

AUTUMN 2026
EDITION

The Prompt Book

*A field guide to asking
AI better questions.*

DREAMPROMPTER

FREE

Contents

00	Start here	3
	The whole guide in sixty seconds, with DREAM and WAKE on one page.	
01	What you're talking to	5
	A working picture of an AI assistant: capable, new to your life, sometimes wrong.	
02	Spells that don't work	7
	Six famous prompting tricks, and what the tests and the model makers say about them.	
03	Anatomy of a dream prompt	10
	The five parts of a strong prompt, and how many you need.	
04	D: Destination	12
	How to say what you want, name its shape, and ask for one thing.	
05	R: Reality	14
	Why your real situation and real documents matter more than wording.	
06	E: Example	17
	How to show what good looks like, and say what to copy.	
07	A: Ask	19
	How to let the assistant interview you before it answers.	
08	M: Measure	21
	How to stop the assistant from simply agreeing with you.	
09	After the prompt	23
	What to say after the first answer, and when to start over.	
10	Wake	25
	Four checks for reading any answer before you act on it.	
11	Health, money, law and the heart	27
	How to use an assistant to prepare for a professional, never to replace one.	
12	Privacy	31
	What never to paste into a chat, and three settings worth finding.	
13	Any model	33
	Four switches that change what an assistant can do, and prompts that work anywhere.	
14	Field recipes	35
	Twelve annotated prompts for twelve parts of life, and six weak ones to fix.	
15	When not to prompt	41
	What to keep for yourself, what to share, and what to hand over.	
16	Practice and library	43
	A seven-day practice plan, a prompt library page, a glossary and the sources.	

Start here

The whole guide in sixty seconds, with DREAM and WAKE on one page.

Two people open the same assistant on the same Sunday evening. Both want help with dinner for the week.

The first types: “meal ideas.”

The second types this:

PROMPT

I want a dinner plan for Monday to Friday for two adults and a six-year-old who won't eat anything spicy. I have 30 minutes a night and about \$90 for the week. Give me a table with the dish, the time it takes and the main ingredients, then one shopping list grouped by aisle. Tell me what you assumed about my kitchen.

The first person gets a list that could have come off the back of a cereal box. The second gets a plan they can shop from tonight. Same assistant, same minute, no secret words. The only difference is what each of them told it.

A good question says where you want to end up and what your life looks like.

The sixty-second version

- 1 Say what you want to walk away with, and its shape.
- 2 Describe your real situation. Paste the real email, numbers or draft.
- 3 Show an example if you have one.
- 4 For anything personal, let the assistant question you first.
- 5 Ask what it assumed and where it is unsure.
- 6 Check the one thing that matters before you act.

You won't need all six every time. OpenAI's guide for ChatGPT users says you need no formula and that a short prompt is often enough. Treat the six as a checklist for when an answer disappoints.

How to use this guide

Read it in order once. After that, go back to the worksheets and to the recipes in chapter 14. Any prompt set apart in a quote block is yours to copy. Words in [square brackets] are blanks for your own details.

Read chapter 11 before you ask about health, money or law. If something is urgent, skip the chat and call your local emergency number.

This is the Autumn 2026 edition. Advice about AI ages fast, so we kept to what the model makers and the researchers who test them agree on.

WORKSHEET

The one-page cheat sheet

DREAM is for the prompt.

	PART	ASK YOURSELF	STARTER
D	Destination	Could a stranger tell when the job is done?	"I want to end up with..."
R	Reality	Would a friend who has never met me have enough to go on?	"Here's my situation..."
E	Example	Did I show it, or only describe it?	"Here's the kind of thing I mean..."
A	Ask	Does it know it may question me?	"Before you answer, ask me..."
M	Measure	How will I know this is good, and how sure is it?	"Before you finish, tell me..."

WAKE is for the answer.

	CHECK	DO THIS
W	Where from	Ask for sources. Open at least one.
A	Assumptions	Ask what it assumed about you.
K	Kick the tires	Check the one claim your decision rests on.
E	Escalate or edit	Take big stakes to a qualified person. Rewrite the rest in your own words.

Quick asks need D alone. Personal ones need D plus R. Anything you will act on gets all five.

The question I most want a better answer to: _____

What you're talking to

A working picture of an AI assistant: capable, new to your life, sometimes wrong.

Picture a temp who arrives at your office at nine. She has read more than anyone you know: law, cooking, tax rules, poetry, plumbing. She writes quickly and well. She has also never met you, has never seen your inbox, and has no idea why you look tired. You say, “Sort out the Henderson thing,” and walk off.

She will do something. It will be neat and confident. It will probably not be what you wanted.

That is roughly what happens each time you type a short line into an AI assistant. The companies that make them use the same picture. Anthropic’s guidance says to think of its model as a brilliant new employee who lacks context on how you work. OpenAI’s guide for workplace users describes the prompt as a handoff to a smart intern in their first week. Ethan Mollick, a co-author of Wharton’s prompting studies, suggests a patient new coworker who forgets everything between chats.

Three things to hold in mind

It is capable. It explains, drafts, compares and summarizes well, on almost any topic you bring.

It doesn’t know you. It knows a great deal about budgets and nothing about yours. Unless you tell it, or it has a memory of earlier chats, it answers an average person who does not exist. Even memory is partial. OpenAI’s help page says memory does not keep every detail of every conversation.

It is uneven, and it sounds the same either way. It can be right about something hard and wrong about something easy, in the same calm voice. Microsoft’s Copilot guide notes that the same prompt can return different answers. Confidence here is a writing style, not a sign of accuracy.

What follows from this

None of the fixes are technical. They are what you would have done for the temp if you had stopped at her desk for two minutes: say where the job ends, tell her what she can't know, show her what good looks like, let her ask, and ask how sure she is. Chapters 4 to 8 take those one at a time.

WORKSHEET

It knows, it doesn't know

Check every line the assistant cannot know unless you tell it or show it.

- ☐ How a fixed-rate mortgage works in general
- ☐ What your lease says about pets
- ☐ Common reasons people put off hard tasks
- ☐ Why you are putting off this one
- ☐ What your sister meant by her last text
- ☐ How much you can spend on the trip
- ☐ What you told a different app last week

Answer: five of the seven, every line with "you" or "your" in it. Those lines are your prompt.

Spells that don't work

Six famous prompting tricks, and what the tests and the model makers say about them.

A coworker forwards you a screenshot with the caption “this changed everything.” The prompt in it reads:

“Act as a world-class career coach with 20 years of experience. Think step by step. I will tip you \$200 for a perfect answer. THIS IS VERY IMPORTANT TO MY CAREER. Tell me how to get promoted.”

It looks like expertise. It is closer to a rain dance. The answer comes back generic, because nothing in those five sentences says who is asking.

In 2025, researchers at Wharton published four Prompting Science Reports that put the best-known tricks through repeated trials. Here is what they found.

The costume

“Act as a world-class expert.” The fourth Wharton report tested expert personas on six models using graduate-level benchmark questions. Matching the expert to the subject brought no significant gain in accuracy, with one model as the exception. Low-knowledge personas, such as a young child, often made answers worse. OpenAI’s page on personalizing ChatGPT says a personality changes style and tone, not capability.

A role is still useful when it is the actual job. “Play the hiring manager and interview me” is a rehearsal. “Write this the way a friendly neighbor would” sets a voice. Neither makes the answer more correct.

The bribe and the threat

“I’ll tip you \$200.” “My job depends on this.” The third Wharton report tried both tipping and threatening. In general, neither had a significant effect on benchmark scores.

Real stakes are still worth stating, as information. “This goes to my landlord tomorrow” changes what a good answer looks like.

Please and thank you

The first Wharton report found that politeness sometimes helped and sometimes hurt, and concluded that no prompting formula is universally valuable. So be yourself. Manners are a choice about who you want to be, not a technique.

“Think step by step”

The second Wharton report found small average gains on models that don’t reason on their own and little effect on models that already do, at a cost in time. OpenAI’s guidance on reasoning models says the phrase is unnecessary and can sometimes hurt. If you want more care, pick the slower thinking mode in your app. If you want to learn, ask for an explanation.

Shouting

ALL CAPS, “CRITICAL,” “NEVER.” Anthropic’s prompting guidance says its recent models can overreact to forceful wording, and advises dialing it back. OpenAI’s developer guide for GPT-5.5 says to keep words like ALWAYS and NEVER for true absolutes, such as safety rules. State one or two real limits, calmly, with the reason.

The mega-prompt

Two thousand words of roles, modes and rules. OpenAI’s GPT-5.5 guide says the model works best when a prompt defines the outcome, and that carrying every instruction over from an old prompt can add noise and make answers mechanical. Anthropic’s blog on prompting describes the same shift: less scaffolding, more of the right context.

What’s left

Take the costume, the bribe and the shouting off that first prompt, and add the one thing it lacked:

 PROMPT

I’ve been a senior nurse on a surgical ward for four years, and I want to move into a charge nurse role within a year. My manager calls me dependable but hasn’t seen me lead. I can give this about two hours a week. Give me the three moves most likely to change how I’m seen, in order, with a first step for each. Tell me what you’re assuming about my hospital.

OpenAI's guide for workplace users compresses the lesson into five words:
"Optimize for clarity, not cleverness."

WORKSHEET

Myth or method

Mark each line Myth or Method.

- 01 "You are a world-class nutritionist." _____
- 02 "Here are the five dinners I actually cook." _____
- 03 "I'll tip you \$50 for a great answer." _____
- 04 "This is for my landlord, and it goes out tomorrow." _____
- 05 "Think step by step." _____
- 06 "Explain your reasoning so I can learn from it." _____
- 07 "NEVER USE JARGON." _____
- 08 "Use everyday words. My reader is new to this." _____
- 09 "Ask me up to five questions before you answer." _____
- 10 "Write it the way a friendly neighbor would." _____

Answers: 1, 3, 5 and 7 are myths. The rest are methods, because each tells the assistant something true about you, your reader or the result you want.

Anatomy of a dream prompt

The five parts of a strong prompt, and how many you need.

In June, Maya will sit down to dinner with her partner's parents, who speak only Spanish. Her last Spanish class was 20 years ago. Here is what she types.

 PROMPT

I want a four-week plan to get my rusty Spanish good enough to chat with my partner's parents over dinner in June. Give it to me as a weekly table I can stick on the fridge. I studied Spanish in school 20 years ago. I can read a menu, and I freeze when someone speaks to me. I have 20 minutes on weekday mornings and a long train commute. I like plans written like this: "Tue: 10 food words, say each one aloud twice." No pep talk. Before you write it, ask me up to four questions, one at a time, if the answers would change the plan. When you're done, tell me what you assumed and which week is most likely to fall apart.

Nothing in it is clever. Look at what each sentence does.

	PART	IN THIS PROMPT	THE JOB IT DOES
D	Destination	“A four-week plan... a weekly table I can stick on the fridge.”	Names the result and its shape
R	Reality	“In school 20 years ago... 20 minutes on weekday mornings.”	Gives the facts and limits
E	Example	“Tue: 10 food words, say each one aloud twice.”	Shows what good looks like
A	Ask	“Ask me up to four questions, one at a time...”	Lets the assistant question you first
M	Measure	“What you assumed and which week is most likely to fall apart.”	Makes it check itself and show doubt

Those five parts spell DREAM. One note on the fourth: Ask means ask back. It is the assistant asking you, not you stating your request. Your request is the Destination.

How much do you need?

Not all five, most of the time.

Quick: D alone. For small asks, a clear destination is plenty.

PROMPT

Give me five subject lines for an email inviting my neighbors to a street cleanup this Saturday morning. Friendly, under eight words each, and one of them a little funny.

Personal: D plus R. The moment the answer depends on you, add your situation.

PROMPT

Suggest three weekend trips within two hours of Denver for two adults and a dog who gets carsick. We have \$600. Give me a short table with drive time, cost and the catch.

Deep: all five. For decisions, sensitive topics and anything you will act on, use the full set, as Maya did.

Where this comes from

DREAM is our name. The ingredients are everybody's. OpenAI's guide for ChatGPT users lists Goal, Context, Output and Boundaries. Google's Workspace guide lists Persona, Task, Context and Format. Microsoft's Copilot guide lists Goal, Context, Source and Expectations, and says only the goal is required.

We left persona out on purpose (see chapter 2). A voice or point of view belongs in the Destination: "written for a reader who has never done this."

WORKSHEET

DREAM builder

D. I want to end up with, in this shape: _____

R. My situation, my limits, and material I'll paste: _____

E. An example, or what I don't want: _____

A. It may ask me this many questions first: _____

M. Before it finishes, it should tell me: _____

D: Destination

How to say what you want, name its shape, and ask for one thing.

You get into a taxi and say, “Somewhere nice.” The driver is skilled, knows every street and wants to help. You will still end up somewhere you didn’t choose.

This prompt is the same request.

WEAK

“Help me plan my mom’s 70th birthday.”

Here is the version with an address on it.

STRONG

I want to end this chat with a one-page plan for my mom’s 70th birthday lunch: a guest list of about 20, a simple timeline for the day, and a budget table that stays under \$800. It’s done when every job on the list has a name next to it. Start with the budget.

The strong version says what the finished thing is, how big it is, and how to tell when it’s finished. That is Destination.

Result first, steps second

Most people describe a process: “First research venues, then compare prices, then draft the invitations.” You rarely need to. OpenAI’s guide for ChatGPT users says to start with the result, and to describe a process only when the process itself matters. The assistant is good at finding a route. It cannot guess where you are going.

A useful test: could a stranger tell when the job is done? “Help with my budget” fails it. “A table with eight categories where every dollar of my pay is assigned” passes.

Name the shape

A destination has a shape: what it is, who it's for, how long, how it's laid out. If the shape matters, say it. Anthropic's guidance adds a rule: tell the model what to do instead of what not to do.

WEAK

"Explain how a heat pump works. Not too long, no jargon."

 STRONG

Explain how a heat pump heats a house to someone who has only ever had a gas furnace. About 150 words, everyday language, one comparison to something in a kitchen. End with three questions to ask an installer.

One finished thing

Ask for five things at once and you get five thin answers. OpenAI's guide for workplace users gives the rule as one prompt, one deliverable. For anything large, work in stages.

WEAK

"Plan my move: movers, budget, packing, address changes and a goodbye party."

 STRONG

Let's plan my move in stages. Stage one only: a packing schedule. I'm moving out of a two-bedroom apartment on [date], I work full time and my weekends are free. Give me a week-by-week table, room by room. Stop there. I'll ask for the next stage when this one looks right.

WORKSHEET

Done means

Finish two or three of these before you type.

I want to end up with _____

It's for _____

It should be about _____ long, laid out as _____

It's done when _____

Start with _____ and stop there.

R: Reality

Why your real situation and real documents matter more than wording.

A 26-year-old has just been put in charge of the kitchen where she was a line cook last month. A 54-year-old runs a team of 40 accountants and has been told his people are afraid of him. On the same morning, each of them types the same seven words.

WEAK

“How can I be a better manager?”

The assistant can’t tell them apart, so both get advice that is true and useless: communicate clearly, give feedback, build trust.

Reality is the part of the prompt that makes the answer yours.

STRONG

I became a manager four months ago. I lead five people who used to be my peers at a 60-person logistics company, and two of them applied for this job. My boss wants weekly reports I don’t have time to write. I have no training budget and about an hour a week for this. Given all that, name the one place to put that hour first, and say why. Tell me what you’re assuming about my team.

What to put in

Think of what a smart friend who has never met you would need.

- Who you are, in one line.
- The facts: numbers, dates, people, places.
- The limits: time, money, energy, rules.
- What you’ve already tried.
- Why it matters. Anthropic’s guidance says that giving the reason behind an instruction helps the model understand your goal.

Sander Schulhoff, lead author of The Prompt Report, a large survey of prompting techniques, calls added context underrated.

Bring the real thing

Don't describe the email. Paste the email. Don't summarize the lease. Paste the clause.

WEAK

"Write a reply to the airline about my lost bag."

 STRONG

Below are the airline's email and my original claim. **[paste both]** Using only these two texts, draft a reply under 150 words asking them to confirm the compensation amount and the date I'll receive it. Quote the line of their email you're relying on. If anything you'd need is missing, tell me instead of guessing.

Two habits are tucked in there. The material comes before the question. With long documents, Anthropic reports that this order improved answers by up to 30 percent in its tests, and Google's developer guidance recommends it too. And "using only these two texts" keeps the assistant inside what you gave it.

Before you paste anything, read chapter 12 on what to leave out.

Relevant, not long

Reality is not your life story. OpenAI Academy's prompting lesson warns that too much extra information can make an answer less helpful. Test each sentence: would the answer change without it? If not, cut it.

Write it once

You will find yourself typing the same facts again and again. Save them.

 PROMPT

Help me write a short "about me" card to paste at the top of any AI chat. Ask me up to eight questions, one at a time, about my work, my household, what I'm working toward, my limits on time and money, and how I like answers written. Then write the card in first person, under 150 words. Leave out anything a stranger wouldn't need, and point out any detail I might prefer to keep private.

About-me card

I am: _____

I live with or look after: _____

This season I'm working toward: _____

My limits. Time: _____ Money: _____

Other: _____

I've already tried: _____

I like answers that are: _____

Always leave out: _____

E: Example

How to show what good looks like, and say what to copy.

Nobody sits down at the hairdresser and says, “Shorter, with some movement, modern but not too modern.” They hold up a photo. The photo does in two seconds what the adjectives never manage.

A prompt works the same way. “Warm but professional” could mean a hundred things. One message you actually wrote means one.

WEAK

“Write a thank-you note to my daughter’s teacher. Make it sound like me.”

STRONG

Write a thank-you note to my daughter’s second-grade teacher for the end of the school year. Here are two messages I wrote that sound like me: [paste two messages]. Match their length and their plain, slightly dry tone, but don’t reuse their phrases. Not like this: “Words cannot express our gratitude.” Under 90 words, and mention the volcano project.

Say what to copy

An example is strong, and that is the catch. Anthropic’s blog on prompting warns that recent models pay very close attention to the details of an example. Hand over a sample without comment and you may get your own sentences back with the nouns swapped.

So add one line on what to take and what to leave. “Match the length and tone. Don’t reuse the phrases.” Or: “Copy the layout of this table, not the content.”

A bad example is as useful as a good one. “Not like this:” followed by one sentence you would never say draws a clear line.

How many, and when

Google’s developer guidance recommends always including examples. Anthropic’s documentation suggests three to five varied ones, while its blog says to start with one. OpenAI’s guidance for reasoning models says to try without examples first.

Our rule: add an example when style or format matters, or when the first answer missed. One is usually enough. A small factual question needs none.

WORKSHEET

Style capture

- ☐ Collect three things you wrote that sound like you: a text, an email, a post.
 - ☐ Run the prompt below.
 - ☐ Correct the description. (“No, I never open with a joke.”)
 - ☐ Save the five bullets next to your about-me card.
-

PROMPT

Here are three things I wrote that sound like me: [paste three samples].
Describe my writing style in five plain bullet points: sentence length, word choice, tone, habits, and anything I never do. Then write one sentence in my style and one that isn’t, so I can check you’ve understood.

What I want copied from my examples: _____

What I never want copied: _____

A: Ask

How to let the assistant interview you before it answers.

WEAK

“What’s my biggest blind spot?”

Back comes something like this: “You may be holding yourself back out of a fear of failure, and you tend to put others’ needs before your own.”

It feels uncanny for a second. Then you notice it would fit your neighbor, your boss and most of your family. It is a horoscope. Your friend’s assistant had a year of her chats to draw on. Yours had nothing, so it described everyone.

The fix is the fourth part of DREAM. Hand over the microphone and let the assistant question you before it answers.

STRONG

I want to know where I’m getting in my own way right now. You don’t know me yet, so don’t answer. Ask me up to six questions, one at a time, each one building on my last answer. Then tell me the one pattern you see, the evidence from what I said, and one small thing to try this week. Don’t flatter me. Say what you can’t know from a short interview.

Why it works

For personal or fuzzy questions, the assistant’s questions are worth more than its first answer. Google’s Workspace guide suggests asking Gemini what questions it has for you. OpenAI’s guide for workplace users recommends having ChatGPT ask you questions and play the task back before it starts.

Three details make an interview go well.

- **A cap.** Three to eight questions, so it knows when to stop and answer.
- **One at a time.** Seven questions at once is a form. One, then the next based on what you said, is a conversation.

— **Permission to guess out loud.** Add “if you have to assume something, say so.”

Knows you, or fresh chat

Some assistants now remember earlier chats. Where that memory is on and well stocked, a shorter version can work:

 PROMPT

From what you know about me, where am I getting in my own way right now? Name one pattern. Show the evidence from our past conversations, separate what you know from what you’re guessing, and give me one small thing to try this week. Don’t flatter me.

In a new account, a temporary chat or an app with memory switched off, that version produces the horoscope again. The interview version works anywhere, so treat it as the default. Either way, read the answer as a mirror, not a diagnosis.

WORKSHEET

Interviews to adapt

For a decision, use recipe 4 in chapter 14.

For a plan:

 PROMPT

I want a realistic plan for **[goal]**. First ask me up to five questions, one at a time, about my time, my budget and what has failed before. Then give me the plan as a weekly table, and tell me what you assumed.

For a thought you can’t put into words yet:

 PROMPT

I can’t explain this clearly yet, so I’ll type everything in my head about **[topic]**. It will be a mess. Ask me up to four questions, one at a time, to find what I’m really asking. Then write my question back to me in two sentences.

The question I’ll try first: _____

M: Measure

How to stop the assistant from simply agreeing with you.

You have mostly decided to get a second dog, and you type this.

WEAK

“Thinking of getting a second dog to keep Biscuit company. Great idea, right?”

You will very likely be told it’s a great idea. You asked for a yes, and assistants are generous with those.

A study by Cheng and colleagues, published in *Science* in 2026, found that 11 leading models affirmed users’ actions nearly 50 percent more often than people did. Participants who got the flattering answers trusted the model more and were less willing to repair a conflict. Agreement feels like help. Often it isn’t.

Measure is the part of the prompt that pushes back. It says how you will judge the answer, and it asks the assistant to check itself.

STRONG

I’m thinking of getting a second dog. I have a four-year-old beagle and a small yard, and three days a week I’m out of the house for nine hours. Don’t reassure me. Give me the strongest case against doing it now, then the strongest case for, then the two facts that would settle it. Say what you’re assuming about my dog and my budget, and say “I don’t know” where you don’t.

Three kinds of friction

The other side. Ask for the case against, or for both cases at equal strength. A useful move is to present your plan as a stranger’s: “Treat this as a plan from someone you have no reason to please.”

The assumptions. Ask what it assumed about you. Often one guess is wrong, and fixing it changes the advice.

The doubt. Anthropic’s blog on prompting says that giving a model permission to express uncertainty reduces made-up answers. OpenAI’s guide for workplace users

recommends asking it to mark uncertainty, list assumptions and give sources. One sentence does it: “If you’re not sure, say so.”

A check line improves the odds. It doesn’t make the answer right, which is why chapter 10 exists.

A test for success

Measure also means saying how you’ll know the answer is good. “It’s a good plan if I can follow it in a week when I’m tired.” “It’s a good email if she can reply in one line.”

WORKSHEET

Check-line library

Twelve lines to add to the end of a prompt. Two per prompt is plenty.

- ☐ Tell me what you assumed about me.
- ☐ Say where you’re guessing.
- ☐ Give me the strongest case against this.
- ☐ What would have to be true for this to be the right call?
- ☐ Which part are you least sure of?
- ☐ If the evidence is thin, say so.
- ☐ Separate what I told you from what you inferred.
- ☐ What can’t you know from here?
- ☐ Don’t flatter me.
- ☐ Link a source for every fact I might act on.
- ☐ Tell me which part to check with a professional, and which kind.
- ☐ What question haven’t I asked that would change the answer?

After the prompt

What to say after the first answer, and when to start over.

You're opening a pottery stall at the weekend market and you need a name. You ask for one. The assistant offers "Clay and Co." You sigh and close the tab.

The woman at the next stall asks the same question and gets the same flat answer. She keeps typing. "Give me 30 more, in five styles." Then: "I like 4, 12 and 23. What do they have in common?" Then: "More like those, but shorter, and leave out the word clay." Ten minutes later she has a name she loves.

Her prompt was no better than yours. She treated the first answer as a draft.

Keep talking

Google's Workspace guide tells users to make it a conversation, and Microsoft's Copilot guide says to expect back-and-forth. Ethan Mollick's advice for beginners is to ask for 30 ideas, not five, and then choose.

Four moves cover most of it.

Widen. Ask for more than you need, in different styles. Choosing is easier than describing. For anything that matters, run the prompt twice and judge the set.

Steer. Point at what you like and ask for more of it.

Push back. Say what's wrong, and quote it. "This line is too stiff: 'I hope this finds you well.' Write it the way I'd say it out loud."

Flip. Ask for the opposite view, a different reader, or a version half the length.

You don't have to know what's wrong. "That's not it, and I'm not sure why. Ask me two questions to find out" is a fine thing to type.

When to start over

Correcting is usually better than restarting, because the chat holds everything you've already said. Start a new chat when the topic has changed, when the conversation keeps circling back to an old idea, or when you want a clean second opinion. Take your context with you, and check the summary before you reuse it.

 PROMPT

Summarize this conversation so I can continue it in a new chat: what I'm trying to do, what we've decided, what's still open, and the details about me that matter. Under 200 words.

WORKSHEET

The follow-up ladder

Each rung is something to type after the first answer.

- ☐ Widen: "Give me 20 more, in four different styles."
- ☐ Steer: "I like numbers _____ and _____. What do they share? Ten more like those."
- ☐ Push back: "This part is wrong for me: _____. Here's why: _____."
- ☐ Flip: "Now argue the opposite."
- ☐ Go deeper: "Why did you put that first? What did you leave out?"
- ☐ Make it mine: "Shrink this to something I can do in 20 minutes today."

Wake

Four checks for reading any answer before you act on it.

It's 11 p.m. A charge you don't recognize has appeared on your card statement, and you've asked the assistant what to do. It answers in three tidy paragraphs. You have 60 days to dispute it, it says, and it adds two links. It sounds like a paralegal. You feel better and go to bed.

Is it 60 days? Maybe. The number could be right, right for another country, right three years ago, or invented, and the reply would read the same. Getting a good answer was half the skill. The other half is reading it.

Four checks cover it. They spell WAKE.

Three ways an answer goes wrong

Made-up facts. A name, a number, a quote or a link that looks right and isn't. Sources alone don't make an answer safe. In a 2025 study coordinated by the European Broadcasting Union and led by the BBC, journalists reviewed more than 3,000 answers about the news from four assistants. They found at least one significant problem in 45 percent of them, and serious sourcing problems in 31 percent.

Stale facts. Prices, rules, deadlines and opening hours change. An answer can be accurate about a world that has moved on.

Flattery. The answer agrees with what you already thought. Chapter 8 covered why that is the default.

W: Where is this from?

Ask the assistant to link a source for each fact, and to say which parts come from general knowledge with no source to point to. Then look.

Treat the reply as a lead, not proof. An assistant can be wrong about its own sources, so open at least one link. Does the page exist, does it say what the answer claims, and what date is on it? OpenAI's guide for ChatGPT users gives the same advice: when an answer depends on current information, ask for a search and sources, then review the result yourself.

A: Assumptions

To answer you at all, the assistant guessed things: your country, your budget, your experience, what you meant. Ask it to list each assumption, correct the wrong ones, and ask which one matters most.

K: Kick the tires

Find the one claim your decision rests on. In the story above, it is the 60 days. Check that claim somewhere outside the chat: the card issuer's own site, the document itself, a person who knows. Asking the same assistant again is not a check. Microsoft's Copilot guide says to verify responses and cross-check them against trusted sources.

E: Escalate or edit

Sort the answer into one of two piles. If it touches health, money, law or safety, take it to a qualified person before you act. The next chapter shows how to arrive prepared. Everything else, edit. Rewrite it in your own words before you send it, sign it or say it. If you can't, you don't understand it well enough to use it yet.

All four at once

 PROMPT

Before I act on this, go back over your last answer and be hard on it. Which facts can you link to a source, and which can't you? What did you assume about me that I never said? Which single claim are you least sure of? Name the one fact I should confirm somewhere else, and tell me where.

WORKSHEET

WAKE checklist

- ☐ W. I asked where the facts came from and opened at least one source.
 - ☐ A. I asked what it assumed about me and corrected what was wrong.
 - ☐ K. The one claim my decision rests on is: _____
 - ☐ K. I checked it outside the chat, here: _____
 - ☐ E. If this is health, money, law or safety, I'm taking it to: _____
 - ☐ E. If not, I've rewritten it in my own words.
-

Health, money, law and the heart

How to use an assistant to prepare for a professional,
never to replace one.

The appointment is tomorrow at 8:40. You waited five weeks for it. You already know that when you sit down, half of what you meant to say will leave your head.

Tonight you could type, “What’s wrong with me?” and get a list of possibilities that keeps you awake. Or you could ask for help turning five weeks of worry into a clear half page to read from.

The first asks the assistant to be your doctor. The second asks it to help you use your doctor well.

If it is urgent

If you or someone else might be in danger, if you are thinking about harming yourself, or if a symptom is severe or getting worse fast, do not open a chat. Call your local emergency number or a crisis line in your country now. An assistant can’t see you, can’t send help, and can miss what a person would catch.

What an assistant is not

An AI assistant is not a doctor, a therapist, a lawyer or a financial adviser. It can’t examine you. It doesn’t know your full history. It holds no license and takes no responsibility for what happens next. OpenAI’s own help page for its health feature says it is designed to support medical care, not replace it, and is not intended for diagnosis or treatment.

What an assistant does well is explain, organize and rehearse. So change the request from “decide for me” to “prepare me.”

The prepare-me pattern

A prepare-me prompt does four things.

- 1 It names the professional you are going to see.
- 2 It tells the assistant not to diagnose, prescribe, pick or decide.
- 3 It asks what the assistant can't know or confirm.
- 4 It ends with questions to bring to that person.

For money and legal questions, add your country or state, and ask which year's rules it is assuming. Both change the answer.

Health. Instead of “What’s wrong with me?”

 PROMPT

I’m seeing my doctor on [day] about [what’s been happening]. Don’t diagnose me or suggest treatment. Interview me first: up to six questions, one at a time, about what I’ve noticed, when it began and what makes it better or worse. Then write half a page I can read from: an opening I can say in two sentences, a timeline with dates, and five questions for the doctor. Tell me which of my answers were too vague, and what only an examination could show. If anything I describe sounds like it shouldn’t wait for the appointment, say so right away.

If the assistant raises no alarm, that is not an all-clear. When in doubt, call the doctor’s office.

A new prescription takes the same shape, with your pharmacist in mind. Ask what this kind of medicine is generally for, and what to ask about the other things you take. Don’t ask a chat whether to start, stop or change a dose. That question belongs to the person who prescribed it.

Mind. Instead of “Be my therapist.”

 PROMPT

I’ve been feeling [what it’s been like] for [how long], and I want to explain it to my doctor or a therapist. Don’t tell me what I have. Ask me up to five questions, one at a time. Then reflect back what you heard in a short paragraph I could read to them, mark anything you’re unsure you understood, and add three questions I could ask. If anything I say suggests I might be in danger, stop and tell me to contact local emergency or crisis services.

This is for finding words before you talk to someone. It is not a way through a bad night. For that, call a crisis line or someone you trust. An assistant may not notice that you need more help than it can give.

Money. Instead of “What should I invest in?”

 PROMPT

I live in [country] and I want to understand my options for [goal]. Treat this as general education, not advice. Explain the main routes, the trade-offs and the typical costs, and say which tax year and rules you’re assuming. Don’t recommend specific products. Finish with what a licensed financial adviser would need to know about me, three questions to ask them, and the figures I should confirm on an official site.

Law. Instead of “Should I sign this?”

 PROMPT

I’m in [country or state] and I’ve been sent this [type of document] to sign: [paste it, with names removed]. Don’t tell me whether to sign. Walk me through it in plain language, and mark anything unusual, one-sided or unclear. Say which points depend on local law you can’t confirm from here, and don’t cite a law or a case unless you can link it. End with the questions to take to a lawyer or a legal aid service before I sign.

The heart. Instead of “Who’s right, me or them?”

 PROMPT

I had a falling-out with [person] and I only have my side of it: [what happened]. Before you say who’s right, write the version they would tell a friend. Then name what I might be missing, and one thing I could say that isn’t about winning. Don’t side with me by default. Say where you’re guessing about them.

One limit. If the other person frightens you, controls you or has hurt you, don’t ask a chat to argue their side. Talk to someone you trust or a local support service.

WORKSHEET

Stop and call a person

Close the chat and contact a person if any of these is true.

- ☐ You or someone else might be in danger, or you are thinking about harming yourself. Call your local emergency number or a crisis line now.
- ☐ A symptom is severe, sudden or getting worse quickly.
- ☐ You are about to start, stop or change a medicine.
- ☐ There is a deadline with legal force: a court date, a notice period, a reply window.

- ☐ You are about to move money you can't afford to lose.
- ☐ You are about to sign something you don't fully understand.
- ☐ The assistant's answer is the only reason you feel sure.

WORKSHEET

Questions for my doctor, adviser or lawyer

I'm seeing: _____ on: _____

Why I'm here, in two sentences: _____

What has happened, with dates: _____

What I've already tried: _____

My five questions: _____

The one thing I must not leave without knowing: _____

Privacy

What never to paste into a chat, and three settings worth finding.

A man wants help writing to his landlord about a broken heater. He pastes in his whole lease: full name, home address, the account his rent comes out of, the landlord's phone number. The question needed one clause.

Share what the question needs and nothing more. Before you paste, take ten seconds to redact. Swap names for roles ("my manager," "my brother"), round the numbers, and delete anything that identifies a person or an account. Whatever your settings, what you type is sent to the company that runs the assistant.

Three settings worth finding

Memory. If your assistant remembers past chats, what you share may be kept. OpenAI's help page on memory says that sensitive information can appear in memory if you share it, and that deleting a chat does not necessarily delete a memory saved from it. Open your memory page and read what's there.

Temporary chats. Use one for sensitive topics. OpenAI's help page says temporary chats do not create memories. Claude's help page says incognito chats aren't saved to its memory or your chat history. Google's blog says Temporary Chats in Gemini aren't used for personalization or training, and are still kept for up to 72 hours. Temporary means not remembered. It does not mean never stored, so the list below applies there too.

Training. OpenAI's help page describes a ChatGPT setting for whether your chats are used to improve its models. Look for the equivalent in the app you use, and choose on purpose.

The never-paste list

- ☐ Passwords, verification codes, security answers
- ☐ Full card or bank account numbers
- ☐ Passport, driver's license, tax or other government ID numbers, or scans of them
- ☐ Other people's medical or private details without their consent
- ☐ Employer or client documents, unless your organization approved the tool
- ☐ Full legal files with the names left in
- ☐ Your home address together with your daily routine
- ☐ Anything that identifies a child
- ☐ Photos or screenshots that show any of the above

Already pasted one of these? Change the password or code, delete the chat, and check the memory page.

Any model

Four switches that change what an assistant can do, and prompts that work anywhere.

Paste the same prompt into three apps and watch. The first replies in one tight paragraph. The second sends back a report with headings, a table and a summary. The third asks you two questions before it will say anything.

Nothing is broken. The companies' own documentation describes different house styles: some models lean short and direct, some lean long and heavily formatted, some ask more questions. The styles shift with each release. So write prompts that don't depend on any of it.

Prompts that travel

- Use plain sentences, with no commands or symbols borrowed from one app.
- Make it self-contained. Assume the assistant knows nothing about you.
- State the shape every time: length, format, reader. That evens out the house styles.
- Write the fallback into the prompt, like this:

 PROMPT

What does it cost to renew a passport in [country], and how long does it take? If you can search the web, link the official page. If you can't, say so and list what I should confirm on the official site.

Four switches

The words are only part of it. The rest is what the app has turned on.

SWITCH	TURN IT ON FOR	IF YOU DON'T HAVE IT
Deeper thinking	Decisions, plans, anything with numbers	Ask for the reasoning and check it yourself
Web search	Prices, rules, news, anything recent	Ask which facts to verify, then verify them
File or photo upload	Questions about a specific document	Paste the part that matters
Memory	Questions that draw on your history	Paste your about-me card

Match the gear to the stakes. Ethan Mollick's 2026 guide to choosing an AI says free models are fine when little is riding on the answer. For high-stakes questions he recommends the most capable model you can get, set to a high thinking level. A stronger setting makes mistakes less likely, not impossible, so the checks in chapter 10 still apply.

The skill is yours

In a study by Jahani and colleagues with 3,750 participants, on a task with a fixed goal, about half of the improvement from a newer model came from people changing how they prompted. On an open-ended creative task, the model did most of the work. Clear asking moves with you from app to app.

WORKSHEET

App settings checklist

Open the assistant you use most and find each of these.

- ☐ The model or thinking picker
- ☐ The web search switch
- ☐ The button for attaching a file or a photo
- ☐ The memory page, and what it has saved about me
- ☐ Temporary or incognito chat
- ☐ The setting for whether my chats are used to improve the model

Field recipes

Twelve annotated prompts for twelve parts of life, and six weak ones to fix.

Tuesday, 7:50 a.m. There are fourteen things on your list and room in the day for maybe five. That is where the first recipe starts.

There are twelve, one per scene in the DreamPrompter library. Under each, Parts maps the sentences by number to DREAM, and Habit names what to carry into your own prompts. Recipes 2 to 7 touch personal ground, so read those answers with the WAKE checklist.

Twelve recipes

01

Focus: the honest to-do list

 PROMPT

Here's everything on my list for today: [paste your list]. I have about [number] hours of real working time. Give me the three things that matter most, the two I could do badly and still be fine, and what to drop or move. Put it in a short table. Tell me what you assumed about my deadlines.

Parts. R: 1, 2. D: 3, 4. M: 5.

Habit. Hand over the whole messy list and one real limit.

02

Work: the raise conversation

 PROMPT

I want to ask for a raise at my review on [date]. I'm a [role] with [number] years in the job, and this year I delivered [three results]. Ask me up to four questions, one at a time, about my manager and how pay gets decided here. Then write a two-minute opening I can say out loud, and the three objections to expect with a calm reply to each. Tell me where my case is weakest.

Parts. D: 1, 4. R: 2. A: 3. M: 5.

Habit. Ask for the weak spot now. It beats finding it in the room.

03

Money: where it went

 PROMPT

Here are my spending totals for last month by category, rounded: [paste totals]. My take-home pay is [amount]. I want a plan for next month: a table with each category, last month's number, a target, and one line on why. Point out the two categories where a small change saves the most. Tell me what you assumed, and which of my numbers look like guesses.

Parts. R: 1, 2. D: 3, 4. M: 5.

Habit. Rounded totals show the pattern without handing over a statement. For debt, tax or investing, use chapter 11.

04

Decide: two doors

 PROMPT

I'm choosing between [option A] and [option B], and I need to decide by [date]. Don't pick for me yet. Ask me up to five questions, one at a time, about what I'd regret either way. Then give me a table: what each option costs, what it protects, and what I'd have to believe to choose it. End with the one question I haven't asked that could change the answer.

Parts. R: 1. A: 2, 3. D: 4. M: 5.

Habit. Ask for a way to decide, not a decision.

05

Know yourself: a week in my life

 PROMPT

I want an honest read on where my energy goes. Ask me up to seven questions, one at a time, about a normal week: what I look forward to, what I put off and what I complain about. Then name one pattern, quote the answers it comes from, and suggest one small experiment for this week. Don't flatter me. Label anything that's a guess.

Parts. D: 1, 3. A: 2. M: 4, 5.

Habit. Ask for evidence in your own words, so you can trace the pattern.

06

People: before the hard conversation

 PROMPT

I need to talk to [person] about [the issue], and I'm dreading it. What I want to come out of it is [outcome]. Give me an opening sentence that names the problem without blame, the two reactions to expect, and a calm reply to each. Then tell me how this might look from their side, what I may be getting wrong, and where you're guessing.

Parts. R: 1, 2. D: 3. M: 4.

Habit. Prepare your first sentence and their point of view. The rest has to be a real conversation.

07

Mind and body: the last 45 minutes

 PROMPT

I want a wind-down routine for the last 45 minutes before bed. I go to bed around [time], I'm on my phone right up until then, and I share a small apartment. Ask me up to three questions first, one at a time, if the answers would change the routine. Then give me five short steps I could start tonight. This is about habits, not medical advice, so tell me what to raise with a doctor if it doesn't help.

Parts. D: 1, 4. R: 2. A: 3. M: 5.

Habit. The last line keeps a habit prompt from posing as medicine.

08

Learn: teach me, then test me

 PROMPT

I want to understand [topic]. What I already know is [my starting point]. Ask me three quick questions, one at a time, to find my level. Then explain it in five short steps, each with an everyday example, and after each step ask me one question and wait for my answer. At the end, tell me the two things I'm still shaky on.

Parts. D: 1, 4. R: 2. A: 3. M: 5.

Habit. Make it check that you followed, not only that it explained.

09

Write and create: a toast that sounds like me

 PROMPT

I'm giving a toast at [the occasion] for [person]. Three true things about them: [three details]. Here's how I sound when I'm relaxed: [paste a message or two you've sent]. Write a 90-second toast in that voice, with one laugh and one sincere line. Not like this: "Webster's dictionary defines love as." Give me two versions, and mark the line you'd cut first if I run long.

Parts. R: 1, 2. E: 3, 5. D: 4, 6.

Habit. Show your voice, and show the tired line you're avoiding.

10

Everyday life: what this letter wants from me

 PROMPT

Here is a letter I received, with names and account numbers removed: [paste the letter]. Using only this text, tell me in plain language what it says, what I'm being asked to do, and by when. Then draft a short reply that asks for [what I want to happen]. If anything is unclear or missing, say so instead of guessing, and tell me who to call to confirm.

Parts. R: 1. D: 2, 3. M: 4.

Habit. Paste the real thing and keep the assistant inside it.

11

AI basics: what should I have asked?

 PROMPT

Before we finish, look back over this whole conversation. What's the most useful question I didn't ask? Which thing I told you would you double-check if you were me? What should I do first when I close this chat? Keep your answer to five lines.

Parts. R: 1, the conversation itself. M: 2, 3. D: 4, 5.

Habit. End any chat that matters with this one.

12

Play: a bowl of questions

 PROMPT

I'm having [number] people over for dinner on [day], and a few of them have never met. Give me ten questions to put in a bowl on the table. Keep them light: nothing about work, and each one answerable in under a minute. Mix funny with thoughtful. Here's one I like: "What's a small thing you're oddly good at?" Mark the three you'd open with.

Parts. R: 1. D: 2, 3, 4, 6. E: 5.

Habit. One example sets the tone better than three adjectives.

WORKSHEET

Rewrite clinic

Six prompts that will get you a horoscope. Rewrite each one. The hint names the missing part.

01 "Help me eat healthier." Hint: Reality. _____

02 "Act as a top career coach and tell me what job I should do." Hint: drop the costume, add Ask.

03 “Write a message to my friend who’s upset with me.” Hint: Reality, plus an Example of how you two talk.

04 “Is this a good deal?” Hint: paste the offer, then add Measure.

05 “Plan my trip to Japan: flights, hotels, food, budget and packing.” Hint: Destination. One finished thing.

06 “I’m right to be annoyed with my sister, aren’t I?” Hint: Measure. Ask for her side.

One way to fix the first:

 PROMPT

I want to eat better on weeknights. I cook for myself and a teenager who lives on pasta, I have 25 minutes, and I won’t give up cheese. Give me five dinners one small step up from what we eat now. Tell me what you assumed about my budget.

When not to prompt

What to keep for yourself, what to share, and what to hand over.

A father is packing his daughter's lunch on the morning of her first big exam. He wants to tuck a note in the box. He opens the assistant, asks for something encouraging, and gets four warm, well-turned sentences. He copies them onto a card.

She will never know. He will. The note took him nine seconds, and the nine seconds are the problem. What she would have kept was the slightly clumsy sentence he came up with himself.

Not everything should be a prompt. A guide to prompting should say so.

Three questions before you type

- **Is the effort the point?** Learning a skill, working out what you think, writing the apology. If the value is in the struggle, a shortcut removes the value.
- **Will I need to do this without help later?** If so, use the assistant as a tutor, not a ghostwriter.
- **Does this belong to a person?** Some things need a friend, a partner or a professional. If you lean on a chat for company or comfort most days, tell someone you trust.

What the evidence says

Ethan Mollick's 2026 essay "Choosing to Stay Human" separates using AI to shortcut thinking from using it to help thinking. He cites an experiment in which high school students who did math homework with plain ChatGPT scored worse on the test than classmates without it, and another in which an AI tutor that tailored practice problems raised exam scores. He also notes that ChatGPT, Claude and Gemini now offer learning modes.

The same tool can make you sharper or duller. The difference is who does the thinking.

When the goal is to get better at something, say so:

PROMPT

I want to get better at [skill], so don't do this for me. I'll make the first attempt. Then tell me what worked, if anything did, and the one thing to fix first, with a hint, not the answer. Ask me to try again before you show me how you'd do it.

Keep, share, hand over

Keep what makes you who you are: your judgment, your taste, your relationships, the words you say to people you love.

Share the work where a second mind helps. You draft and it critiques. It quizzes you. It argues the other side.

Hand over the chores where only the result matters: reformatting a list, summarizing a long document you would never read, turning notes into a table.

Our line for this: think better, not less.

WORKSHEET

Keep, share, hand over

Sort the things you did with an assistant last week, or wanted to.

Keep. I'll do these myself: _____

Share. We'll do these together: _____

Hand over. It can have these: _____

One thing I handed over that I want back: _____

Practice and library

A seven-day practice plan, a prompt library page, a glossary and the sources.

Tomorrow you will type something into an assistant anyway: a question about lunch, a reply to an email, a plan for the week. Before you send it, add one sentence that says what you want to end up with. That is day one.

Reading about better questions does little. Asking them does a lot. Ethan Mollick's advice to beginners is to skip the formulas and put in about ten hours of real use.

WORKSHEET

Seven days of practice

- ☐ Day 1, Destination. Take a prompt you'd normally type. Add what you want to end up with and its shape. Compare the answers.
- ☐ Day 2, Reality. Build your about-me card (chapter 5) and paste it above a question.
- ☐ Day 3, Example. Run the style-capture exercise (chapter 6).
- ☐ Day 4, Ask. Run one real, personal question as an interview (chapter 7).
- ☐ Day 5, Measure. Add two check lines (chapter 8) to a decision you're weighing.
- ☐ Day 6, Wake. Run the WAKE checklist (chapter 10) on one answer from this week.
- ☐ Day 7, Your call. Try your best prompt in a second app. Then fill in the keep, share, hand over list (chapter 15).

What I noticed this week: _____

My prompt library page

Copy this page for each prompt worth keeping. Three good ones beat a folder of 300.

Title: _____

When I use it: _____

The prompt: _____

[illegible]

Parts it uses (circle): ☐ D ☐ R ☐ E ☐ A ☐ M

Needs (circle): ☐ nothing ☐ a file ☐ web search ☐ deeper thinking ☐ memory

Worked on (app and date):

Follow-ups that helped:

Glossary

TERM	PLAIN MEANING
Assistant	An AI chat app such as ChatGPT, Claude, Gemini, Copilot, Perplexity or Grok
Context	Everything the assistant has to go on: your words, your files, its memory of you
Fresh chat	A conversation in which the assistant knows nothing about you
Memory	A feature that carries facts about you from one chat to the next
Temporary chat	A chat that isn't saved to memory. Some apps call it incognito
Thinking mode	A slower setting that reasons for longer before it answers
Made-up fact	A confident statement with nothing behind it. Often called a hallucination
Flattery	The pull toward agreeing with you. Researchers call it sycophancy

Sources

The factual claims in this guide come from these pages and papers, read in October 2026. Vendor pages change, so check the date.

From the companies:

OpenAI, ChatGPT Learn, Prompting.

<https://learn.chatgpt.com/docs/prompting>

OpenAI, ChatGPT Learn, Personalize ChatGPT.

<https://learn.chatgpt.com/docs/personalize>

OpenAI Cookbook, ChatGPT Enterprise: Practical prompt engineering for everyday work.

https://developers.openai.com/cookbook/examples/chatgpt/chatgpt_prompt_guide/chatgpt_prompt_guide

OpenAI, Using GPT-5.5 (developer guide).

<https://developers.openai.com/api/docs/guides/prompt-guidance?model=gpt-5.5>

OpenAI, Using GPT-6 (developer guide).

<https://developers.openai.com/api/docs/guides/prompt-guidance>

OpenAI, Reasoning best practices.

<https://developers.openai.com/api/docs/guides/reasoning-best-practices>

OpenAI Academy, Prompting.

<https://academy.openai.com/public/clubs/work-users-ynjqu/resources/prompting>

OpenAI Help, Memory in ChatGPT.

<https://help.openai.com/en/articles/8590148-memory-in-chatgpt>

OpenAI Help, Health in ChatGPT.

<https://help.openai.com/en/articles/20001036-health-in-chatgpt>

Anthropic, Prompting best practices.

<https://platform.claude.com/docs/en/build-with-claude/prompt-engineering/claude-prompting-best-practices>

Anthropic blog, Best practices for prompt engineering for 2026.

<https://claude.com/blog/best-practices-for-prompt-engineering>

Claude Help, Use Claude's chat search and memory to build on previous context.

<https://support.claude.com/en/articles/11817273-using-claude-s-chat-search-and-memory-to-build-on-previous-context>

Google, Google Workspace with Gemini Prompt Guide.

<https://workspace.google.com/learning/content/gemini-prompt-guide>

Google, Gemini API, Prompt design strategies.

<https://ai.google.dev/gemini-api/docs/prompting-strategies>

Google blog, Gemini adds Temporary Chats and new personalization features.

<https://blog.google/products-and-platforms/products/gemini/temporary-chats-privacy-controls/>

Microsoft, Get started writing prompts in Microsoft Copilot.

<https://support.microsoft.com/en-us/topic/learn-about-copilot-prompts-f6c3b467-f07c-4db1-ae54-ffac96184dd5>

From independent teachers:

Ethan Mollick, Getting started with AI: Good enough prompting.

<https://www.oneusefulthing.org/p/getting-started-with-ai-good-enough>

Ethan Mollick, Choosing to Stay Human.

<https://www.oneusefulthing.org/p/choosing-to-stay-human>

Ethan Mollick, An opinionated guide to which AI to use to do stuff (July 2026).

<https://www.oneusefulthing.org/p/an-opinionated-guide-to-which-ai-b22>

Lenny's Newsletter, AI prompt engineering in 2025, with Sander Schulhoff.

<https://www.lennysnewsletter.com/p/ai-prompt-engineering-in-2025-sander-schulhoff>

From research (the four Prompting Science Reports are from Wharton):

Prompting Science Report 1 (politeness).

<https://arxiv.org/abs/2503.04818>

Prompting Science Report 2 (step-by-step prompting).

<https://arxiv.org/abs/2506.07142>

Prompting Science Report 3 (tips and threats).

<https://arxiv.org/abs/2508.00614>

Prompting Science Report 4 (expert personas).

<https://arxiv.org/abs/2512.05858>

Cheng and colleagues, Sycophantic AI decreases prosocial intentions and promotes dependence, Science, 2026.

<https://doi.org/10.1126/science.aec8352>

Cheng and colleagues, free preprint of the same study.

<https://arxiv.org/abs/2510.01395>

Jahani and colleagues, Prompt Adaptation as a Dynamic Complement in Generative AI Systems.

<https://arxiv.org/abs/2407.14333>

BBC Media Centre, Largest study of its kind shows AI assistants misrepresent news content 45% of the time (October 2025).

<https://www.bbc.co.uk/mediacentre/2025/new-ebu-research-ai-assistants-news-content>

What would you ask if you knew how?

THE BOOK

The Prompt Book. A field guide to asking AI better questions. Autumn 2026 edition. Published by DreamPrompter. Free to read, print and pass on.

POWERED BY

 Digital Kitchen, the studio behind DreamPrompter: dkstudio.ai

LICENSE

Creative Commons Attribution 4.0 International (CC BY 4.0). Copy it, share it and adapt it for any purpose, as long as you credit DreamPrompter and link to the license: creativecommons.org/licenses/by/4.0

ON THE WEB

The library of prompts, the lessons and the newest edition of this book are at dreamprompter.app

A CAUTION

Advice about AI ages fast, and an assistant is not a doctor, a lawyer or a financial adviser. Chapter 11 says how to use one to prepare for the person who is.

TYPE

Set in Newsreader, Geist and Departure Mono, each under the SIL Open Font License.
