

The AI Bootcamp

14

How to Show AI What Good Looks Like

In this chapter you'll learn

One finished sample beats a paragraph of adjectives, and how to pick the sample

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 A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes.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 fix a format that describing has failed to fix, choose an example that teaches the pattern instead of its own surface, judge how many to give for the kind of job in front of you, and recognise the point where an example has started writing the answer for you.

You can write a perfectly clear instruction and still get back something shaped wrong, over and over, because "concise" and "structured" mean one thing in your head and an average of everyone else's in the reply. Every extra adjective you add costs you a round trip. There is a move that skips the whole argument, and most people only ever use it on voice.

CONTENTS	14.1	Telling it what you want is not the same as showing it
	14.2	The sample nearest to hand is usually the wrong one
	14.3	The number you need is set by the kind of job, not by your patience
	14.4	An example can quietly take over the answer

TOPICS REQUIRED

Chapter 13. You need the 5-part instruction framework, and one prompt you have written with it that still is not returning the shape you wanted. That failing prompt is the raw material for this chapter.

AFTER THIS CHAPTER

Chapter 15 assumes you can steer the shape of an answer. It moves on to steering how the answer gets reached, which is a different problem with different tools.

14.1 Telling it what you want is not the same as showing it

A description of an output and the output itself are two different objects, and only one of them is unambiguous. Most format arguments end the moment you stop describing and paste.

A description

Words about the output you want. It has to rebuild the thing you are picturing out of your adjectives.

TRAITS

- Fast to write, and often enough
- Every word is open to interpretation
- You find out what it heard only after it answers

PROS

- + Costs nothing to try first
- + One word changes and you rerun

CONS

- "Concise" and "formal" do not mean to it what they mean to you
- Silent about everything you forgot to mention

An example

One finished output of the kind you want, pasted in whole. It answers questions you never thought to ask.

TRAITS

- Settles format, length, order and register at once
- Carries your house conventions without naming them
- Costs you the time of finding a good one

PROS

- + Removes most of the guesswork about shape
- + Gives you something specific to point at afterwards

CONS

- It copies what you show it, including your mistakes
- A badly chosen one is worse than none, see 14.4

Fig 14.1 Two ways to specify the same output. The description is cheaper to write, the example is far harder to misread. The professional habit is to start with the description and reach for the example the moment the description has failed twice.

14.1

The name. Giving examples is called few-shot prompting. None is zero-shot, one is one-shot. The jargon adds nothing to the doing.

You have met it once. Chapter 10 fixed AI-sounding writing by giving it a sample of yours. That was one use of a general technique.

Where it sits. In Chapter 13's format slot. An example is what you put there when words have run out.

Here is the moment this chapter exists for. You wrote the instruction Chapter 13 taught you, all 5 parts of it. You described the format you wanted in careful detail. What came back is still not shaped the way you meant, and nothing in the prompt is actually wrong.

The problem is that every adjective you use to describe a format is an act of faith. Concise, formal, punchy, structured, executive: you know exactly what each of those means, because you have one specific document in your head as you type it. The assistant has an average of everyone else's. The distance between your document and that average is the distance you keep trying to close by adding more words, which is why the third rewrite of an instruction so rarely works better than the second.

There is a faster move. Show it one finished output of the kind you want. Paste a real one above your request and say that this is the shape you are after. That single paste settles format, length, ordering, level of detail, house punctuation and register at the same time, including all the parts you would never have thought to specify because you have never had to say them out loud to a colleague.

You have already done this once, on one narrow thing. Chapter 10 told you that the cure for writing that sounds like AI is to hand it a sample of your own writing, and it works for exactly the reason above. This chapter takes that move off voice and points it at the rest: how a record should be laid out, how a borderline case should be decided, how much detail a summary is supposed to carry, what to write in a field when there is nothing to put in it.

The technique has a name, few-shot prompting, and you will meet it in every vendor document you read. The name matters less than the reflex: when you have described the output twice and it is still wrong, stop describing and start showing.

Two things it will not fix, so you know the edges. An example cannot make it know something it does not know, and it will not stop it being confidently wrong about a fact, which is Chapter 8's problem and Chapter 23's taxonomy. Examples fix shape. They do not fix truth.

14.2

The sample nearest to hand is usually the wrong one

Most people paste in whatever is open, or whatever they are proudest of. Both choices teach the wrong lesson, and the second one teaches it more convincingly.




Fig 14.2 Four tests an example has to pass before it earns its place in a prompt. The bar heights are visual rhythm, not a ranking: an example that fails any one of the four will quietly teach something you did not intend.

14.2

Not your best work. An exceptional sample teaches it to aim for exceptional and miss. A typical one sets the standard you actually want held across the batch.

Include the awkward one. The example showing what to do when a field is empty is worth 3 more clean ones.

Label the reject. A counter-example only works if you say it is one, and say in a line why. Unlabelled, it is just a bad example.

The example you reach for first is almost always the wrong one, because you reach for the one you are proudest of. Pasting in the single best client memo you ever wrote sets a standard nothing in the batch will meet, and worse, it teaches the particular flourishes of a piece that was unusual on purpose. Pick the one you would describe as fine. Fine is the standard you actually want held across 40 records.

Take a real job. You have a quarter of meeting notes, 40 of them, sitting in a folder in various states, and you want each one turned into the same record: client, date, what was agreed, what is outstanding, who owns the next step. Three clean examples of well-run meetings will get you 30 good records and 10 improvisations, because 10 of those meetings ended with nothing agreed, or with two owners, or with a decision that quietly reversed one from the month before.

So choose the examples so at least one of them is awkward. Show the record for the meeting where nothing was decided, with "none agreed" written into the field rather than the field deleted. That one sample does more work than the other two put together, because it is the only one that tells it what to do at the moment it would otherwise invent something. This is the difference between examples that are relevant and examples that are diverse, and the vendor guidance for all of these tools asks for both.

Send the example exactly as it would go out. Do not trim it to save space, do not swap the real names for placeholders unless placeholders are what you want back, do not clean it up first out of embarrassment. Every edit you make to an example is an instruction you did not mean to give, and the edits you make out of tidiness are the ones you will not think to look for later.

The most under-used example is the one you turned down. Paste a rejected version, label it clearly as rejected, and give one line of reason: this one is wrong because it summarises the discussion instead of recording the decision. That single line draws a boundary three more good examples cannot draw, because good examples only ever show the inside of the space, never its edge. Keep it to one, and keep the label attached to it.

None of this needs volume. One well-chosen example beats several near-identical ones, and the testing keeps landing in the same place: what improves the output is which examples you picked, not how many. Three redundant samples showing the same easy case teach one lesson, three times.

14.3

The number you need is set by the kind of job, not by your patience

Examples are not uniformly useful. They do their heaviest work on formats and categories, and their lightest on open writing, which is the opposite of where most people spend them.




Fig 14.3 Read down until a line describes the job in front of you. The last line is the one people skip: past about 5 examples the returns flatten and can reverse, so a job still failing at 5 is failing on the instruction, not on the count.

14.3

The vendors converge.

Anthropic's guidance says include 3 to 5. Google's warns that too many make it overfit to your samples. OpenAI's says add them only when a plain instruction has failed.

Order counts. It weights what it read last most heavily. Put your strongest example at the bottom of the stack.

Keep them identical. Same headings, same punctuation, same field order in every example. Variation between your samples reads as permission to vary.

How many examples to give is decided by the kind of job, not by how much effort you feel like putting in, and the honest answer is uneven enough to be worth learning properly.

For open-ended writing where you want it to sound like you, one example is usually the right number, and a second one often makes things worse rather than better. What you are transferring is register, and one real piece carries register completely. Stack up 4 and you start transferring the specific moves of those 4 pieces instead, which is section 14.4's problem arriving early.

For a fixed format you will run repeatedly, the meeting records from 14.2, a standard client update, a weekly summary that always has the same 5 fields, 2 or 3 examples is the working range. One establishes the shape. The second and third establish which parts of that shape are fixed and which are allowed to move, which one example on its own can never tell it.

For sorting things into set categories, 3 to 5, and try to spend them on one per category rather than 3 on the easy category and none on the two that are hard to tell apart. This is where examples earn the most: putting things into buckets is the task where showing it 4 decided cases outperforms any amount of describing where the line between buckets sits. It is also where the biggest gap lies between people who use this technique and people who do not.

Past about 5, the returns do not just flatten. Independent testing and the vendors' own documentation both report accuracy going down again, because a long stack of samples starts teaching their surface rather than their pattern. If you are at 5 examples and the output is still wrong, adding a sixth is the least likely thing to fix it. Go back and read your instruction: nine times in ten the examples are fine and the task sentence is ambiguous.

Two mechanical points that cost nothing. It weights what it read most recently, so put your strongest example last. And keep the formatting between examples strictly identical, down to the punctuation and the order of the fields, because any variation between your own samples is read as a signal that variation is allowed.

14.4

An example can quietly take over the answer

This technique has a real failure mode, and every professional using it hits the failure mode within a month. It looks like success, which is what makes it expensive.




Fig 14.4 Three ways a good example goes wrong. Heights are rhythm, not frequency. The tell is the same in all 3 cases: an output that resembles your sample more closely than it answers your question.

14.4

Name the pattern. Write one line saying what the example demonstrates and what it does not. Without it, everything visible in the sample is fair game to copy.

Vary the incidentals. If all 3 of your examples run to 4 bullets, you have taught length, not structure. Make them different lengths deliberately.

Try it without. On open creative work, run the request once with the example and once without, then compare. Sometimes the plain version is broader and better.

Here is the thing the enthusiastic write-ups leave out. An example is the most concrete object in your prompt, and concrete objects get copied. Not the pattern underneath the example, which is what you meant, but the surface of it: the sentence rhythm, the vocabulary, the particular turn of phrase, sometimes the name of the specific client who happened to be in the sample.

The most costly version is anchoring on a number you never intended to set. If the example you pasted pulled 4 risks out of a contract, you have taught it that this task returns about 4 risks, and the next contract comes back with 4 risks in it whether the document contains 3 or 11. Nothing looks wrong on the page. You get a clean, well-formatted, confident answer that is missing 7 things, and you will only find that out by reading the source yourself.

The third version shows up on creative work. Ask for 10 subject lines with one sample attached and you will often get 10 variations of the sample rather than 10 options, because the example has fixed the range as well as the format. That is a fair trade when you want consistency and a bad one when you want choice, and knowing which of the two you are asking for is the whole skill.

The fixes are small. Say in one line what the example is there to demonstrate and what it is not: this shows the field order and the level of detail, not the length or the subject matter. Deliberately vary everything you do not want copied, so if length is not the lesson, make your examples visibly different lengths. And on anything open-ended, run it once with the example and once without, and actually compare the two rather than assuming the version with more effort in it won.

One more, worth knowing before you meet it. On the deeper thinking settings you met in Chapter 6, a stack of examples has been reported to do worse than a single clean instruction, apparently because the effort goes into matching the shape of the samples instead of working the problem. That is Chapter 15's territory and it will pick it up there. For now, treat a long example stack and a heavy reasoning mode as two tools that do not always want to be used together.

Go and find the prompt you brought from Chapter 13, the one still coming back wrong. Add one example to it, chosen against the 4 tests in 14.2, and write the line saying what it demonstrates.

READING ENDS HERE

Everything before was reading. Everything after is doing.

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

Do these before Chapter 15

3 exercises. About 15 minutes

Open the folder where one of your recurring outputs lives: status reports, client records, meeting notes, whatever repeats. Pick the example you would use, then write down why it is typical rather than your best.

Take that example and write the two lines that have to travel with it. Being able to fill the right-hand box is what stops it copying the wrong thing.

WHAT THIS EXAMPLE DEMONSTRATES

WHAT IT DOES NOT DEMONSTRATE

Run your Chapter 13 prompt twice: once as it stands, once with your example and its two lines pasted in above it. Write one sentence on what changed, and one on anything the example copied that you did not want.

Chapter 14 on one page

THE MOVE

Paste one finished output of the kind you want, and say what it is demonstrating. That is the whole technique, known as few-shot prompting

CHOOSING ONE

Typical rather than exceptional · includes the awkward case · complete, exactly as you would send it · labelled with what it shows

HOW MANY

1 for open writing in your voice · 2 to 3 for a fixed format · 3 to 5 for sorting into categories · past 5 the returns flatten and can reverse

THE COUNTER-EXAMPLE

One version you rejected, labelled as rejected with a one-line reason, draws a boundary that another good example cannot

NEEDED BEFORE CHAPTER 15

One example you have actually tested, plus the line saying what it demonstrates and what it does not

BRING TO CHAPTER 15

The example you tested and the two lines that travel with it. Chapter 15 stops working on the shape of the answer and starts working on how the answer gets reached.

BOTTOM LINE

You can now fix with one example what a paragraph of adjectives could not, and spot the point where the example starts writing the answer.

NEXT CHAPTER

Make AI Think Before It Answers

Chapter 15 moves from the shape of the answer to the route it takes to get there: asking for the working, breaking a question into steps, and telling the difference between real reasoning and a plausible account of it written afterwards.

BRING WITH YOU


- ☐ The example you tested in this chapter, and the line saying what it demonstrates
- ☐ One question you have asked that came back confident and wrong

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

GAMMA

IIElevenLabs

guidde.

 Outskill

 **Luma AI**

HubSpot

 **attio**

S Synthflow

 Notion

 **Replit**

We turn attention into pipeline.

visit thecentral.ai

