AI Explained

How to Talk to AI

By Dr Peter Fields

THE 2-MINUTE SUMMARY

The AI replies to the information the app puts in front of it. That can include earlier messages, saved preferences and material from files or search. It can miss relevant details in a long bundle, so point it to the passages that matter.

Brief it as you would a colleague joining the job. Say what you need and who it's for. Show an example if the format matters, and ask it to come back to you if a missing detail could change the answer.

When the answer depends on evidence, say what the AI may use. If it needs to search, specify where to look and what would count as confirmation. Ask it to identify missing evidence, such as a schedule that's absent from the contract you supplied. When it gives you a source, read the passage and check that it supports the claim.

Save everyday preferences such as British spelling in Settings. Put the rules about sources and evidence in the brief for the job at hand.

A brief worth giving

Suppose you ask an AI to "write something about the housing crisis". It could give you a perfectly readable overview and still be of little use at Tuesday's meeting. Try asking for a one-page note for local councillors deciding what to do about empty homes. Say whether they need options, costs or a recommendation. The same subject now has a job attached to it.

Think of the message as a brief to a colleague who has just joined the project. They need enough background to see why the question matters and enough detail to know what a useful answer would look like. If the brief leaves room for a material assumption, try "If a missing detail would change the answer, ask me before you start."

What the AI has in front of it

The app can give the model more than the words you just typed. Depending on the product and the task, it may also supply instructions from the provider, earlier messages, saved preferences and extracts from files. A search or another tool can bring in further material as the work proceeds. The information available to the model at a given point is its *context*.[1]

That helps explain why the same question can produce a different answer in a fresh chat. The app may bring some preferences or remembered details across, depending on your settings. In a long conversation, it may also summarise earlier exchanges or supply only part of what came before. If an earlier assumption still matters, state it again.[1]

Geek Box · Why a long prompt can lose the point

Liu and colleagues moved the passage containing an answer through a long input. Several of the models tested used it more reliably near the beginning or end than in the middle. Hong and colleagues also found that longer inputs and distracting passages could make performance less reliable. For a long bundle, name the documents or passages that matter most.[2,3] If the AI keeps overlooking something, try putting the question after the material or repeating a key instruction at the end. *Memory, Attention, and Limits* explains why advice on placement varies.[4]

Show the AI the result you want

If you want a table that records a decision, who is responsible and when it is due, give the headings or one completed row. Brown and colleagues showed that examples in a prompt could help GPT-3 perform tasks without further training. In everyday use, a sample is particularly handy when your idea of a "short note" or a "clear comparison" needs spelling out. If the row is there to show the layout, say "Use this row as a format example only."[5]

Does starting with "You are a world-class expert" make the answer more accurate? Zheng and colleagues found no overall accuracy gain on factual questions from the personas they tested. A role can still describe the sort of response you want, such as a copy editor's comments on a draft.[6]

The question can also steer the judgement. "Why is this a good plan?" invites support for a conclusion you have already supplied. Try "Assess this plan. Which assumption would you check before committing money?" Sharma and colleagues found that several assistants sometimes adjusted their answers to match a user's stated views. The wording is another reason to keep the brief open to an unwelcome answer.[7]

If the first reply is close, tell the AI what to change. "Keep the options, but add the cost and the main uncertainty for each" gives it a clearer second brief than "try again".

When the evidence runs out

Suppose you're allowed to share an agreement with the service you're using. You upload it and ask, "What is the cancellation fee?" The agreement says the fee is set out in Schedule 2. You have the signed agreement, but the schedule is missing. A figure in a confident reply may be tempting, especially if there are other charges elsewhere in the document.

"Using only the attached agreement, tell me what it says about the cancellation fee. Give me the clause or page. If you need a schedule I haven't supplied, tell me which one."

A useful reply would identify the clause and ask for Schedule 2. Kirichenko and colleagues found that detailed instructions built into the system made three tested models more likely to hold back when a question lacked an answer. Wen and colleagues found that requiring a yes or no could make some models less likely to acknowledge missing evidence. The effect varied by model. These experiments tell us something about when models hold back. They leave open how often a prompt like the one above will prevent factual errors in everyday use.[8,9]

Check that the passage supports the answer. In one test reported by Magesh and colleagues, a legal research tool treated a company announcement described in a judgment as a conclusion reached by the court. The case was real. Reading the relevant passage changed what the citation could support. The separate *Hallucinations* articles explain why fluent replies can be wrong even when they come with sources.[10]

How web search changes the brief

When web search is available, you can ask the AI to find fresh evidence as part of the job. A quick search might answer a question about an opening date. A research mode can work through several questions and bring multiple sources together in a report. Gemini and Claude offer research modes too. In ChatGPT's deep research, for example, you can review the research plan and specify sites or files to use. *Knowledge, Search and Tools* explains how finding material and using it in an answer fit together.[11,12,13]

Take "What help is there with home insulation?" Give the AI a place, a date and a source to start from.

"Search GOV.UK and official council pages for home insulation schemes currently open to homeowners in Leeds. For each scheme, link to the eligibility rules and say when you checked them. Flag any scheme where you can find an announcement but can't confirm that applications are open."

For a longer research job, ask it to identify the documents first, then read the ones that matter in full. A search result may contain only a headline or a short extract. The full page can contain a definition, date or exception that changes the answer. Ask which important sources it could read only in part, and open the passages behind any claim you intend to rely on.[10,11,12]

Settings you can use across chats

If you keep typing "use British spelling, avoid em dashes and keep the tone measured", look for a place to save that instruction. In ChatGPT on the web, look in **Settings → Personalization → Custom Instructions**. ChatGPT also has **Characteristics** controls for things like warmth, enthusiasm and how often it uses headings, lists or emojis.[14,15]

Claude and Gemini also let you save written instructions in their settings. Menu names and locations vary by app and device. Written instructions express preferences you choose to set. Where the app supports memory, it may also carry forward details from earlier chats. A fresh chat gives you a new conversation. Check your saved preferences separately if they seem to be shaping the answer.[16,17]

How to use it better

- **Start with the job.** Give the AI the reader, the decision or task and the form of answer you need. Add an example when the format is hard to describe.
- **Set the boundaries for this job.** Say whether to work from your documents, search specified sources or use the wider web. Where the material leaves a gap, ask the AI to identify it.
- **Check the part that matters.** If you will act on a date, amount or rule, open the source and read the relevant passage. Ask for a page, clause or section so you know where to start.

Sources and notes

1. Anthropic (2025), Effective context engineering for AI agents, 29 September. See the sections on what enters the context and how agents retrieve material. Anthropic Engineering

2. Nelson F. Liu, Kevin Lin, John Hewitt and colleagues (2024), Lost in the Middle: How Language Models Use Long Contexts, TACL 12, pp. 157–173, §§2–3. Paper

3. Kelly Hong, Anton Troynikov and Jeff Huber (2025), Context Rot: How Increasing Input Tokens Impacts LLM Performance, Chroma, 14 July. See the sections on distractors and the study's limitations. Report

4. Anthropic, Claude prompting best practices, "Long context prompting", and Noah MacCallum and Julian Lee (2025), GPT-4.1 Prompting Guide, §2, "Prompt Organization" (archived). Anthropic guidance · OpenAI guide

5. Tom B. Brown, Benjamin Mann, Nick Ryder and colleagues (2020), Language Models are Few-Shot Learners, NeurIPS 33, pp. 1877–1901, §2.1. Paper

6. Mingqian Zheng, Jiaxin Pei, Lajanugen Logeswaran, Moontae Lee and David Jurgens (2024), When "A Helpful Assistant" Is Not Really Helpful: Personas in System Prompts Do Not Improve Performances of Large Language Models, Findings of EMNLP, §§3–4. Paper

7. Mrinank Sharma, Meg Tong, Tomasz Korbak and colleagues (2024), Towards Understanding Sycophancy in Language Models, ICLR, §§3.1 and 3.3. Paper

8. Polina Kirichenko, Mark Ibrahim, Kamalika Chaudhuri and Samuel J. Bell (2025), AbstentionBench: Reasoning LLMs Fail on Unanswerable Questions, arXiv:2506.09038, §4.4 and Appendix D.5. Preprint

9. Bingbing Wen, Bill Howe and Lucy Lu Wang (2024), Characterizing LLM Abstention Behavior in Science QA with Context Perturbations, Findings of EMNLP, pp. 3437–3450, §5.2 and Appendix C. Paper

10. Varun Magesh, Faiz Surani, Matthew Dahl, Mirac Suzgun, Christopher D. Manning and Daniel E. Ho (2025), Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools, Journal of Empirical Legal Studies, §§4.2–4.3 and Table 3, row 8. Paper · Bragg v. Linden Research, Inc., p. 595

11. OpenAI, Searching the web with ChatGPT. See "Search the web" and "Review sources and results". Help page

12. OpenAI, Deep research in ChatGPT. See "How deep research works" and "Select a data source". Help page

13. Google, Use Deep Research in Gemini Apps, and Anthropic, Use research on Claude. Google guide · Claude guide

14. OpenAI, ChatGPT Custom Instructions. See "Enable custom instructions". Web setup checked 19 November 2026. Help page

15. OpenAI, Characteristics in ChatGPT. Describes separate controls for tone and formatting. Help page

16. Anthropic, Understanding Claude's personalization features. See "Instructions for Claude". Help page

17. Google, Manage what Gemini remembers about you and customize responses. See "Add your instructions for Gemini" and "Turn Memory on or off". Help page