

## WRITING EFFECTIVE PROMPTS

*The key ideas from the video, with a few clarifications on the trickier points.*

### THE BIG IDEA

If you've ever written a clear assignment sheet, or a clear request to a colleague, you already know most of what makes a good prompt. It says **what to do, who it's for, how it'll be judged, and what to hand back**. The robbers in the bank-vault opening asked a generic question ("How do we get out of this vault?") and got generic advice that would fit almost any problem. As the detective put it, there was no context.

### WHY VAGUE PROMPTS GET VAGUE RESULTS Video 3:18

A **prompt** is whatever you type or paste into the chat box, and that can include a whole document. The video's running example is a department's draft Findings section from its annual assessment report (fictional: 17 of 24 students, or 71%, met the standard against an 80% benchmark; last year's 65% came from different projects and an older rubric). Paste it in, add "Revise this," and two different things go wrong.

#### "Revise this."

Two words. The model fills every blank itself, with the most typical guess.

Who reads this?

What do the numbers support?

What are the rules?

What should come back?

#### What comes back

- Smoother, more formal, and generic
- Calls 71% an "improvement" over a year that used a different rubric. It made that up.

#### A prompt with all four pieces

You fill the blanks, so there's much less left for it to invent.

Context: chairs, committee

Task: clear and accurate

Rules: only the draft's claims


Format: section + questions

#### What comes back

- Says plainly the benchmark was missed
- Says the two years can't be compared directly
- Asks bracketed questions instead of guessing

- **AI leans generic by default.** These tools predict the most likely response to what you give them. With two words, the model has to fill in everything else, like who reads this and what the numbers can support, and it reaches for the most typical version of each. Most of the time it won't stop to ask; it hands you a confident draft either way.
- **In the absence of context & oversight, AI is more prone to make things up.** The revision calls 71% an improvement over last year, even though the draft notes that last year's 65% came from different projects and an older rubric. A good prompt shrinks the first problem. Only your own reading catches the second.

### THE FOUR PIECES OF A GOOD PROMPT Video 4:58



In Demo 1, we added the pieces one at a time to the same request, starting a new chat each time so each step's effect shows on its own. The task and constraints step is the one that stopped the invented claim:

*"Revise this section so it's clear and accurate for chairs and the assessment committee. Use only the numbers and claims in the draft. Where something's unclear or unsupported, add a bracketed question for the department instead of fixing it. Under one hundred fifty words, plain language."*

#### 1. "Revise this."

Generic, and invents an "improvement" the data can't support

#### 2. + Context

Knows who's reading and hedges, but still calls it an increase and promises changes the department never mentioned

#### 3. + Task and constraints

"Use only the numbers and claims in the draft ... add a bracketed question instead."

**Nothing invented. Says the benchmark was missed and the years can't be compared.**

#### 4. + Format

"The revised section, then the questions as bullets, then a list of every change."

**Questions ready to send back, and a change list that shows exactly what to check**

"Faculty felt" became "Faculty observed." Small word, big difference.

### CLARIFICATION: WHY A WORD LIMIT IS A TARGET, NOT A PROMISE (VIDEO 6:59)

The model doesn't count words the way you do. It works in tokens, chunks that can be a word, part of a word, or punctuation, so "under 150 words" steers the length without guaranteeing it. Check the length yourself.

### SHOW AN EXAMPLE OF WHAT GOOD LOOKS LIKE Video 7:39

It works from whatever's on the page, so a real sample is the most specific instruction you can give. "Clear and concise" is vague; a Findings section your committee held up as a model isn't. Two cautions from the video: it will copy the example closely, including parts you didn't intend, so remove the other department's name and numbers first; and say what the example is for, such as "match the structure of this, not the content."

**Clarification:** using examples this way is often called *few-shot* or *multishot* prompting. The ability of large models to pick up a task from a few examples in the prompt was a central finding of Brown et al. (2020), and Anthropic's current guidance calls examples one of the most reliable ways to steer format, tone, and structure, while warning that they should vary enough that the model doesn't copy unintended patterns (Anthropic, n.d.). Even one good example helps; the guidance suggests several varied ones when you have them.


### TWO MORE EXAMPLES Video 8:28

| Task                                                     | Weak prompt             | Stronger prompt, built the same four ways                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Action plan</b> (the department missed its benchmark) | "Write an action plan." | Gives the findings, and asks for two or three actions the department could take in one semester with no new budget. For each: who would do it, when, and what evidence would show it worked. Leave names and dates blank for the department, and flag any action the findings don't support.                                                                                                           |
| <b>Data summary</b> (yours to finish in the video)       | "Analyze this data."    | Context and task: a department meeting; capstone rubric scores for 24 seniors on four rubric rows, no names; summarize for faculty with one chart. Rob's constraints and format: compare each row to the 80% benchmark; a bar chart with the benchmark drawn as a line, in colorblind-safe colors; under 150 words; flag any row where one or two students would change whether the benchmark was met. |

- **Read first, decide yourself.** Read the findings before opening the tool. What comes back is a first draft for the department to react to; the actions are the department's call. The evidence piece is what closes the loop, and it's an easy one to leave out.
- **Why that flag matters.** With 24 students, each one is worth about four percentage points. 19 of 24 is 79%, just under an 80% benchmark; 20 of 24 is 83%. You can ask the tool to help you check, not just to produce.
- **Before any student data goes in** (10:26), take out names and ID numbers, and use the account your campus has approved. Another module in the series covers this in depth.

### FOLLOW-UPS: WHERE THE REAL WORK HAPPENS Video 10:44

Almost nothing good comes out of a single prompt, and every follow-up is more context for the next prediction. "Make it shorter" is a revision request; "Cut the sentence about the meeting and open with the benchmark" is a much better one. Say what's working, too: "Keep the bracketed questions exactly as they are." **Why that matters (11:09):** each time you send a follow-up, the model rereads the whole conversation and writes the reply fresh. Nothing from the last draft is held in place unless you say so, so name what stays.



Demo 2 (11:23): setting up a study helper for an intro statistics course.

## WHEN NOT TO PROMPT, AND KEEPING WHAT WORKS Video 12:42

Some writing tasks probably shouldn't be assigned to AI at all: a recommendation letter, a note to a colleague who's had a hard week, the syllabus paragraph about what he values in a classroom. **A test: if the writing is how you figure out what you think, you should write it.** Some judgments stay yours, too, like whether a department's results are good enough.

For everything else, **keep the prompts that work.** Save them in a plain document with brackets where the details change, such as “[audience]” or “[paste findings here].” Most tools also have a spot for standing instructions (your campus version may have it off), and a saved template is worth rereading each term: rewrite it or throw it out if it no longer fits.

### CHECK YOURSELF

1. Name the four pieces of a good prompt (the video's own check, 13:50).
2. Why does “Revise this” come back generic?
3. Which problem can a good prompt shrink, and which one only your reading catches?

*Answers: (1) context, task, constraints, format. (2) It fills every blank with the most typical guess. (3) Generic output; invented claims.*

### TRY IT THIS WEEK

Pick one real task and write a prompt with **all four pieces**. Try it in the tool your campus provides.

If what comes back is useful, save the prompt in a document called **My Prompts**.

*Same draft, very different result: the only thing that changed was what you told it.*

## SOURCES

Anthropic. (n.d.). *Prompting best practices*. Claude Platform Docs. Retrieved October 1, 2026, from <https://platform.claude.com/docs/en/build-with-claude/prompt-engineering/claude-prompting-best-practices>

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., . . . Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

*This handout was created in collaboration with Claude AI (Anthropic).*