

The AI Bootcamp

5

Write an instruction it can actually follow

In this chapter you'll learn

The 3 things a good ask contains, one before-and-after, and the tricks you can stop typing

The AI Bootcamp

THE PROBLEM

Most AI education is a pile of disconnected tips. A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes. None of it is built to be read in order, so nothing compounds. You finish twelve videos and can recite twelve facts, not perform one new skill.

THE FIX

The AI Bootcamp is built the other way. 90 skills, packed 2 to 3 per chapter, 36 chapters, 3 tracks. Every chapter states what it needs from you and what it hands to the chapter after it, so skill 47 only works because skill 12 already landed. Read it in order and by the end you are not collecting facts about AI, you are running work through it.

LETTER FROM THE EDITOR

We built this because we kept meeting smart people who felt behind on AI for no reason other than nobody had organized the material for them.

This book does not assume you are technical, and it does not water anything down either. Every chapter earns its place: a real skill, a real exercise, a one-page companion you can flip back to in month three when you have forgotten the details but still need the answer.

If a chapter loses you, or a chapter changes how you work, I want to hear about it either way.

Kris

Editor, AI Central

Write to me at editor@thecentral.ai

THE 3 TRACKS

AI 101

Get running

Pick a tool, set it up, write a clear instruction, use it safely at work.

AI 202

Prompt properly

The framework behind every good instruction, and catching AI when it is wrong.

AI 303

Build without watching it

Connect AI to your tools, and automations that run while you are not looking.

IN THIS LESSON YOU'LL LEARN

By the end of this chapter you can look at an answer that came back flat and name the exact piece of information that was missing, rewrite the ask in under a minute so it lands, and stop spending effort on prompt tricks that no longer do anything.

Chapter 2 promised you this chapter, because by now you have had a few conversations and the results have been uneven. Some answers were genuinely useful and some were beige, and from where you are sitting that looks like luck. It is not luck, and the difference is almost never the tool.

CONTENTS	5.1	A vague ask gets an average answer
	5.2	A good ask carries 3 things
	5.3	The same task, badly and then well
	5.4	You can stop typing the magic words

TOPICS REQUIRED

Chapters 2, 3 and 4. You need your account open and a few real conversations behind you. Most useful of all: one answer that disappointed you, still sitting in your history. Go and find it before you read on.

AFTER THIS CHAPTER

Chapter 6 assumes your asks are already clear, so that when an answer is still thin you know the question itself is hard rather than badly worded.

5.1 A vague ask gets an average answer

Nothing broke. You asked a question with several reasonable readings, and it answered the most common one, confidently, as though you had meant that all along.




Fig 5.1 Three blanks that every under-specified ask leaves open. It fills all 3 for you, silently, with the most ordinary version, which is why the reply reads as though it were written for nobody. The bar heights here are visual rhythm, not a ranking.

5.1

Not a refusal. It will always answer. That is the problem: a guess and a good answer arrive looking identical.

Average, not wrong. Nothing in the reply is false. It is just written for the average person who might have typed those words.

The tell. If the answer would suit any company in your industry, it was written for none of them.

Here is the problem this chapter exists to fix. You have had a few conversations by now, some answers were genuinely useful, some came back flat, and you have no way to tell which one you are about to get.

Take the ask almost everybody types at some point: write me something about marketing. There is nothing unreasonable about it. It is also a question with many reasonable readings, and the assistant has no way of knowing which one you meant, so it produces the version that would satisfy the largest number of people who might have typed those words. That is the beige paragraph you got back.

Every under-specified ask leaves the same 3 blanks: who the answer is for, what you are actually trying to get done, and what shape you want it in. It fills all 3 rather than asking you, and it fills them with the most common answer available.

Notice what did not happen. It did not tell you the question was too vague. It did not ask who this was for. It did not hedge or slow down. It answered a guess with exactly the confidence it uses when it is right, which is the same unnerving property Chapter 8 is entirely about.

The tell is easy once you know to look for it. Read the reply and ask whether it would suit any company in your industry. If it would, it was written for none of them, and specifically not for yours.

So before you decide the tool is overrated, go back and read what you actually typed. Nearly every disappointing answer in your history is a question that had more than one right answer.

5.2 A good ask carries 3 things

This is not a formula to memorise. It is 3 blanks to fill, written in whatever order they come out of your head.




Fig 5.2 The 3 parts, in the order that is usually easiest to write. The order makes no difference to the answer. Leaving one out does.

5.2

Context. Everything it cannot see. Your company, your quarter, the meeting this is going into, the person reading it.

Task. One verb, one job. Two jobs in one ask gets you a worse version of both.

Format. The part people skip and the part that saves the most editing. Say the shape or it picks one.

A good ask carries 3 things, and none of them is a magic word. Context, the task itself, and the shape you want the answer to arrive in.

Context is everything the assistant cannot see. It does not know your company, your customers, last quarter, the meeting this is going into, or the person who will read it. Anthropic's own guidance puts the point better than most: treat it as a brilliant new hire on their first morning, who has no idea how your place works. You would not send that person off with write me something about marketing either.

The task is the single thing you want done, with a verb in front of it. Summarise. Draft. Compare. Rewrite. List. If you catch yourself typing and also for the second time, you have two tasks in one ask, and you will get a diluted version of both. Splitting a large job into steps is Chapter 16's subject; for now, one ask, one job.

Format is the part most people skip and the part that saves the most editing. Say how long, in what form, and where it is going to be read: 6 bullets for a meeting, a 3-line reply to a customer, a table with one row per option. Leave it out and it will choose for you, and it tends to choose medium-length prose with headings, which is rarely what you needed.

All four of the major assistants publish their own version of this advice and they agree on the substance, which is unusual and worth taking seriously. Some of those guides give the idea a name, an acronym and 2 more parts. That fuller version is real and worth learning, and it is Chapter 13's job, at the start of the next track. Three parts fixes most of what is going wrong for you right now.

The test to use is not whether your ask sounds professional. It is whether a competent colleague, handed only those words and none of your context, could go and do the job. If they would come back with a question first, put the answer to that question into the ask.

5.3 The same task, badly and then well

One real piece of work, written both ways. Everything that changes between them is information you already had in your head and had not typed.

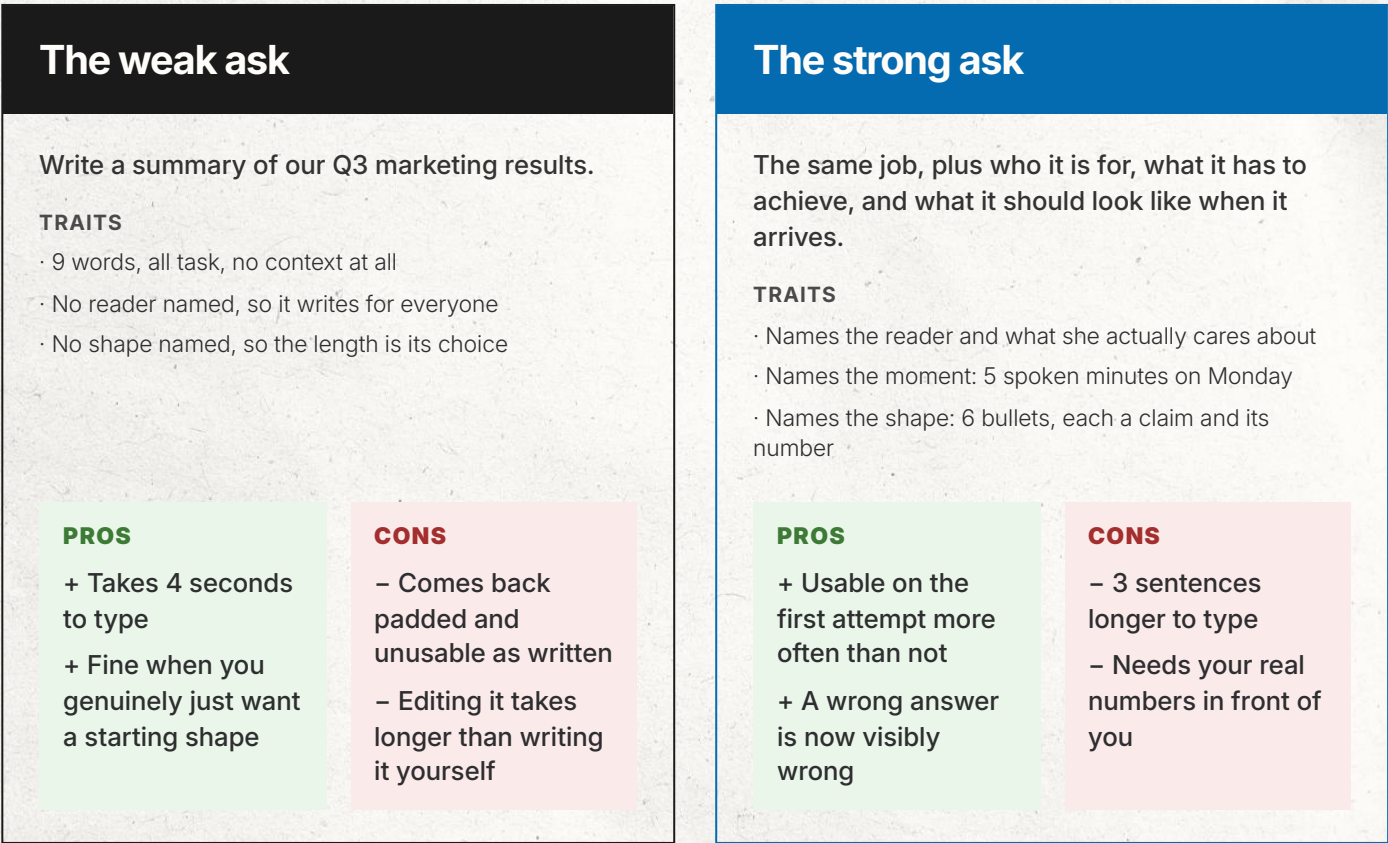


Fig 5.3 Same task, same assistant, same day. Every addition on the right is a fact you already knew. None of it is clever wording.

5.3

The weak one. "Write a summary of our Q3 marketing results." Reasonable, and it will get you something beige.

What changed. Nothing about the task. Only the context around it and the shape asked for.

The cost. About 3 sentences of typing, once. That is the whole investment.

Here is one real piece of work, asked both ways. Read the weak version first, and notice that there is nothing wrong with it.

Write a summary of our Q3 marketing results. Nine words, and every one of them defensible. What comes back will be competent and unusable: a structure it invented, an audience it invented, and far more words than the moment you need it for will allow.

Now the same job, asked properly. "I run marketing at a 200-person B2B software company. On Monday I get 5 minutes to update our CEO verbally at the leadership meeting. She cares about pipeline and cost per qualified lead, not campaign activity, and she will interrupt. Here are the Q3 numbers, with Q2 alongside for comparison. Write me 6 bullets I can read off the page, each one a single claim with the number that supports it, in plain sentences rather than marketing language."

Read the two again and count what actually changed. Not the task: it is the same summary in both. What changed is that the reader now exists, the moment now exists, the thing she cares about is stated, and the answer has a shape to arrive in. All of that was already in your head. You simply had not said any of it out loud.

Two honest caveats before you copy this. The strong version assumes you can get your actual numbers in front of it, and pasting or uploading your own material is Chapter 7's subject rather than this one. And what you are allowed to put in that box at work is a real question with a real answer, which is Chapter 9. Neither changes the shape of the ask, so keep reading.

Go and do this with the disappointing answer you found earlier. Do not start a fresh chat and do not rewrite the whole thing. Add the reader, add what you are trying to get done, add the shape, and send it again.

5.4 You can stop typing the magic words

A good deal of the prompt advice in circulation was true of weaker models and has quietly stopped paying. Knowing which parts saves you effort you can spend on the parts that still work.




Fig 5.4 Four widely repeated tricks, each tested at scale across several models through 2025. The bars are near-identical on purpose: measured against a plain, clear request, none of them reliably changes the answer you get.

5.4

Tested properly. The persona result comes from 6 leading models run thousands of times each on graduate-level questions. No reliable gain.

Where roles still earn it. A role can shift tone and point of view. It just does not make the facts more accurate.

Be polite anyway. It costs nothing and it keeps you writing like a person. Simply do not expect a better answer in exchange.

Now the part most prompt guides will not tell you: a good deal of the advice you have already been handed has stopped working, if it ever worked.

Start with the most repeated one, the expert role. You have been told to open with you are a world-class marketing strategist. In December 2025 a research group at Wharton tested exactly that across 6 leading models, thousands of runs each, on graduate-level questions. Personas matched to the subject produced no reliable improvement. Personas from the wrong field sometimes made things worse, and one model started refusing to answer. Personas suggesting limited knowledge, telling it that it was a beginner, consistently made the answers worse.

The same group tested think step by step. On older models that do not reason on their own it still helps. On the assistants you are most likely using now, which already work through a problem before replying, it bought about 3% more accuracy while adding somewhere between 20% and 80% to how long you waited, and on one of the models tested it came out slightly negative. When the question is genuinely hard, there is a better move than a magic phrase, and it is Chapter 6.

Offering a tip, promising a bonus, applying pressure: tested, and no significant effect on the answer. Politeness is the interesting one, because it turns out to help on some questions and hurt on others and net out at roughly nothing. Say please if you want to. Just do not file it under technique.

What still pays is unglamorous and it is the whole of this chapter: the context it cannot see, one clear task, and the shape you want back. There is an opposite failure worth naming too. OpenAI's own guidance warns that too much extra information can make an answer less useful, and that is the trap people fall into once they learn that detail helps. Everything true and relevant goes in. Nothing decorative does.

Roles have not become worthless, incidentally. They still change tone and point of view, which is a real thing to want, and Chapter 13 puts them where they belong. What they do not do is make it more right. Spend the 30 seconds on the reader and the format instead.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 6.

Do these before Chapter 6

3 exercises. About 15 minutes

Open the disappointing answer you found before reading this chapter. Do not rewrite anything yet. Just fill in the 3 things you never told it.

WHO THE ANSWER IS FOR

WHAT YOU ARE TRYING TO
GET DONE

WHAT SHAPE YOU WANT IT
IN

Now write the ask out in full, using all 3 boxes above, in plain sentences. Send it in the same conversation, not a new one.

Write one sentence on what changed in the answer. Then write down which of the 3 boxes made the biggest difference: that is the one you are habitually leaving out.

Chapter 5 on one page

THE 3 PARTS

Context, everything it cannot see · Task, one verb and one job
· Format, how long, what form, where it lands

THE COLLEAGUE TEST

Could someone competent, handed only these words and none of your context, do the job? If they would ask a question first, answer it in the ask

THE GENERIC TELL

If the answer would suit any company in your industry, it was written for none of them, and a blank was filled in for you

SAFE TO STOP TYPING

"You are an expert" · please and thank you as technique · offers of a tip · "think step by step" on an assistant that already reasons

NEEDED BEFORE CHAPTER 6

One ask you rewrote, and the answer it gave you

BRING TO CHAPTER 6

The ask you rewrote and what came back. Chapter 6 starts from a question that is already clear, so that when an answer is still thin you know the question is the hard part.

BOTTOM LINE

You can say why an answer came back generic, name the piece of information that was missing, and rewrite the ask in under a minute.

NEXT CHAPTER

Let it think harder when the question deserves it

Chapter 6 handles the other reason an answer disappoints: the question was genuinely hard. Modes, reasoning effort and model pickers, what they actually change, and the many times the default was already enough.

BRING WITH YOU


- ☐ The ask you rewrote in this chapter, and the answer it produced
- ☐ One question from your work you suspect is genuinely hard, not just badly worded

GO FURTHER

The Ultimate AI Library

Every guide, prompt pack and playbook AI Central has published, in one place. The reference shelf this course keeps pointing you to.

explore the library



WHO MAKES THIS



AI Central Media is a London-based editorial team. Our flagship publication covers practical AI for senior professionals, and 100+ AI companies, SaaS platforms and growth teams have advertised with us since 2023.

100+

AI companies, SaaS platforms, education brands and growth teams have advertised with us since 2023


613K

Accounts reached every month, across the 7 channels we operate

151

Countries our readers are in, plus all 50 US states

PROUD PARTNER OF



We turn attention into pipeline.

visit thecentral.ai

➤