The AI Bootcamp

13

The 5 Parts of a Prompt That Works

In this chapter you'll learn

The 5 parts of a reliable instruction, which of them do the actual work, and the famous advice that has been tested and does not

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 A LinkedIn carousel here, a YouTube tutorial there, a $2,000 cohort course that goes out of date before the refund window closes.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 name the 5 parts an instruction is built from, assemble one part by part on a real piece of your own work, say which parts change the answer and which only change the voice, and look at a disappointing reply and name the part you left out.

You can already get useful work out of an assistant. What you cannot yet do is say why it worked on Tuesday and not on Thursday, which means you cannot repeat it on purpose, cannot fix it quickly when it fails, and cannot hand it to anyone else. Naming the parts is what turns a knack into a method.

CONTENTS

13.1 5 parts, and what each one buys you

13.2 Watch it improve, one part at a time

13.3 Role is the part everyone overrates

13.4 When it comes back wrong, name the missing part

TOPICS REQUIRED

Chapter 12, which closed AI 101. Bring the recurring uses you named there, and the plain-English asks you have been writing since Chapter 5. You need real instructions of your own to take apart, not hypothetical ones.

AFTER THIS CHAPTER

Chapter 14 assumes you can build a 5-part instruction without looking it up, and adds the one thing that changes an output more than any amount of wording: showing it a finished example of what good looks like.

13.1 5 parts, and what each one buys you

Chapter 5 gave you three of these in plain English and deliberately refused to name them. Here they get names, two more join them, and each one gets held to account for what it actually delivers.

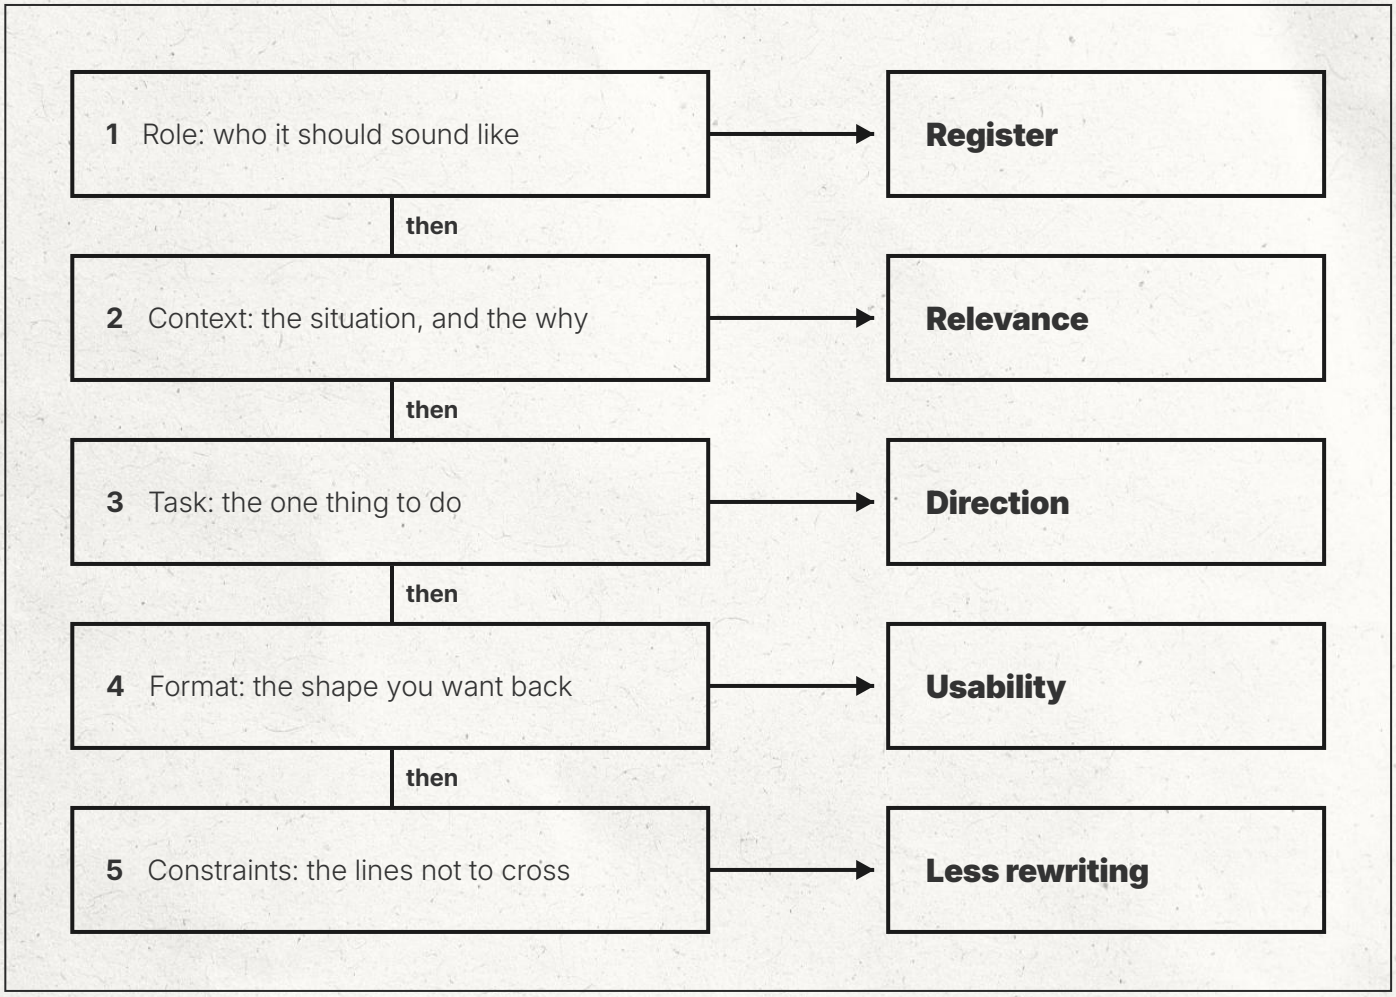


Fig 13.1 The 5 parts, in the order they are conventionally written down. That is a filing order, not a ranking. Section 13.2 builds a real instruction and the order of impact comes out close to reversed.

13.1

Google Publishes Persona, Task, Context, Format for Gemini, and adds that you do not need all four every time.

Microsoft Publishes Goal, Context, Expectations, Source for Copilot. Only the goal is required.

Anthropic Says be explicit, explain the motivation behind an instruction, and be specific about format and constraints.

OpenAI Tells ChatGPT users to identify the task, provide context, and set the tone and format they want back.

Chapter 5 taught you that a good ask contains the situation, the thing you want done, and the shape you want it back in. It gave you a before and an after, it used no jargon, and it was right. Everything here is built on that. What changes now is that the parts get names, two more get added, and each one is judged on what it independently buys you.

The 5 parts are role, context, task, format and constraints. Role is who the answer should sound like and who it is for. Context is the situation and, just as importantly, why you are asking. Task is the single thing you want done, stated as an action. Format is the shape you want back: length, structure, medium. Constraints are the boundaries: what to leave out, what not to assume, where the answer must stop.

The vocabulary is not ours, and that matters more than it sounds. Google publishes Persona, Task, Context and Format as its framework for Gemini. Microsoft publishes Goal, Context, Expectations and Source for Copilot. Anthropic tells you to be explicit, to explain the motivation behind an instruction, and to be specific about format and constraints. OpenAI tells ChatGPT users to identify the task, give context, and set the tone and format. Four companies with every commercial reason to differentiate, describing the same object in four dialects.

The reason to name the parts is not tidiness. An unnamed thing cannot be checked. Right now, when a reply comes back thin or off, your only move is to rewrite the whole request and hope. Once the parts have names, "that came back wrong" becomes "I never told it who this was for", which is a one-line fix rather than a fresh start.

None of this means every message needs all 5. Microsoft says it plainly in its own documentation: all that is required is a clear goal. Most of what you type will be two or three parts, typed in one breath, and that is correct. The framework is a 10-second mental checklist for the asks that matter, not a form to complete.

Before you read on, go and find one instruction you actually typed this week and count how many of the 5 it contained. Almost nobody scores above three, and the two that go missing are usually the same two.

13.2 Watch it improve, one part at a time

A real task, built in passes. Each pass adds exactly one part, and you can see what that part fixed and what it left alone.

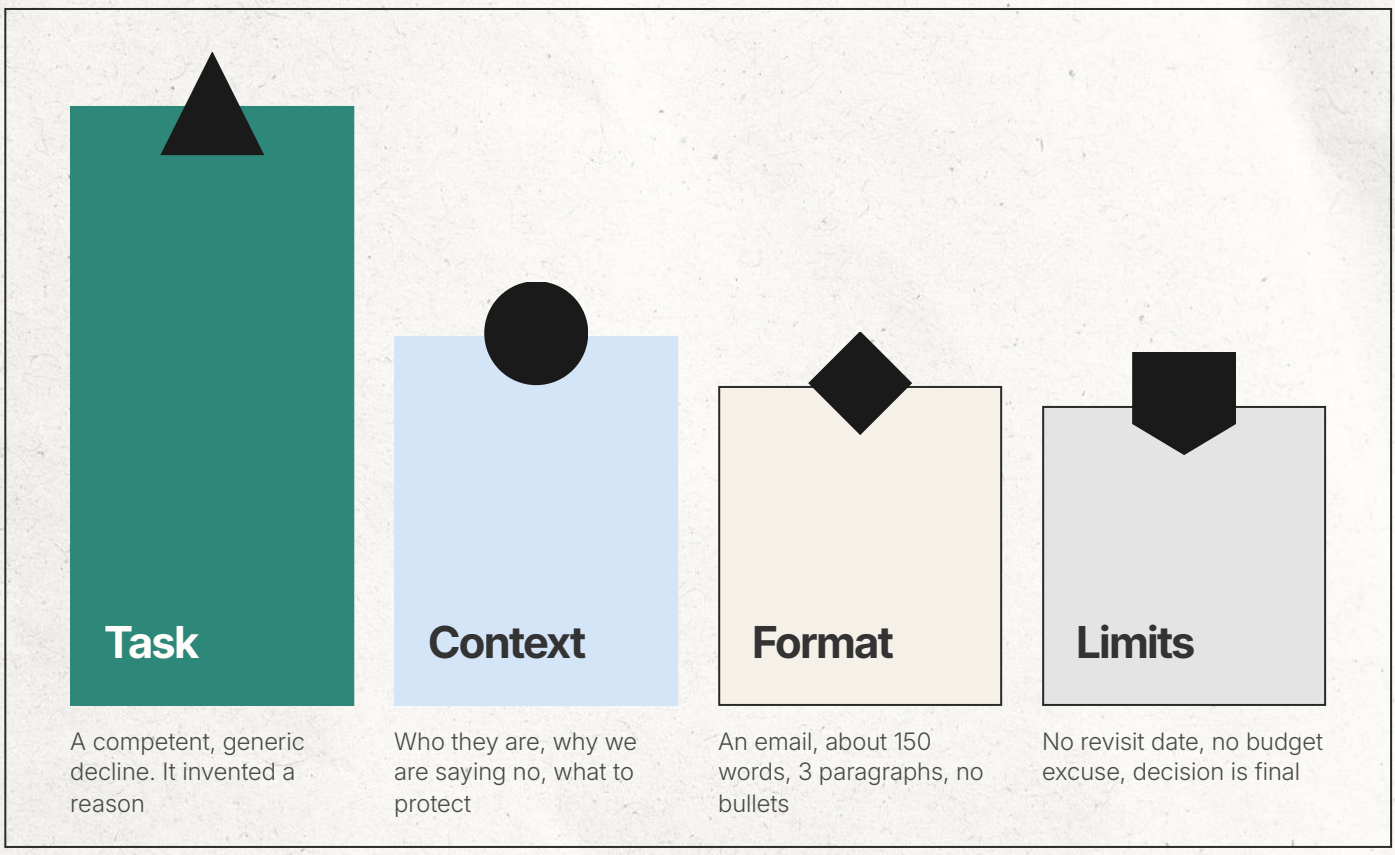


Fig 13.2 One real instruction, built in 4 passes. The heights are how much of the reply still needed rewriting by hand, judged by eye on this single example, not a measured average across many. Role was added fifth and did not move the bar at all, which is section 13.3.

13.2

Context did most of it. One paragraph of situation stopped it guessing. Nothing else moved the answer nearly as far.

Constraints save the edit. Every line you would have deleted by hand is a constraint you could have written once.

Length is approximate. Ask for about 150 words and you land near it. No assistant can count its own words reliably, so treat any number as a target.

Here is the task. A company you already buy from has proposed a partnership. You want to say no to the partnership and keep the supplier relationship intact, and the person who sent it is someone you have met twice and will meet again. This is the kind of message that takes 20 minutes because you keep rewriting the second paragraph.

Pass one, task only: "Write a reply declining this partnership proposal." What comes back is competent and completely unusable. It is twice as long as you would ever send, it thanks them three times, it offers to revisit in the new year, and it has quietly invented a reason for the decline that is not your reason. Nothing here is a mistake by the assistant. It filled every gap you left, because you left them.

Pass two, add context: you are a 40-person software company, they are your hosting supplier and you want to stay their customer, the real reason is that the integration needs engineering time already committed elsewhere this year, and it has nothing to do with price. The invented reason disappears. The reply now argues your actual position instead of a plausible one. This single addition moves the output further than the other three combined, and it is the part most people skip because it feels like typing something the assistant should be able to work out.

Pass three, add format: an email, about 150 words, three short paragraphs, no bullet points, no subject line. The substance does not change. The thing you get back is now the shape of a thing you would send, rather than a document you have to compress first. Format almost never changes whether an answer is right. It changes whether the right answer is usable, which on a busy day is the same problem.

Pass four, add constraints: do not propose a date to revisit, do not use budget as the reason, state that the decision is settled for this year. These are the three sentences you were going to delete by hand every single time. Writing them once as boundaries is the difference between editing an output and approving one.

Pass five, add role: write as a founder who is direct, warm, and does not pad. The email now sounds like you. Read the two versions side by side and you will find that the facts, the structure, the length and the boundaries are identical. Only the voice moved. Keep that observation, because the next section is about what it means.

13.3 Role is the part everyone overrates

The most repeated advice in prompting is to open by naming an expert. It has now been tested carefully, several times, and it does not do the thing people think it does.

What is claimed

That naming an expert makes the assistant know more, and that pushing harder makes it try harder.

TRAITS

· "You are a world-class strategist" before every task
· A tip or a threat to raise the effort
· Politeness treated as a performance lever

PROS

+ Costs nothing to type
+ Feels like you are steering something

CONS

− Tested across models and benchmarks, and does not raise accuracy
− Detailed expert personas measurably lowered it

What holds up

That the role line sets register: who this should sound like, and who is going to read it.

TRAITS

· One clause, aimed at tone and audience
· Naming the reader beats naming the expert
· Dropped entirely when the voice does not matter

PROS

+ Reliably changes voice, pitch and level of detail
+ Cheap, and it never costs you accuracy

CONS

− Will not make it know anything it did not already know
− Easy to mistake for the part doing the work

Fig 13.3 Two claims about the same opening sentence. One of them survives testing, and it is the more modest one. Note that the surviving version is also the cheaper one to write.

13.3

Personas, tested. 6 models, graduate-level questions, expert personas matched and deliberately mismatched to the subject. No reliable gain over no persona at all.

Threats and tips, tested. 9 prompt variants including a $1,000 tip and a $1 trillion one. No meaningful change in overall accuracy.

Why the stories survive. On any single question the swings are enormous in both directions. Averaged over hundreds, they cancel out.

In December 2025 a team at Wharton ran the experiment everyone had been asserting the answer to. They took 6 different models and two sets of graduate-level questions, and compared no persona at all against an expert persona matched to the subject, an expert persona deliberately mismatched to it, and a persona chosen to know nothing. Matched expert personas did not reliably beat having no persona. Mismatched ones sometimes made things worse. The low-knowledge personas usually did.

Separate work in 2026 found the same shape with a sharper edge: as the persona got more elaborate, accuracy on a standard knowledge benchmark fell, from roughly 72% with no persona to 68% with a light one to 66% with a detailed one, while tone, formatting and style all improved. The researchers read this as the model optimising for sounding like the expert rather than for being right. That reading is theirs and you do not have to accept it. The measurements stand on their own.

This is not a contrarian position, which is the part worth noticing. Anthropic's own published guidance now tells you not to over-constrain the role, and says that "you are a helpful assistant" is often better than "you are a world-renowned expert". Google still leads its framework with Persona and Anthropic still recommends setting a role, and both are entirely consistent with the research, because what they claim it changes is behaviour and tone. Neither of them ever claimed it makes the thing smarter. That claim came from the internet.

The same team tested the other famous shortcuts. They tried 9 variations on the same questions across 5 models, 25 runs each, including offering a $1,000 tip, offering $1 trillion, and several threats of the kind that circulate as insider technique. There was no meaningful overall improvement from any of them. An earlier report in the same series found that swapping "please" for a flat command swung results by as much as 60 percentage points on individual questions, in both directions, and cancelled to nothing across the full set.

That last finding explains why the folklore is immortal. Everyone who tried a trick once and got a better answer was telling the truth about what they saw. The swings are real. They are just not the trick, and averaged over a few hundred questions they vanish. Be polite to it because you would rather be a person who is polite, not because it pays a dividend.

So: keep the role line, hold it to one clause, and point it at register rather than expertise. "For a non-technical board" earns its place. "You are a world-renowned McKinsey strategist" does not. Then take the sentence you were going to spend on the job title and spend it on context, which is the part the same body of evidence keeps pointing at.

13.4 When it comes back wrong, name the missing part

The framework pays for itself at the moment something fails, not at the moment you write it. A bad reply is not a mystery. It is a specific part you did not supply.

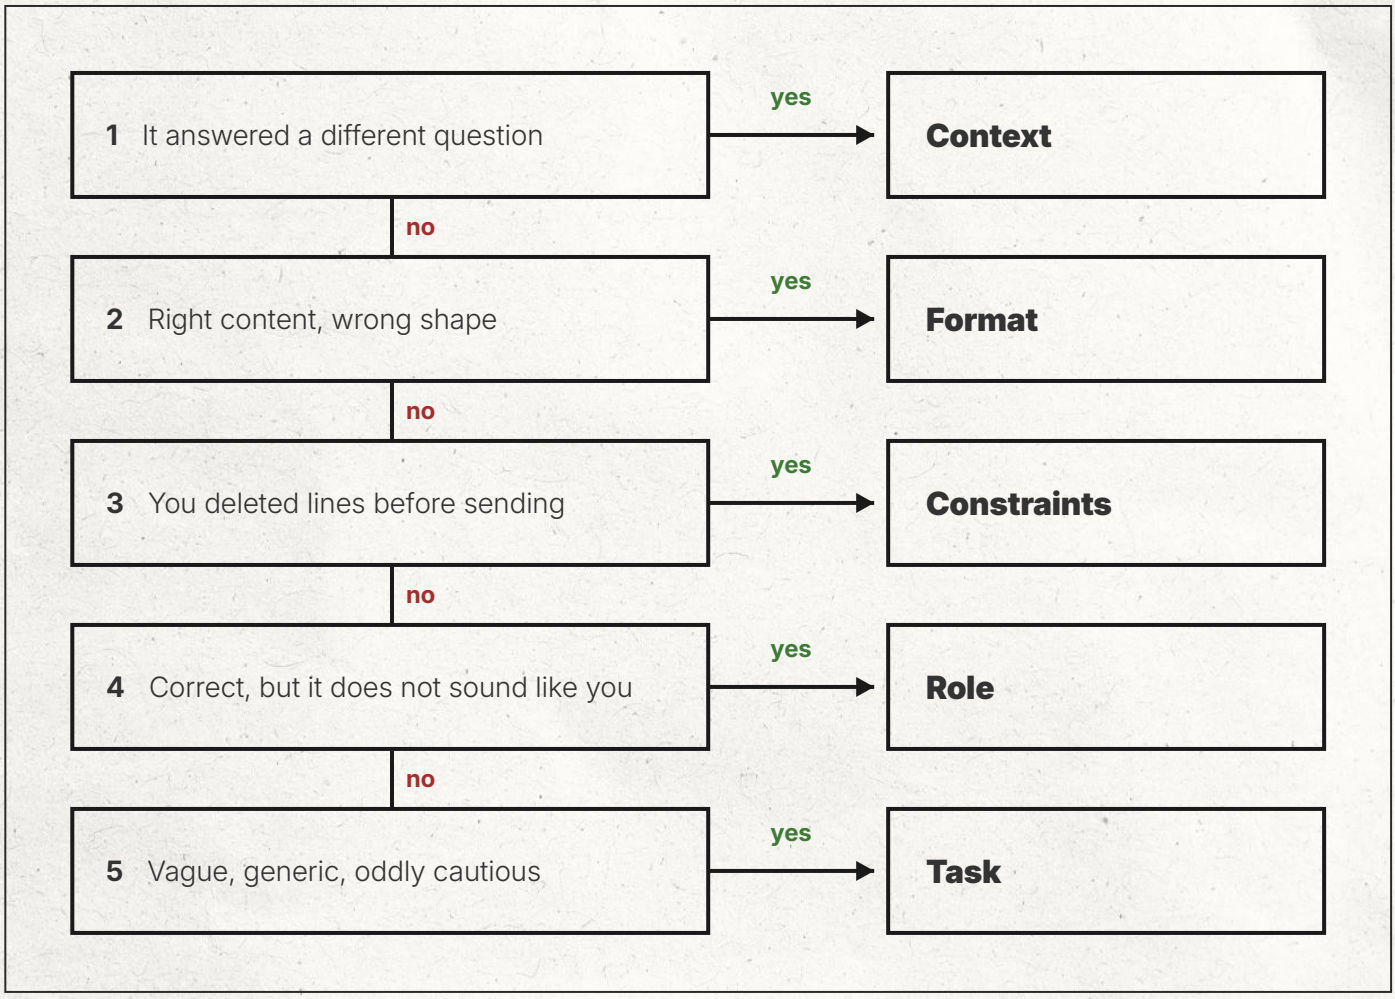


Fig 13.4 Read down until a line describes what you actually got back. The word on the right is the part you left out. Most people respond to a poor answer by rewriting everything: you only ever need to add one thing.

13.4

One part at a time. Add the missing part and ask again. Rewriting the whole instruction throws away the parts that were working.

Say what to do. "Write in flowing paragraphs" lands more reliably than "do not use bullet points".

Attach the reason. A constraint with a why on it generalises to the cases you did not think to name.

Writing a 5-part instruction is the easy half. The half that compounds is what you do 15 seconds later, when the reply is not what you wanted. Run the figure above: read down until one line describes what you got, then add that one part and ask again. Most people meet a bad answer by rewriting the whole request, which discards the parts that were working and teaches you nothing about which ones were.

Two things about constraints are worth getting right, because constraints are where most instructions are weakest. The first is that boundaries land better stated positively. Anthropic's own documentation makes the point with an example: instead of "never use ellipses", say that the reply will be read aloud by a text-to-speech engine, which cannot pronounce them. The second is why that works. A constraint with its reason attached extends itself to the cases you forgot to list, while a bare prohibition covers only the exact thing you banned.

Restraint matters too. Not every ask needs 5 parts, and over-specifying has a real cost: pin down the output shape too tightly on a question that needs judgement and you can flatten the judgement along with the mess. If you want the thinking and the tidy version, ask for the thinking first and the shape second. And treat every number you give as a target rather than a contract. No assistant can reliably count its own words while writing them, so "about 150 words" is an instruction and "exactly 150 words" is a wish.

The 5 parts are the floor that the rest of AI 202 is built on. Chapter 14 adds an example of a finished output for it to imitate, which changes results more than any rewording of the instruction itself. Chapter 15 covers making it reason before it answers. Chapter 16 covers what to do when a job is too big for one instruction, however well built. Chapter 17 turns the instructions you keep retyping into reusable ones. All four assume you can write the 5 parts without thinking about it.

Getting there takes about a week of being slightly annoying to yourself. Before you send anything that matters, run the 5 parts in your head and notice which one you were about to skip. It will usually be context, occasionally constraints, and almost never task. Once you have caught yourself four or five times, the checking stops being a step and becomes how you type.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 14.

Do these before Chapter 14

3 exercises. About 15 minutes

Take one real instruction you sent an assistant this week and pull it apart. Fill in what you left out, using the exact words you would have needed.

CONTEXT YOU LEFT OUT	FORMAT YOU ASSUMED IT KNEW	LIMITS YOU FIXED BY HAND

Now rebuild that instruction with all 5 parts and send it. Write one sentence on what changed, and one on what did not.

Diagnosis drill. Think of the last time an assistant handed you something you could not use, run it down the figure in 13.4, and name the single part that was missing.

Chapter 13 on one page

THE 5 PARTS	Role (who it sounds like) · Context (the situation, and why) · Task (the one action) · Format (the shape back) · Constraints (the lines not to cross)
WHICH TO ADD FIRST	Context, almost always. It changes the answer. Role changes only the voice, so it goes last and stays to one clause
DIAGNOSING A BAD REPLY	Answered a different question, add context · right content wrong shape, add format · you deleted lines, add constraints · sounds wrong, add role · vague, sharpen the task
TESTED AND NOT WORTH IT	Grandiose expert personas, tips, threats, and politeness used as a performance lever. Real swings on single questions, nothing at all on average
NEEDED BEFORE CHAPTER 14	One instruction of your own rebuilt in 5 parts, and the diagnosis of one that failed

BRING TO CHAPTER 14

The rebuilt instruction from exercise 2, and the one you diagnosed in exercise 3. Chapter 14 works on a real instruction of yours, not a demonstration one.

BOTTOM LINE

You can name the 5 parts of an instruction, build one deliberately, and tell which parts change the answer from which only change the voice.

NEXT CHAPTER

How to Show AI What Good Looks Like

Chapter 14 adds the one thing that moves an output further than any rewording: a finished example to imitate. Describing the standard you want is slow and imprecise. Showing it takes one paste.

BRING WITH YOU

- ☐ The 5-part instruction you rebuilt in exercise 2
- ☐ One piece of finished work, by you or anyone, that is the standard you want matched

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

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

GAMMA · IIElevenLabs · guidde. · Outskill · Luma AI

HubSpot · attio · Synthflow · Notion · Replit

We turn attention into pipeline.

visit thecentral.ai ↗