A Useful AI Request Without Private Information

One small question. No personal details. Your judgment stays in charge.



ONE SMALL FICTIONAL TASK

An inviting first exercise for anyone curious about conversational AI, with three fictional stories, practical ways to examine a fluent answer, and two optional private paper cards.

You do not need to create a new account to learn the ideas in these pages. You do not need to give an AI your name or tell it what has happened in your life. You may try the example later with an AI service you have independently chosen—or use it as an ordinary pencil-and-paper thinking exercise.

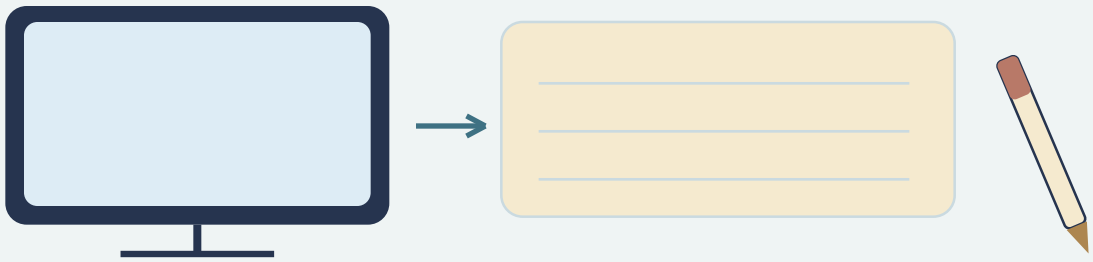
EXPLORE INSIDE

Privacy first
Three examples
Check facts
Private cards

October 2026 · All people, clubs, places and example outputs are fictional

Your judgment stays in charge.

The helpful thing is a first draft



There is a pleasant moment at the beginning of many projects when nothing is decided yet. The note has not been written. The outing has not been arranged. The small thing you might learn is still a possibility. You have room to think without committing yourself.

That is a good place to meet an AI writing assistant.

Imagine sitting down with a notebook to plan a message for a fictional neighborhood reading group. You want the note to sound friendly and brief. A conversational AI might provide six possible openings in a few seconds. That speed can be useful. It also means you can be handed six polished sentences before you have settled on what you actually want to say.

The distinction matters. A sentence that sounds finished is not necessarily a sentence you endorse. An attractive itinerary is not a confirmed arrangement. A neat explanation is not proof that the facts are right.

For our first encounter, give the tool a modest job: *suggest*, *rephrase*, *organize*, or *compare*. Keep the last action for yourself: choose, check, reject, or leave it alone.

This short guide asks you to practice on invented material. If you never open an AI service, the exercise still offers a useful way to distinguish an idea from a decision. You can write a first draft for yourself just as well. The aim is not to become dependent on a new technology; it is to learn one way to use it, with a clear boundary.

A first principle: A helpful draft creates options. It does not create obligations.

Put the real-life details away first



An AI service does not need to know where you live to help draft a general invitation. It does not need to know your birthday to suggest hobbies. It does not need the name of a relative, a family photograph, the contents of a private letter, or a story about somebody's illness to help you practice a courteous sentence.

A safe beginner's exercise begins by **inventing** its people, situation, and place. Not by copying real information and removing only the name. A rare job, a specific date, and a distinctive event can identify somebody even when the obvious name is gone.

Compare these two requests:

Not suitable for a beginner's public AI exercise: “Rewrite the letter I received yesterday from my neighbor at 17 Elm Road about our disagreement. Here is the whole message.” This exposes another person's words and a specific address.

Suitable for our fictional practice: “Write three warm opening sentences for an imaginary neighborhood reading circle inviting newcomers to an informal discussion. No personal names or real locations.”

The second request leaves the work of composing a welcoming sentence intact. Almost everything sensitive has vanished because it was never inserted.

Before you type, pause over four questions. Is this about a real identifiable person? Does it include any information that was entrusted to me? Would I mind if these words were retained or seen beyond this conversation? Can I replace the whole example with a fictional one and still get help?

If the last answer is yes, make the replacement before entering anything. If a task requires confidential information, do not assume the service is appropriate. Check the specific provider's terms, privacy controls, workplace rules and consent requirements—and consider not using it.

A useful boundary is not “I hope it stays private.” It is “I did not put the private detail in.”

Choose a task you can discard



Some tasks seem impressive because they matter a great deal: interpreting medical information, reviewing a contract, making a financial choice. They are precisely the wrong tasks for a beginner's casual experiment. The consequences may be important, the input may be private, and a confident-looking answer can be wrong.

Choose something that can go into the recycling bin without consequence. A fictional welcome note. A list of creative uses for a blank notebook. Five titles for an imaginary local-history display. An alternate first paragraph for a story you wrote purely for this exercise.

A good practice task has three characteristics.

It is small. You know what one acceptable response might look like. Three sentence openings are easier to judge than an elaborate month-long plan.

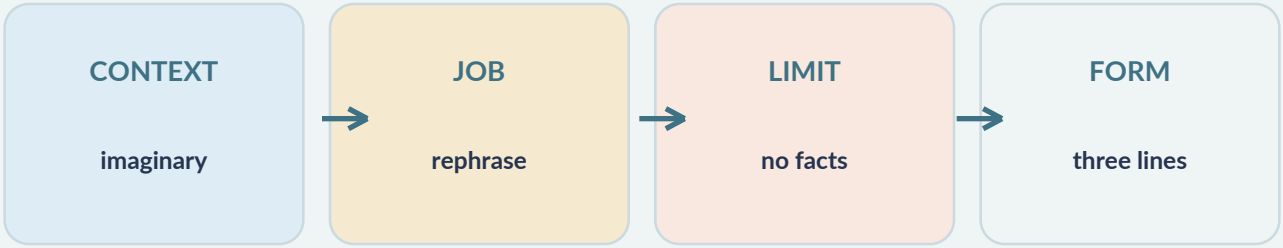
It is invented. No member of your family, workplace, community organization, or client appears under a real or lightly disguised identity.

It is reversible. Nothing is purchased, booked, published, delivered, sent to another person, or decided by the model. You can close the page.

Try this imaginary assignment on paper first: An invented group called *The Lantern Reading Circle* wants one friendly sentence for a casual gathering. You want to avoid sounding either exclusive or overly enthusiastic. Make up a sentence yourself: “If you enjoy talking about stories, you are welcome to join an informal conversation.”

It is already usable as a starting point. The AI's job is not to rescue you from your own words. It is simply to offer some alternatives. With a small comparison like that, you can discover whether the tool is worth your attention at all.

Four parts make a clear request



You do not need special vocabulary, a secret formula, or pages of instructions. One useful request often needs just four parts: imaginary context, the job, a limit, and the form of the answer.

Context: “For an imaginary neighborhood reading circle…”

Job: “Suggest three different welcome sentences…”

Limit: “Do not invent a real location, date, host, fee, or registration process.”

Form: “Use plain, warm English. Keep each sentence under twenty words.”

Put together, the example becomes:

" For an imaginary neighborhood reading circle, suggest three different welcome sentences. Do not invent a real location, date, host, fee, or registration process. Use plain, warm English. Keep each sentence under twenty words.

The request tells the tool what kind of help would count as help. It also deliberately denies permission to fill missing logistical information with plausible details.

If you have a particular tone in mind, say so. “Gentle rather than promotional” is a useful constraint. If you prefer a few short alternatives to an essay, say so. If the response is too ornate, you can revise your instruction: “Keep the meaning but use simpler words.”

Notice what the prompt does **not** include. There is no personal account, no address, no existing private invitation copied in full, and no need for a model to contact anyone. The more deliberately the practice task is bounded, the easier it becomes to judge the output.

Keep the first request short enough that you can remember what you asked.

Three answers can all be different—and none may be yours



An AI asked for three alternatives might reply:

A — welcoming: “You are warmly invited to share a few thoughts about stories with a small, informal reading circle.”

B — direct: “If you like books and conversation, come explore a story with us.”

C — quiet: “There is room for one more voice in our next conversation about books.”

All three are made-up examples for this guide. They are not an actual response from a named service, and they do not announce a real event.

Which is best? The answer depends on the voice you want. A may feel too formal, B a little pushy, and C intimate in a way you did not intend. You are allowed to borrow only “books and conversation” and write something entirely new.

One possible human revision is: “Anyone who enjoys books and easy conversation would be welcome.” That sentence is not mathematically better; it just might fit your invented aim more comfortably. A person can recognize a tone that a machine cannot know from an abstract instruction.

A helpful way to inspect a draft is to circle a useful phrase, underline an unwanted claim, and cross out anything that sounds unlike you. Do this with a pencil on a copied, non-sensitive practice answer. If nothing works, discard the whole response.

The speed of the tool can make the choice feel urgent. It is not. You can take longer over your decision than the software took to produce the options. The decision is the part that belongs to you.

A polished answer can smuggle in facts



Suppose you asked for a welcome sentence. Instead, the answer reads:

" Join us at the Lantern Library, 14 Willow Street, this Sunday at 2 p.m. Admission is free, and no booking is required.

Every detail sounds ordinary. The trouble is that our imaginary request contained **none** of them. The tool supplied a place, address, day, hour, price, and booking policy on its own. They might be wholly invented. The tone makes the information easy to overlook.

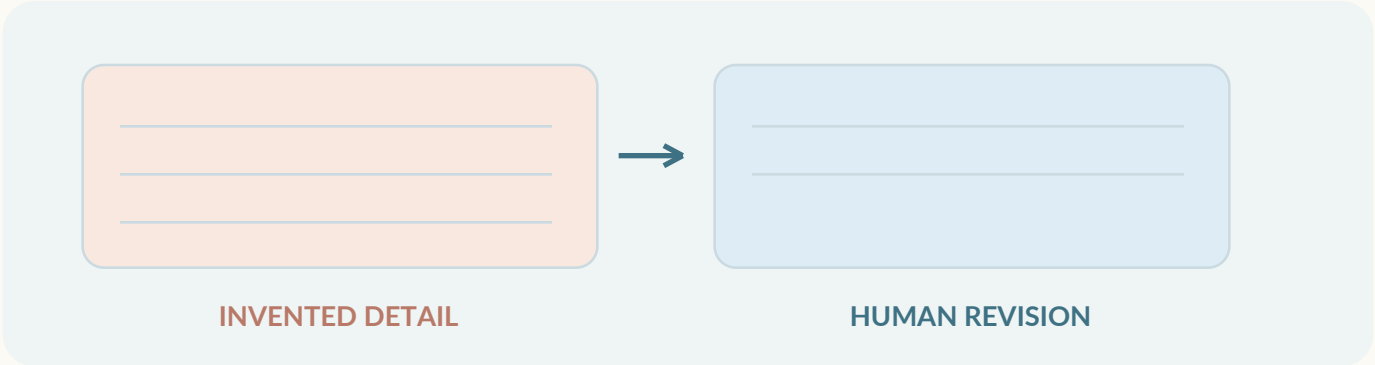
This is one form of the risk people call an AI *hallucination*: the generation of unsupported information that can sound authoritative. A model can also omit a condition, misunderstand a question, or mix facts from different sources. The precise mechanism need not be your first concern. Your useful skill is spotting that something which sounds definite has not been verified.

Take a highlighter to the fabricated announcement. Highlight the venue, street number, Sunday schedule, free admission, and booking statement. These are the parts you would have to verify with an actual responsible organization before using them. You cannot make them true by asking the AI “Are you sure?” A second fluent answer is not independent evidence.

You can reply: “Remove all logistical facts. This is a fictional tone exercise. Keep only the welcome sentence.” But even after that revision, read the result yourself. An instruction to correct an error is not proof that every error is gone.

For real opening hours, prices, route information or rules, consult an up-to-date official source or the actual provider directly. If you cannot verify a necessary detail, leave it out or mark it unknown rather than printing a confident guess.

Elaine's welcome note: a draft, not a sent message



Fictional example one. Elaine used to help with community events. She now wants to write an imaginary welcome note for practice. This is not her actual club's correspondence, and she uses no real group names or personal details.

She wrote: “For a fictional book group, suggest a friendly invitation opening in no more than thirty words. Do not add time, place, fees, or contact information.”

Her imagined AI response was pleasant but went beyond the request: “Come meet us every Thursday at the charming downtown library, where all are welcome.” Elaine immediately noticed *every Thursday* and *downtown library*. The request had not established either fact. She crossed out the invented schedule and venue.

Elaine then wrote her own sentence: “A relaxed conversation about books can be a lovely way to meet new voices.” She kept it in her practice notebook. She did **not** send it, sign up a member, confirm an event, or claim to represent an actual library.

What changed? Not much, deliberately. She explored a tone and saw how easily logistical claims can appear. She also discovered that her own version suited her better than any of the proposed ones.

One reader might call the exercise a success because it produced a usable sentence. Another might call it a success because it exposed a limitation before a real message was involved. Both are reasonable. The result is not a published invitation. It is a tested way of thinking about a first draft.

Elaine's stopping point: one sentence she likes, saved privately on paper; no real-world action.

Martin's three possibilities: ideas are not bookings



Fictional example two. Martin wants to find something interesting to do on a quiet afternoon. He is curious but has not selected a class, venue, seller, or schedule. He keeps the request fictional: “Suggest three low-equipment creative activities a made-up person could try at home. No local business names, product prices, or guarantees about suitability.”

The imagined suggestions are drawing the shapes of household objects, collecting observations about light at different times of day, and writing a short scene from the point of view of a houseplant. All three are plausible starting ideas. None requires a commercial recommendation from the tool.

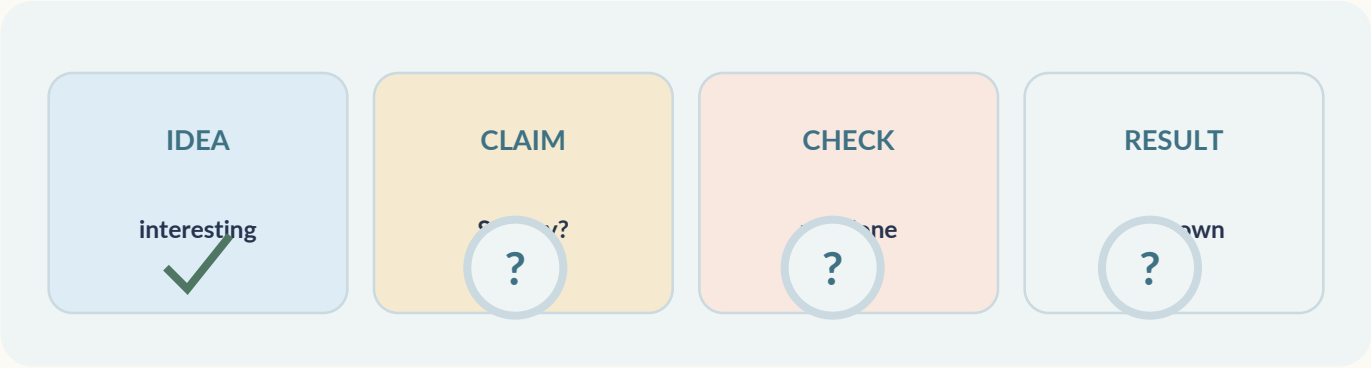
Martin considers the drawing idea. He has a pencil and scrap paper available in this fictional scene, so he tries sketching a mug for ten minutes. His first sketch is not accurate; that is irrelevant. He enjoyed noticing the handle's changing curve. The AI did not draw it for him, and it did not decide whether the activity was worthwhile.

He does not buy a kit or book a course. There is no calculated saving from not buying something he had not planned to buy. His only deliberate experiment is ten minutes with supplies already on hand. If he later considers a paid class, he will check the actual provider's cost, eligibility, location, accessibility, and timetable himself.

The remaining two possibilities stay on paper. The point of the exercise was not to fill a whole week with obligations. It was to obtain a few ideas without sharing personal facts, and to choose one optional activity from them.

Martin's stopping point: one ten-minute home experiment; two untested ideas, zero bookings or purchases.

Noor's useful mistake: the false certainty test



Fictional example three. Noor is fascinated by old maps. To practice questioning an AI, Noor invents a place called *Map House at Copper Square* and asks for a short paragraph describing why someone might enjoy visiting an exhibit about local cartography. The request specifies that the venue is fictional.

The imagined AI answer contains one appealing line about seeing a town change across decades. But it also adds: “Open every Sunday, with free guided tours.” Noor sees the problem: this fictional place has no established operating calendar and cannot be verified as a real visitor attraction.

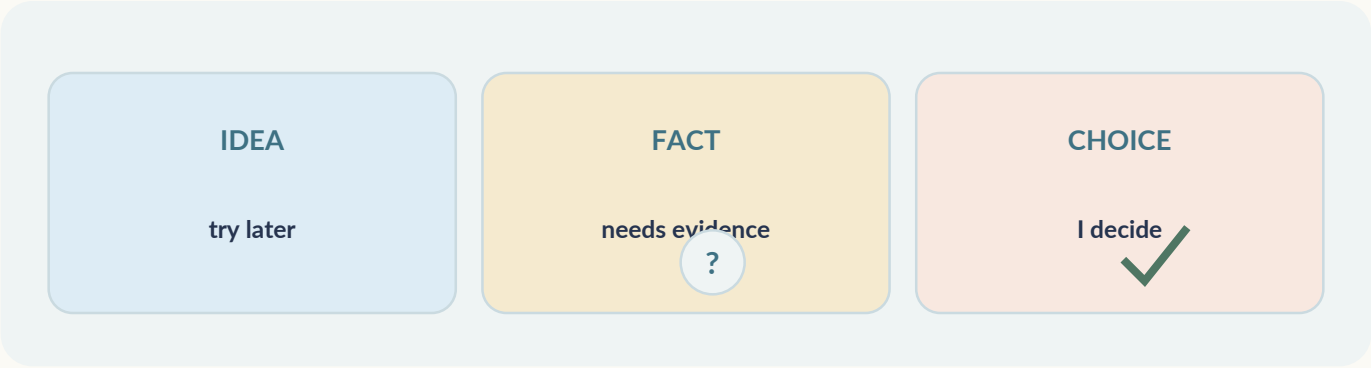
Noor makes a two-column note. On one side: “Interesting idea—a series of maps could show how streets changed.” On the other: “Unsupported claim—Sunday opening and tours.” Noor refuses to carry the second sentence into any real recommendation.

Instead of asking the same AI to confirm its invented claim, Noor practices the next sensible question: “If I wanted to visit a **real** museum, which facts would I need to confirm with its official site or staff?” The useful answer would name items such as current opening days, the exhibit's actual existence, admission terms, transport and access arrangements. Even that checklist would not establish any of those facts.

Noor writes no public review, does not travel to the imaginary venue, and pays nothing. The lesson is not that every AI statement is wrong. It is that an appealing description and an independently verified fact are different kinds of things.

Noor's stopping point: a clearer checklist and one fabricated claim correctly left unverified.

Sort what you have been given



After even a harmless request, an AI response may mix different sorts of statements in one confident paragraph. Sorting them into three groups can make it easier to think.

Idea: “You could practice by writing about the view from an imaginary train window.” This is a creative option. It need not be true about the world to be useful. You might still decide it is dull.

Checkable fact: “The town library holds a writing group at 4 p.m. tomorrow.” That is a statement about a real-world condition. An AI answer alone cannot establish it. You need a reliable current source, such as the library's own announcements or direct staff confirmation. If you have no source, label it unconfirmed.

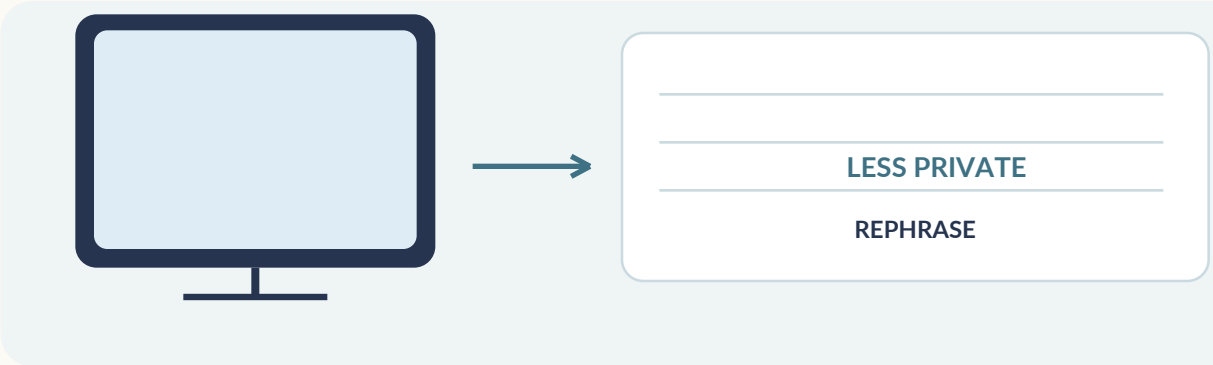
My choice: “I would enjoy writing in the morning rather than joining a group.” This depends on you. Neither a model nor another person's routine has the final say.

When an answer claims to summarize a source, check that the source exists, that the quoted material is accurate, and that the conclusion actually follows. A convincing-looking citation can be wrong or point somewhere irrelevant. This is especially important before passing advice to other people.

A quick three-mark pencil system can help. Draw a light bulb beside an idea, a question mark beside a fact needing verification, and a small heart or check beside a personal choice. You do not have to use those exact symbols. The important thing is to avoid treating all the words as one category.

If the answer is too difficult to classify, reduce the task. Ask the AI to list its suggestions separately from any claims about the real world, then independently inspect what it produces. The separation is convenient, not a guarantee.

A second request can be smaller, not more personal



When the first result disappoints, people sometimes respond by supplying their full situation: family names, calendar details, addresses, correspondence, bills. It feels natural: perhaps the system needs more context to do better. More context can indeed change an answer. It can also disclose more than the task justifies.

For our imaginary invitation, a safer follow-up is about the writing itself: “The tone is too promotional. Please give two calmer versions. Keep the made-up group and omit all event logistics.” This adds a style preference, not a biography.

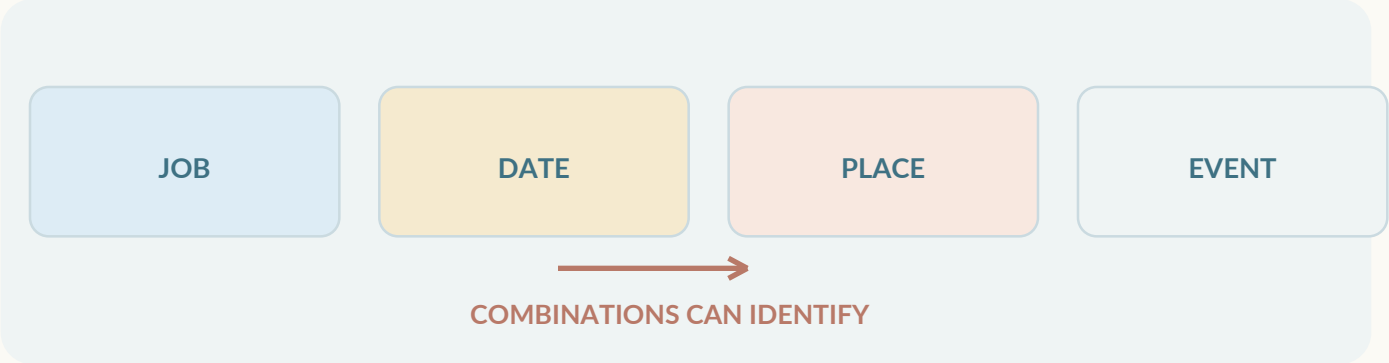
For Martin's fictional hobby ideas: “Compare the activities by what a beginner could do with ordinary paper and a pencil. Do not suggest brands or invent class availability.” Again, the change concerns the task rather than the person's private life.

For Noor's imaginary museum paragraph: “Separate imaginative description from statements that would need independent confirmation. Mark unknown details as unknown.” It is reasonable to ask, though Noor still needs to review the outcome.

Some systems can ask clarifying questions or use connected tools. Be deliberate about permissions. An app connection, access to an address book, microphone, photo collection, inbox or browser is not necessary for this basic fictional exercise. A permission dialog deserves its own consideration; do not grant it automatically because the draft seems convenient.

If the model repeatedly ignores an important boundary, stop using that answer. You do not owe the tool a fourth or fifth chance. The point of practicing is to exercise judgment, not to perfect the system.

Leaving out a name may not be enough



Privacy is not only about obvious passwords or account numbers. A short description can identify somebody when it combines unusual details: a tiny town, a rare volunteer role, a date, a family relationship, a distinctive incident. Taking out the name does not necessarily make the story anonymous.

Suppose someone wants help writing about a conflict in a very small community organization. A prompt that says “the only retired lighthouse keeper in our town, who spoke at the meeting last Tuesday” may point to a real person even if that person's name has been removed. Sharing such facts may be unnecessary, especially when a simple fictional scenario would let you practice the same courteous wording.

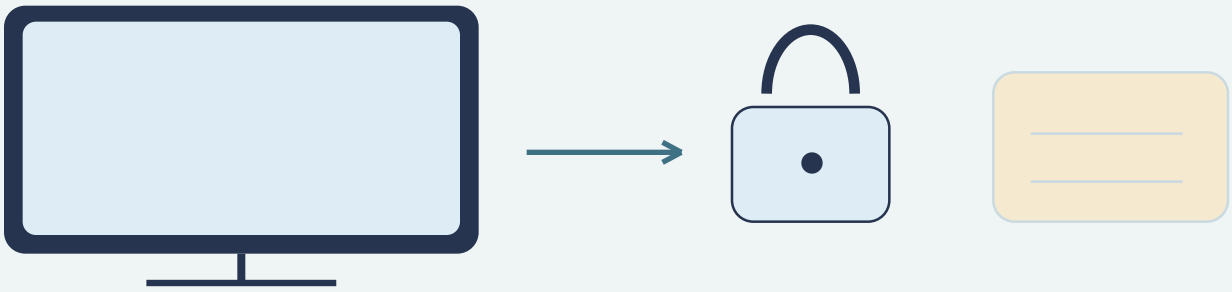
The same caution applies to copied text. A letter may include identifying turns of phrase, signatures, background and private concerns. A screenshot can reveal notifications, account names or tabs at the edge. A photograph can contain other people's faces, documents or location clues. Do not upload such material casually just to save yourself retyping.

There is also a difference between **fictional** and **lightly disguised**. An invented character with no real-world counterpart is appropriate for our small exercise. Changing “Joan” to “Anna” while preserving Joan's exact circumstances may not be.

Before using an AI, ask: “Could I illustrate the writing problem with a scene from a book that does not exist?” If yes, try that. If the work truly involves someone else's confidential or personal information, check their rights and any relevant professional obligations rather than assuming a clever rewrite makes disclosure safe.

A shorter request is sometimes a more respectful request.

Know that privacy settings are not the same as silence



CHECK PROVIDER CONTROLS

Different AI services and account types have different policies for retaining, reviewing, using and sharing what people enter. Some offer controls over whether eligible conversations help improve models. Some offer temporary modes or account-managed rules. Details can change, and a feature in one service is not a promise about another.

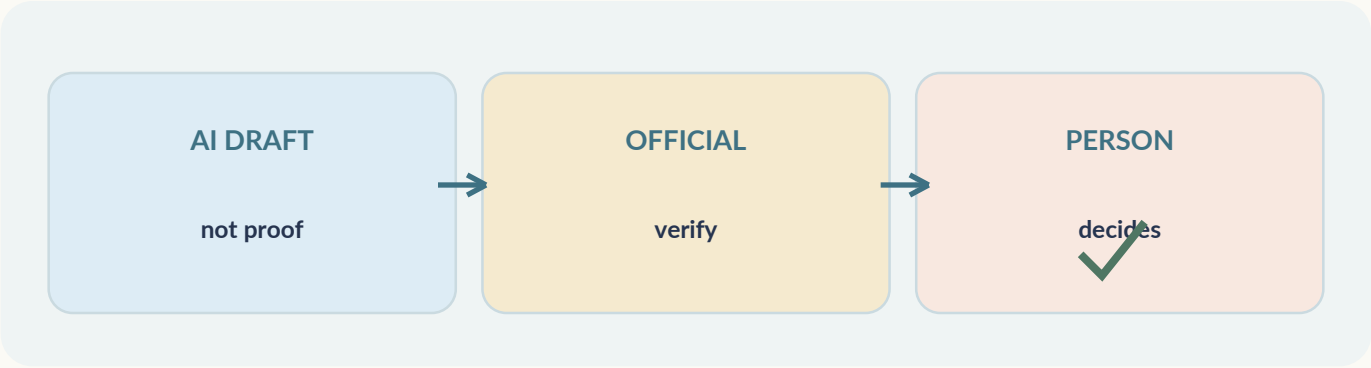
For example, at the time this guide was prepared, OpenAI's official help materials described Data controls for model-improvement preferences and a Temporary Chat option with different history and retention behavior. Those are **product-specific options**, not a universal guarantee that every AI service—or even every type of conversation—deletes information immediately. Even when model training is turned off, that choice is not the same thing as deleting an existing chat or ensuring zero retention.

You do not need to memorize a settings menu to complete this exercise. Use no real sensitive information, and you avoid depending on a privacy setting you may not fully understand.

If you later choose a particular service for real work, read its current official privacy explanation and any privacy notices or workplace rules that apply. Pay attention to whether you are connecting third-party apps, giving a tool permission to act, or sharing an answer beyond the conversation. Do not assume an AI is a private journal, a confidential professional adviser or a secure vault merely because the interface is friendly.

The most durable privacy control for our beginner task is wonderfully ordinary: use invented material in the first place.

When the answer mentions the world, leave the chat to check



An AI can suggest that you visit a gallery, attend a workshop, take a bus or borrow equipment. None of that establishes whether the suggestion is possible where you are, today, at an affordable price or with appropriate access.

For a real public program, go to the organizer's current official information—or call the organizer. For a real transport option, consult the actual operator. Check the date, service area, full return trip, fare conditions and accessibility details that matter to you. For a real purchase, use the seller's current terms. A source from years ago may no longer describe the situation.

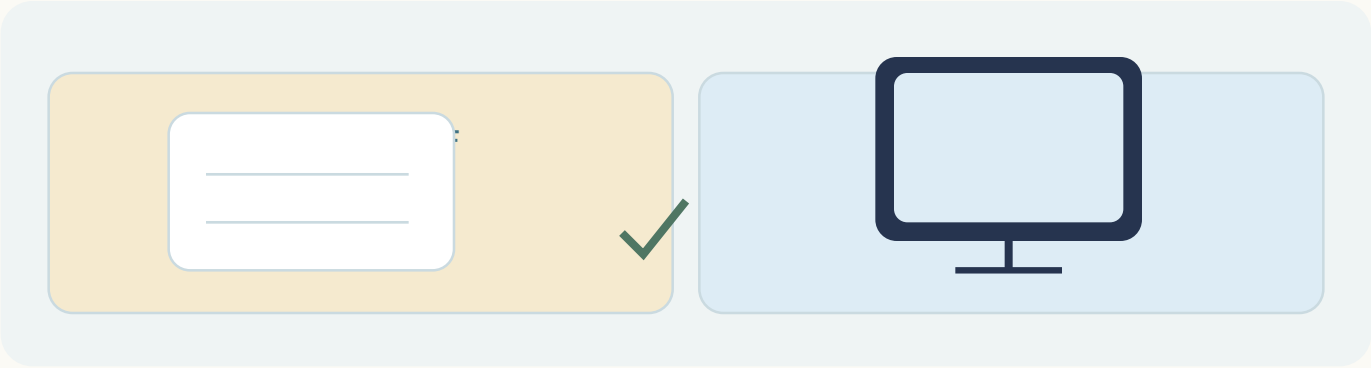
Do not treat a list of links supplied by an AI as already vetted. A URL may be invented, outdated, unrelated or legitimate but insufficient for the particular claim. If you use one, independently confirm that it belongs to the right organization and supports what you are relying on.

For health, legal, financial or emergency matters, a casual language-model response is not a substitute for appropriate qualified advice and verified information. Our first practice avoids those topics entirely.

In Noor's fictional exercise, the venue was made up. No amount of searching for it would turn the fabricated Sunday opening hours into reality. The honest response is simply: “This exercise has no actual venue.” If Noor later researches a real museum, that is a *new* task with different information needs.

A useful four-word reminder is: **Draft here. Verify there.** The final decision remains yours, including the decision that the tool has not earned a place in this activity.

Was it worth asking at all?



Not every use of AI saves time or improves the day. You may spend longer correcting a draft than you would have spent writing your own three sentences. That is not a failure. It is useful information about the tool's value for you.

Imagine two approaches to the same imaginary welcome sentence. In the first, you jot an opening in one quiet minute. In the second, you enter a request, read three suggestions, notice an invented detail, ask for a revision, and decide that your original sentence is clearer. Neither approach deserves a prize. One may simply suit you better.

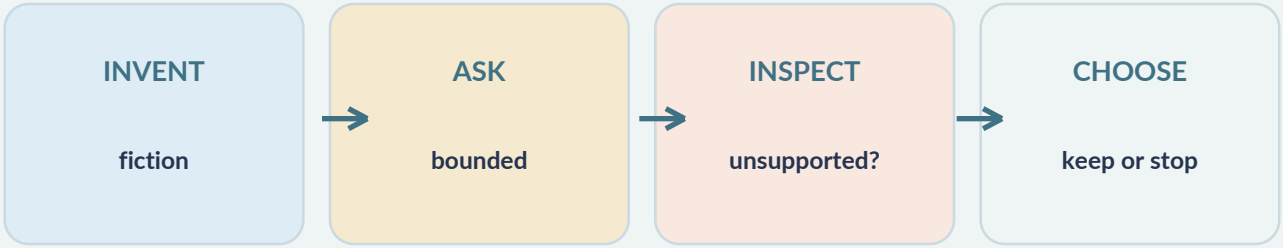
A small task is worth trying when the alternatives help you start, reveal a word choice you had not thought of, or provide a new perspective that you can evaluate. It is less worthwhile when you are repeatedly correcting confident inaccuracies, struggling with a cluttered interface, or feeling compelled to turn every private thought into a prompt.

Sometimes a blank page is a better companion. A friend may offer a more thoughtful response. A real librarian or local organizer may know something the AI cannot verify. Those are not obsolete ways of doing things.

The purpose of the exercise is not to persuade you that AI belongs in every part of life. It is to put a new option beside older, reliable ones. Keeping control includes being free to decline the option.

Your first success can be nothing more dramatic than this: “I now know what sort of small, non-private request I would try—and what I would check before using the response.”

The whole exercise, on one page



You can try the following on paper now and, if you wish, with a conversational AI later. Nothing in it requires a real account, a contact list, or a purchase.

Invent. Make up a gentle task. “An imaginary neighborhood group would like a one-sentence invitation for people who enjoy gardening stories.” No real organizer, address or membership list is needed.

Ask. “Offer three warm, plain-English openings, each under twenty words. Do not add a real venue, date, fee or link.”

Inspect. Mark one phrase you like. Circle any factual claim, hidden assumption, or awkward promise. Notice whether the answer quietly shifts from *fictional draft* to *real event*. If it does, don't use that detail.

Choose. Keep, rewrite, independently check a claim, or throw the whole response away. If nothing is useful, the experiment is complete anyway. You have not lost an obligation or spent any money.

A response to the imaginary task might be: “If stories and growing things interest you, join us for an easy conversation.” Another might be: “Bring a memory of a garden and enjoy exchanging stories.” A third might use a tone you dislike. These sample lines are original examples for this guide, not evidence of an actual group.

Once you have seen the exercise, repeat it only if you want to. One careful encounter teaches more than ten unexamined answers pasted into real messages. The valuable habit is the pause between reading what an AI suggests and doing something in the world.

A useful tool is one you can put down



DRAFT / VERIFY / DECIDE

An excellent first AI request does not need to be ingenious. It needs to be *bounded*: a small invented situation, one clear job, explicit omissions, and an answer you can examine before using.

You may discover a helpful phrase. You may discover that the system offers plausible details you did not give it. You may discover that you prefer your own opening. All of those are legitimate outcomes. None requires you to reveal your life to a service or to treat a fluent answer as an authority.

Keep three distinctions close: **invented practice is not personal disclosure; a suggestion is not a verified fact; a draft is not your decision**.

If you choose to use AI beyond this first exercise, confirm the current privacy options of the exact service you have chosen. Even an official settings description is about a specific product and date—not a promise that every tool handles information alike. For independent real-world facts, go to a reliable source that actually has responsibility for them.

Optional official reading (links verified while preparing this October 2026 edition):

- Office of the Privacy Commissioner of Canada, *Principles for responsible, trustworthy and privacy-protective generative AI technologies*: https://www.priv.gc.ca/en/privacy-topics/technology/artificial-intelligence/gd_principles_ai/
- OpenAI Help Center, *Data controls in ChatGPT*: https://help.openai.com/en/articles/7730893-data-controls-in-chatgpt

These links are offered to illustrate privacy principles and **one** service's published controls. They are not a requirement to use that service, a guarantee of confidentiality, or an endorsement of entering sensitive information.

The choice to write a sentence yourself, after all, is still available. That is part of what makes this the beginning of a new chapter rather than another thing you must master.

Optional card: My harmless first request

INVENT	LIMIT	STOP

Use this privately. You may leave every line blank. Write only fictional or non-sensitive practice material; no real names, contact details, addresses, passwords, account numbers, or personal stories. Do not upload a filled card to an AI service.

1. An imaginary situation I can use:

2. A single small job (suggest / rephrase / organize / compare):

3. Details that must NOT be added:

4. Tone or length I would enjoy:

5. One safe sentence I might ask:

6. A stopping point that requires no real-world action:

You may carry the idea, not the filled-in sheet, into a practice session. Recheck your words before typing. A fictional writing task is enough.

Optional card: What I kept, checked, or left behind

KEEP | CHECK | DISCARD

Keep this card offline and private, or leave it unfilled. It is not a scorecard. Do not record sensitive details or paste a private conversation into it.

1. One phrase or idea that was genuinely useful:

2. One unsupported detail, if any:

3. The kind of independent source needed to verify it:

4. What I changed in my own words:

5. My actual decision (keep / check / discard / do nothing):

6. Would I use this approach for another harmless task? Why or why not?

A thoughtful “no” counts. You are not being graded on using a new tool. You are deciding what belongs in your life.