- It cannot act on its own. It cannot move money, send an email, or touch your accounts unless you specifically connect that account and specifically tell it to do that one thing. Nothing happens in the background while you are not looking.
- It does not "learn" your secrets and sell them. Anthropic (Claude), OpenAI (ChatGPT), and Google (Gemini) all have privacy policies that govern this, and none of them are in the business of using your bank statement to scam you. You are not the product here in the way a Facebook ad is built around you.
- Connecting an account (like Gmail or a calendar) is like handing someone a key to one specific room, not the whole house. You choose what it can see, and you can revoke that access at any time in your account settings.
- It will show you before it does anything consequential (sending an email, booking something, spending money) and ask you to confirm. It does not quietly click "buy."
- It can be wrong. It cannot be malicious. The realistic risk with these tools is a bad guess or a sloppy draft, not a con artist wearing a robot suit.

> **The one rule that matters most.** If something ever asks you for a password, a social security number, or to "verify" an account outside the app, stop. That is not how these assistants work, and it is a sign someone is impersonating one. Real AI assistants never ask for that.

Spotting Scams and Fakes

This is, honestly, one of the most useful things an AI assistant can do for you, and it cuts both ways, so it's worth taking slowly.

Using it as a second pair of eyes

If a text, email, or letter feels off, paste the actual wording in and ask directly:

"Does this look like a scam? Here's exactly what it says: [paste the message]"

It's very good at spotting the classic warning signs: urgency ("act within 24 hours"), secrecy ("don't tell anyone"), requests for gift cards or wire transfers, and links that don't quite match the company they claim to be from. Treat its answer as a strong second opinion, not the final word. If it says something looks safe but your gut still says no, trust your gut and call the actual company directly using a phone number you look up yourself, not one from the message.

The flip side: AI can now fake a voice

Here is the part worth taking seriously. The same technology can clone a voice from a short clip (sometimes just a few seconds pulled from a social media video) convincingly enough that a phone call "from" a grandchild or family member in trouble, asking urgently for money, can sound completely real even though it isn't.

> **Set up a family safe word.** Agree on a word nobody outside the family would know, to use on the phone if anyone ever calls in a panic asking for money or gift cards.
>
> If you get one of these calls, hang up and call that person back directly on their known number before doing anything else. A real emergency will still be real two minutes from now.
>
> An AI assistant can help you think through whether a message matches a known scam pattern after the fact, but it can't verify a live phone call for you in the moment. That's on the safe word and the callback, not the app.

None of this means the technology is something to fear. It means the same healthy skepticism you'd apply to an unexpected knock at the door now applies to an unexpected voice on the phone, too.

Part Two: The Basics

Everything you need to start a conversation today, with no accounts to connect, nothing to risk.

What These Things Actually Are

Forget the robot-uprising movies. Think of an AI assistant as an extremely well-read employee who is available 24 hours a day, never gets annoyed by a repeated question, and forgets everything about you the moment you close the conversation, unless you tell it to remember.

You type a message. It writes back. You keep going back and forth, the way you would with a smart coworker, until you have what you need. That is the entire mechanic. Everything else in this guide is just detail on top of that.

The three you will hear about most:

- Claude, made by Anthropic.
- ChatGPT, made by OpenAI.
- Gemini, made by Google.

They all work the same basic way. Where they differ, this guide calls it out by name so you can match the instruction to whichever one you actually have open.

Getting Started: The Universal Steps

1. Open the app or website (claude.ai, chatgpt.com, or gemini.google.com) and sign in, or create a free account with an email address.
2. Find the message box at the bottom of the screen. That is the only control you truly need.
3. Type what you want in plain English, the way you would ask a person. "Help me write a birthday card for my granddaughter turning 8" works better than a string of keywords.
4. Read the answer. If it is not quite right, do not start over. Just reply again. "Shorter, please" or "make this warmer" or "she likes horses, can you work that in" all work.
5. When you are done, you can simply close the window. Nothing is saved unless you asked it to remember something, or the app has a memory feature turned on.

That loop (ask, read, refine) is 90 percent of using any of these tools well. Everything below is optional polish.

Try This Today: Five Low-Stakes Things to Try First

Before you connect a single account, get comfortable with the back-and-forth on something that can't go wrong. Pick one of these and just try it. There is no way to mess this up.

- "I have chicken, rice, and whatever's in a normal fridge. Give me something easy for dinner tonight."
- "Can you explain what this means in plain English?" Then paste in a confusing paragraph from a doctor's note, an insurance letter, or a bill.
- "Scale this recipe up to feed 12 people instead of 4." Paste in any recipe and let it do the math.
- "Help me write a birthday text to my granddaughter turning 10. Something warm but not corny."
- "Explain this news story to me like I just walked in. What happened and why does it matter?" Paste in a headline or a short article.

Notice what didn't happen: you didn't need an account connected, a credit card, or any special setting. This is the entire tool at its simplest, and it's also most of what you'll ever actually use it for.

A Few Things Worth Knowing

It only remembers what you tell it to

Claude, ChatGPT, and Gemini all now offer an optional memory feature, off by default in some cases, that lets the assistant recall things across conversations: your granddaughter's name, your preferred tone, an ongoing project. You can turn this on or off in Settings, and you can always ask it directly: "what do you remember about me?" It will tell you honestly.

Uploading files

All three let you attach a document, a photo, or a spreadsheet directly in the message box (usually a paperclip or plus icon). You can hand it a photo of a confusing letter and ask, "what is this asking me to do?"

It can make things for you, not just talk

Ask for a checklist, a schedule, a flyer, a spreadsheet, or a printable one-pager, and it can hand you back an actual file rather than just words on the screen. If you do not specify a format, ask: "can you make that into something I can print?"

Connecting your accounts (optional, and never required)

Each service has a Settings or "Connected Apps" area where you can link things like your calendar or email. This is entirely optional. The assistant is fully useful with nothing connected at all. If you do connect something, you decide exactly what it can do with it, and you can disconnect it in one click at any time.

Talking Instead of Typing

If typing on a phone screen is more trouble than it's worth, skip it entirely. All three assistants let you just talk.

- **ChatGPT:** look for a small soundwave or headphone icon in the message box. Tap it, talk normally, and it answers out loud, and you can interrupt it mid-sentence just like a phone call.
- **Claude:** tap the microphone icon inside the message box on the phone app, speak your question, and it turns your speech into text before answering. You'll see what it heard, so you can fix anything it misunderstood before sending.
- **Gemini:** tap the microphone icon, or use "Gemini Live" for a real back-and-forth spoken conversation, closer to talking on the phone than typing messages.

If your voice gets misread (everyone's does occasionally), just try again a little slower, or switch to typing for that one message. Nothing is lost either way.

What It Actually Costs

The free version of all three is genuinely useful, not a stripped-down teaser. Here's what changes if you pay:

- Claude, ChatGPT, and Gemini are all free to sign up and use with no credit card required.
- What a paid plan typically adds: more messages per day (free plans cap how much you can ask before waiting a few hours), access to their most capable model, faster responses when the service is busy, and the ability to upload more or larger files.
- Connecting accounts like Gmail or Google Drive, the way the demo projects in this guide show, does not cost extra on any of the three. It's a built-in feature, not an upsell.
- Paid plans run roughly $20 a month per service as of this writing, billed and cancelable the same way as any other subscription, directly in that service's account settings, with no phone call or retention pitch required.

Start free. You will likely bump into the daily limit only if you're using it heavily every day, and if that happens, you'll know exactly why you might want to upgrade, rather than guessing.

Know Your Assistant

Same idea, three different personalities. Here's the quick read on each.

> **Claude (Anthropic).** Best known for careful, thoughtful writing and for being cautious about anything that could be unsafe or misleading. Memory and file-creation live under the settings menu (usually a gear or your name in the corner). Visit: claude.ai

> **ChatGPT (OpenAI).** The most widely used and the most flexible for casual back-and-forth. "Custom Instructions" in Settings lets you tell it once how you'd like it to talk to you, so you don't have to repeat yourself. Visit: chatgpt.com

> **Gemini (Google).** Built into Google's own ecosystem, so it connects most smoothly with Gmail, Google Calendar, and Google Docs if you already use those. Look for a "Google Workspace extensions" toggle to link them. Visit: gemini.google.com

Part Three: The Bigger Picture

For when the practical questions are settled and the curious ones start. Read this part whenever you like.

How Smart Is It, Really?

This is the question everyone actually wants answered, so let's not dance around it. The honest answer is: it's smart in a very particular way, and not smart at all in some others, and knowing the difference is what keeps you using it well instead of either dismissing it or trusting it too much.

What it's genuinely good at

- Standardized tests and licensing-style exams. Since around 2023, general-purpose AI models have consistently scored in the top tier on the bar exam, medical licensing-style questions, the SAT, and the GRE, often better than most of the humans who take those tests for a living. This isn't a fluke or a one-time headline; it's held up across every major model since.
- Reading enormous amounts of material fast. What would take you a weekend (a 400-page contract, a decade of tax returns, a stack of medical records), it can read start to finish in under a minute and tell you what matters.
- Writing and rewriting at a professional level. Drafting a letter, tightening a paragraph, matching a tone. This is closer to its home turf than almost anything else you'll ask of it.
- Code, math shown step by step, and translating between languages. These are all areas where there's a clear "right answer" it has seen a version of millions of times before.

Where that intelligence quietly runs out

- It doesn't actually know things the way you do. It predicts the most likely next word based on patterns in enormous amounts of text it was trained on. Most of the time that produces something true and useful. Sometimes it produces something confident-sounding and wrong, with zero change in tone to warn you. This is the single most important thing to remember about all of it.
- It has no lived experience, no common sense built from actually being a person, and no idea what happened after its training cutoff unless it looks it up for you. Ask it who won last night's game and it genuinely doesn't know unless it checks.

- It can't tell the difference between a well-documented fact and a popular myth unless the truth clearly won out in what it was trained on, so on genuinely contested or obscure topics, treat its answer as a first draft, not a verdict.
- It has no judgment about your specific life. It doesn't know your family, your finances, or your health history unless you tell it in that conversation, so its advice is only ever as good as the context you hand it.

So how should you think about it?

Not as a genius, and not as a gimmick, but as a brilliant research assistant with an outstanding memory for facts and zero memory for your life, who never gets tired, never rolls their eyes at a repeated question, and occasionally states something wrong with the exact same calm confidence as something right.

That last part is the whole ballgame. Use it the way you'd use a very sharp but unsupervised intern: fantastic for a first draft, a summary, or a second opinion; not the final word on anything where being wrong actually costs you, like money, health, or a legal deadline. Read what it gives you. If something matters, double-check it. That one habit is worth more than understanding how any of this actually works under the hood.

What About the Headlines Saying AI Threatens Humanity?

Fair question, and it deserves a straight answer rather than either dismissal or panic.

Start with the distinction that gets lost in most coverage: the chatbot on your phone answering questions and the scenario in those headlines are not the same thing. What you're using has no ability to act unless you direct it, no access to anything you haven't connected, and no capacity to do anything at all between the moments you're actively typing to it. This whole guide has been built around that fact. The "threat to humanity" stories are about something else: much more advanced, largely hypothetical future systems, and about how this technology gets governed at a societal scale: things like automated weapons, large-scale cyberattacks, job markets shifting faster than people can adapt, or a handful of companies or governments gaining outsized power from it.

Who actually believes what

This is a genuinely contested question even among the people who build this technology for a living, and it's worth knowing the real range rather than picking one headline as the final word:

- Some prominent AI researchers, including leadership at Anthropic, OpenAI, and Google, have publicly said that as these systems keep getting more capable, the long-term risks deserve serious, sustained attention, up to and including the possibility of AI systems acting in ways their creators didn't intend.

- Many other researchers and technologists think those long-term fears are overstated or premature, and argue the systems available today are nowhere near capable enough for that scenario, while agreeing that nearer-term problems (scams, misinformation, job disruption, biased or wrong answers) deserve the actual attention.
- Governments in the US, UK, EU, and elsewhere have started treating this seriously enough to hold AI safety summits and draft regulation, which tells you the concern isn't limited to any one side of the debate, even as they disagree sharply on what to actually do about it.

This is exactly why Anthropic, the company behind Claude, structured itself the way the next section describes: as a Public Benefit Corporation with a written Responsible Scaling Policy and safety checks that get stricter as each new model gets more capable. Whatever you make of the long-term debate, it's the reason that framework exists at all.

What this means for you, practically

Nothing about this changes how you should use the tool day to day. The same habit this guide keeps repeating (read what it gives you, and check anything that actually matters before you rely on it) is exactly the right amount of caution for the assistant in your hands right now. If the bigger debate interests you, treat it the way you'd treat any other complex, unsettled topic in the news: read a few different sources rather than taking one alarming headline as the whole picture, since credible, informed people land in very different places on this one.

Straight from Anthropic: Facts About Claude

Since Claude is one of the three assistants this guide covers, here is what its maker, Anthropic, actually says about it: plain facts, not marketing copy, so you know exactly what you're working with.

The company behind it

- Anthropic was founded in 2021 in San Francisco by siblings Dario and Daniela Amodei, along with several colleagues who had previously worked on AI safety research elsewhere.
- It is structured as a Public Benefit Corporation, a company legally required to weigh the public good alongside profit, with an independent Long-Term Benefit Trust that oversees whether it's living up to that.
- Anthropic's stated mission is to build AI that is "helpful, honest, and harmless," and to treat AI safety as a rigorous science rather than an afterthought.

How big a conversation it can hold

- Claude can read and remember roughly 200,000 "tokens" in a single conversation. Tokens are small chunks of text, and 200,000 of them works out to somewhere around 500 pages. In practice, that means you can hand it an entire long contract, a full book manuscript, or years of email and it can still refer back to the beginning.
- Certain versions can go even further, up to 1 million tokens (around 2,500 pages) for especially large jobs.

What it can actually read and produce

- Documents: Word files, PDFs, plain text, and web pages.
- Spreadsheets: Excel files it can read, calculate on, and hand back.
- Presentations: PowerPoint files, built or edited.
- Images: photos and scanned pages. It can describe what's in a picture or read text out of a photographed letter.
- Computer code, if that's ever relevant to your work.

Where you can use it

- A website (claude.ai), desktop programs for Mac and Windows, and phone apps for iPhone and Android.
- A browser extension (Claude in Chrome) and a version built into Microsoft 365 for those who live in Word, Excel, and Outlook.
- Optional connections to outside tools (Gmail, Google Drive, Slack, and others) that only work once you turn them on yourself, exactly as this guide's demo projects walk through.

How it was trained to behave

- Anthropic uses a method it calls "Constitutional AI." Instead of leaving Claude's sense of right and wrong purely to trial and error, it is trained against a written set of principles about how to act helpfully and safely, similar in spirit to a code of ethics.
- Anthropic publishes a "Responsible Scaling Policy" that lays out graduated safety checks new Claude models must clear before release, with the checks getting stricter as the models get more capable.

Who it's for, and who it isn't

- Claude is restricted to users 18 and older. There is no version built for or marketed to children.

- On the other end, Anthropic has been expanding into education. Claude is fully deployed across entire campuses at universities including Northeastern, the London School of Economics, and Champlain College, and Anthropic's Teach For All partnership has trained more than 100,000 educators in 63 countries.

None of this makes Claude infallible. It is still a tool that can misread something or get a detail wrong, which is exactly why this guide keeps coming back to the same habit: read what it gives you, and check anything that matters before you rely on it.

Part Four: Putting It to Work

Two real, start-to-finish projects, plus a translation guide for anyone running a business solo.

For Small-Business Owners: Running the Shop Without a Staff

If you run a small business on your own, you don't need to hire an assistant, a project manager, and a scheduler. You need to know which questions to ask an AI assistant to get much of the same result. Here is the direct translation:

- **Instead of an assistant:** "Draft a polite email to this vendor asking for a revised quote by Friday," or paste in a rambling voicemail transcript and ask, "turn this into three clear action items."
- **Instead of a project manager:** "Here are five jobs I have going. Put them in order of what's most urgent and tell me what I'm missing," or ask it to build you a simple weekly checklist from a list of tasks.
- **Instead of a scheduler:** paste your week's commitments and ask, "where are my gaps, and when should I call this client back?" If you connect your calendar (Google Calendar works well with Gemini, and Claude and ChatGPT can connect to it too), it can check real availability for you.
- **Instead of a bookkeeper's first pass:** upload a spreadsheet or a photo of receipts and ask it to organize totals by category, then still hand the final numbers to your actual accountant.

The pattern in every case: hand it the messy raw material (an email, a voicemail, a list, a spreadsheet) and ask for the organized version. That is the job it is good at.

Demo Project: Cleaning Up a Messy Gmail Inbox

This is a real, start-to-finish project you can do this afternoon. It shows you exactly what "connecting an account" looks like and proves to yourself that nothing scary happens when you do it.

What you are about to do

You are going to let an AI assistant look at your Gmail inbox (just look, at first) and help you sort years of clutter into a system, without deleting a single thing until you say so.

Step 1: Connect Gmail (this is the only new part)

1. Open your assistant (Claude, ChatGPT, or Gemini) in a web browser, not the phone app. Connections are easier to set up on a computer.
2. Look for a settings icon (usually a gear, or your name or photo in a corner) and find a section called "Connected Apps," "Integrations," or "Extensions." Gemini calls this "Google Workspace extensions," and since Gmail is a Google product, it is often already switched on there.
3. Find Gmail in the list and turn it on. You will be sent to a familiar Google sign-in screen. This is Google itself confirming the connection, not the AI company.
4. Google will show you exactly what it is about to allow, for example, "read, compose, and organize your email." Read it, then approve it. You are not giving away your password; you are approving a specific, revocable permission.

Step 2: Ask it to look before it touches anything

Type something like this into the message box:

"Look through my inbox and tell me the biggest categories of clutter you see: old newsletters, receipts, one-time notifications, things from people I actually know. Don't change anything yet, just report back."

It will come back with a plain-English summary, something like "you have 400+ promotional emails from 30 different senders, 60 shipping confirmations, and a stack of unread newsletters going back two years." This is the moment you realize it actually looked, and that looking did nothing destructive.

Step 3: Build a simple system together

Now ask for a plan, not action:

"Based on that, suggest 4 or 5 folders or labels I could use to keep this organized going forward. Keep it simple, I don't want more than five."

A realistic result: Bills & Receipts, Family, Newsletters, Shopping, and Everything Else. Adjust the names until they sound like your life, not a corporate filing cabinet.

Step 4: Let it do the sorting, with a leash on

Only now do you ask for action, and you ask for it in a small, reversible batch:

"Create those as labels in my Gmail, then apply them to my last 50 emails as a test. Don't delete or unsubscribe from anything. Just label. Show me what you did when you're finished."

It will report back with what it labeled. Open Gmail yourself and look. If it is right, say "good, now do the rest of my inbox the same way." If something is off, say so ("that newsletter from the hardware store should be Shopping, not Family") and it will adjust and remember the correction for the rest of the cleanup.

Step 5: The optional, satisfying part

Once you trust the sorting, you can ask it to go further, always one category at a time:

- "Show me every promotional email older than a year so I can decide whether to unsubscribe. Don't unsubscribe me from anything automatically."
- "Draft an unsubscribe click-through list for the ones I approve, but don't send anything until I say go."
- "Find any emails that look like bills or important account notices from the last month and flag them so I don't miss them."

Notice the pattern: it proposes, you approve, then it acts. That is the whole safety model in one sentence, and it works exactly the same way for every task you hand it, in every assistant.

Undoing it, if you ever want to

Gmail labels can be renamed or removed in a couple of clicks inside Gmail itself, and nothing about this process deletes an original email unless you specifically asked it to. You can also revoke the entire connection at any time from the same "Connected Apps" screen where you turned it on. One click, and the assistant can no longer see your inbox at all.

Demo Project: Organizing a Google Drive

Same idea, different mess: a Google Drive full of years of documents, photos, and files with names like "Scan001" and "Final_Final_v3."

Step 1: Connect Google Drive

1. In your assistant's Connected Apps or Extensions settings (the same place you found Gmail), locate Google Drive and turn it on.
2. Approve the Google permission screen. Again, this is Google confirming access, and you can see exactly what is being allowed before you approve it.
3. If you use Gemini, Drive is often available automatically since it is a Google product; Claude and ChatGPT connect to it as a separate extension.

Step 2: Get the lay of the land

Start with a look, not a change:

"Look at my Google Drive and tell me roughly what's in there: how many files, what kinds, and whether you notice a lot of duplicates or oddly named files. Don't move or delete anything yet."

Step 3: Propose folders, don't impose them

"Suggest a simple folder structure for me based on what you found, something like Taxes, Family Photos, Home & Insurance, and Old Work Stuff. Keep it to 5 folders or fewer."

Rename anything that doesn't match how you actually think about your own files. There is no wrong answer here except one that doesn't make sense to you.

Step 4: Sort in a small, visible batch first

"Create those folders, then move just the files from the last 3 months into the right ones as a test, and show me a list of what you moved."

Check the result in Drive itself. Once it looks right, say "good, now do the rest, oldest files last since those are riskiest to misfile."

Step 5: Handle duplicates and clutter carefully

This is the one spot worth extra caution, because deleting is harder to walk back than labeling an email:

- "Find likely duplicate files (same name or same content) and just list them for me. Don't delete anything."
- Review that list yourself. Anything you don't recognize or aren't sure about, leave alone.
- Only after reviewing: "Move these specific duplicates I've approved to a folder called Review Before Deleting, instead of deleting them outright."

That extra folder is your safety net. It keeps anything questionable one step away from gone, where you can empty it yourself in a week once you're confident nothing was missed.

The rule that makes both projects safe

In both demos, the assistant never deletes, sends, or finalizes anything on the first pass. It looks, it proposes, you approve a small test batch, you check the real result yourself, and only then does it scale up. That is not a special setting you have to remember to turn on; it is simply how you should phrase every request like this, in any assistant, for any account you ever connect.

Part Five: When You Need a Hand

Quick answers, fixes, and a glossary: the part you'll come back to most.

Common Worries, Answered Plainly

"Will it steal my identity if I tell it personal things?"

No. It has no capability to act on your information the way a scammer would. Its purpose in that conversation is to answer you, not to file the information away for misuse.

"Will it lock me into a subscription I can't escape?"

All three have a genuinely usable free version. Paid tiers exist for heavier use, and you can cancel them the same way you would any other subscription, in account settings.

"What if I break something?"

You cannot. There is no button in a chat conversation that deletes a file, empties a bank account, or breaks your computer. The worst outcome of a wrong click here is a message you did not want to send, and even that requires your confirmation first for anything connected to a real account.

"Is it listening to me all the time?"

No. It only sees what you type or upload into that conversation, and only while you are actively using it.

When Something's Not Working

Before you call your kid, try these. They solve almost everything.

"I can't remember which app or how to log back in"

Look for the app icon on your phone's home screen, or type the website directly into a browser: claude.ai, chatgpt.com, or gemini.google.com. Sign in with the same email you used the first time. If you're not sure which one, try the email you check most often first.

"It's frozen" or "nothing's happening"

Close the app completely and reopen it, or if you're on a website, refresh the page. This fixes the great majority of freezes, the same as with any app.

"It stopped in the middle of an answer"

Just type "keep going" or "you got cut off, please continue," and it will pick up where it left off. This happens sometimes on very long answers and is not an error on your part.

"It didn't understand what I said" (using voice)

Try again a little slower in a quieter room, or switch to typing that one message instead. Background noise and speaking quickly are the usual culprits, not you doing anything wrong.

"The answer feels wrong or off"

Say so, plainly, the way you'd correct a person: "that's not right" or "that doesn't sound like me, try again." It will not get offended, and it genuinely does better on the second try when you tell it what missed the mark.

A Few Words You'll See

None of these are required reading. They're here so a word you don't recognize never has to stop you cold.

- **Prompt:** the plain-English message you type or say to the assistant. Asking a question is a prompt; there's no special way to phrase one.
- **Chatbot / assistant:** other names for the same thing this whole guide is about: Claude, ChatGPT, and Gemini are all chatbots, or AI assistants.
- **Model:** the specific "version" of the assistant doing the answering, the way a car has a model year. Newer models are generally more capable, but the free version of any of them is still genuinely useful.
- **Hallucination:** the industry's word for when it states something wrong with total confidence. Worth knowing the term, since you'll see it in news coverage; the fix is always the same one this guide keeps repeating: check anything that matters.
- **Memory:** the optional feature that lets it recall things about you across separate conversations, off unless you turn it on.
- **Connector / integration:** the technical name for linking an outside account, like Gmail or Google Drive, the way the demo projects in this guide walk through.

Quick-Start Cheat Sheet

Keep this part taped by the keyboard if nothing else.

1. Type what you want in plain language.

2. If the answer is off, just tell it what to change. Do not start over.
3. It never touches money, email, or accounts without connecting them first and confirming with you every time.
4. Nothing is remembered unless you turn memory on or ask it to remember something.
5. If anything asks for a password or a code, that is not the real assistant. Stop and ask for help.
6. For any connected account: ask it to look first, propose a plan second, test on a small batch third, and only scale up once you've checked its work.