

Prompt Engineering Best Practices

Discussion Questions and Exercises

Kevin Crook

KevinCrook.com

Copyright © 2026 by Kevin Crook (KevinCrook.com). Free to use and adapt with attribution, with full terms at the end of this document.

Start Simple with Interactive Chat

1. **(estimated 1 to 2 minutes)** Name a time an AI gave you a useless answer. What was missing from the question rather than from the answer?
2. **(estimated 5 to 7 minutes)** Take a deliberately vague request and write six prompts that build on each other, each one adding exactly one thing. Report which single addition you think would change the result the most, and defend it.
3. **(estimated 10 to 15 minutes)** Pick a real task. First, write the one perfect prompt you would attempt if you only had one shot. Then write the six-step interactive version. Trade both with another group and have them judge which would produce better results. Finish with a paragraph on which was easier to write, and whether easier and better were the same answer.

The Secret Sauce of Prompt Engineering

4. **(estimated 1 to 2 minutes)** Three techniques combined. Before we go further, which do you predict does the most work?
5. **(estimated 5 to 7 minutes)** Write what you expect each of the three to contribute. Then rank them by importance. Report your ranking and defend whichever one you put last, since that is the one you are claiming matters least.
6. **(estimated 10 to 15 minutes)** Find a task where this combination would be overkill or would actively make things worse. Describe it. Then write the simplest prompt that would do the job properly. Finish with a paragraph on how you tell in advance whether a task deserves structure.

Meta-Prompting

7. **(estimated 1 to 2 minutes)** Would you rather be handed a good answer or a good question? Say why in one sentence.

8. **(estimated 5 to 7 minutes)** Write a messy, plain-language description of something you actually want. Then write the single sentence that converts it into a meta-prompt. Report your sentence and compare wordings with another group.
9. **(estimated 10 to 15 minutes)** Write your messy description, then write your own best attempt at the expert prompt it should produce. Trade with another group and have them write what they would expect an AI to produce from your description. Compare all three. Finish with a paragraph on what the other group included that you left out. If devices are available, run it and add the AI's version to the comparison.

Role + COSTAR

10. **(estimated 1 to 2 minutes)** Seven fields. Name the one you think most people forget.
11. **(estimated 5 to 7 minutes)** Take a request and fill in all seven fields for it. Then remove one field and describe specifically how the answer would change. Report the field whose removal did the most damage.
12. **(estimated 10 to 15 minutes)** Write the same request three times, changing only the Audience field: once for a child, once for an expert, once for someone who has to make a decision today. Predict each response in a few sentences. Finish with a paragraph on how much of the output is decided by that one field alone.

In Context Learning

13. **(estimated 1 to 2 minutes)** Name a document you would want an AI to know that it could not possibly have been trained on.
14. **(estimated 5 to 7 minutes)** You have a document the AI has never seen. Write five questions it could only answer with the document in hand, and five it could bluff its way through. Report the pattern that separates the two lists.
15. **(estimated 10 to 15 minutes)** Write a one-page briefing that would let an AI answer questions about something you know well. Then trade, and have another group write three questions your briefing cannot answer. Finish with a paragraph on what you left out, and on whether you left it out because you forgot or because you assumed.

Manage Context Rot and Task Rot

16. **(estimated 1 to 2 minutes)** Name the first sign that would tell you a long conversation has drifted off track.
17. **(estimated 5 to 7 minutes)** Write the summary request you would use. Then list everything a good summary has to contain for the work to restart cleanly. Report your must-have list and see what other groups included that you did not.
18. **(estimated 10 to 15 minutes)** Take a long project you know well and write the summary you would carry into a fresh chat. Trade with another group and have them attempt to

describe the next step using only your summary. Finish with a paragraph on what broke, and on why you thought it was safe to leave out.

Handle Lost in the Middle Issues

19. **(estimated 1 to 2 minutes)** You have a hundred-page report. Where would you cut it into chunks, and what is your rule?
20. **(estimated 5 to 7 minutes)** Design a full chunking plan for a specific document: where the boundaries go, what you ask about each chunk, and how you combine the results. Report your boundary rule and defend it against a group that chose differently.
21. **(estimated 10 to 15 minutes)** Take something real and long. Chunk it, summarize each chunk yourself, then combine your summaries into one. Then summarize the whole thing in a single pass without chunking. Compare the two. Finish with a paragraph on what the chunked version caught that the single pass missed, and on what it cost you.

Ask for Chain of Thought

22. **(estimated 1 to 2 minutes)** Have you ever reached a right answer through wrong reasoning? What happened?
23. **(estimated 5 to 7 minutes)** Take a multi-step problem and write out the steps to solve it. Then go back and insert one plausible error partway through, carrying it forward so everything after it stays internally consistent. Trade and see whether another group can find it. Report where yours hid.
24. **(estimated 10 to 15 minutes)** Take an answer that shows its reasoning, either from an AI or one your instructor provides. Audit it one step at a time and mark every step that is asserted rather than supported. Finish with a paragraph on whether the reasoning shown is necessarily the reasoning used, and on how you could possibly tell.

Try Different LLM Versions and Vendors

25. **(estimated 1 to 2 minutes)** How many different AI chatbots have you actually tried? Count hands at one, two, three, and more.
26. **(estimated 5 to 7 minutes)** List the reasons the same prompt might produce different answers on different systems. Then rank those reasons by how much each should affect your trust in any single answer. Report your top reason.
27. **(estimated 10 to 15 minutes)** Design a fair test. Write one prompt you would use to compare vendors, chosen so that real differences would actually show up, and write down exactly what you would measure. Then trade and critique another group's test. Finish with a paragraph on whether better is even one thing, or several things that disagree.

Build or Use a Prompt Library

28. **(estimated 1 to 2 minutes)** Name a prompt you would want to keep forever.
29. **(estimated 5 to 7 minutes)** Design how your library would be organized: what the categories are, and what each entry stores besides the prompt itself. Report your categories and see whether anyone found one you missed.
30. **(estimated 10 to 15 minutes)** Write five prompts you would genuinely reuse, in full, with blanks where the details change. Trade with another group and have somebody try to use one of yours for their own situation. Finish with a paragraph on what they had to fix, because that is what your prompt was assuming without saying.

Share Chat Links

31. **(estimated 1 to 2 minutes)** Name one person you would send a chat to, and say what they would get out of it.
32. **(estimated 5 to 7 minutes)** Compare sending somebody the answer against sending them the whole conversation. List what each one teaches the person receiving it. Report the biggest gap between the two lists.
33. **(estimated 10 to 15 minutes)** Pick something you know how to get from an AI that somebody else does not. Write out the conversation you would want them to see, prompts included, as though you were sharing the link. Then trade and have them describe what they learned. Finish with a paragraph on what the conversation taught that a finished answer never could.

Copyright and Permissions

Copyright © 2026 by Kevin Crook (KevinCrook.com)

Anyone may learn from these materials.

They are free to everyone, with no permission needed.

Anyone may teach from these materials, as long as the people learning pay nothing for it.

Schools, colleges, universities, libraries, nonprofits, community groups, faith organizations, and employers training their own people are all welcome. In a school or college, standard tuition is fine, but no additional fee for a course or workshop built on this material. Nobody may sell access to it, charge admission, or use it as paid content in any form.

Anyone may adapt these materials under the same terms, with attribution.

Add the notice "Adapted from content by Kevin Crook (KevinCrook.com)" prominently under the title of the document, or under the title of the section it belongs to, and at the bottom of any slide or group of slides adapted from the original. Adapted versions carry these same permissions forward, so what you build from this stays free for others too.

If you would like to use this material in a way these terms do not cover, contact me on LinkedIn. The link is at the top of KevinCrook.com.

AI Use Disclosure

AI was used to help prepare this content. All of it was reviewed and edited by the author.

Disclaimer

I am a lecturer at the University of California, Berkeley. This document reflects my own views and opinions, which are not necessarily those of UC Berkeley.