The AI Bootcamp

15

Make AI Think Before It Answers

In this chapter you'll learn

Why "think step by step" stopped earning its place, and the 3 instructions that replaced it

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 A LinkedIn carousel here, a YouTube tutorial there, a $2,000 cohort course that goes out of date before the refund window closes.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 tell the difference between a setting that buys depth and an instruction that buys direction, write the 3 asks that still change how an assistant approaches a problem, and read the reasoning it hands back for the parts that are genuinely checkable.

"Think step by step" is the most repeated piece of prompt advice in circulation, and it has quietly stopped earning its place: the assistants now do that work before they show you anything, and their own vendors publish guidance telling you to stop asking for it. What never got absorbed into the product is the part that matters to you, which is seeing what an answer assumed, what it ruled out, and why. This chapter replaces a dead ritual with the technique that took its place.

CONTENTS

15.1 The setting is theirs. The instruction is yours

15.2 The advice you have heard 3 times has quietly expired

15.3 3 asks that still change the answer

15.4 Read the working, do not admire it

TOPICS REQUIRED

Chapter 13's framework and Chapter 14's worked examples. Bring one prompt you have already written in that shape, ideally one where the answer came back confident and you had no way to check it. Chapter 6's tour of the Think button and the effort settings is assumed here, not repeated.

AFTER THIS CHAPTER

Chapter 16 assumes you can get one high-quality answer out of a single instruction, and moves to the jobs that are too big for one answer at all.

15.1 The setting is theirs. The instruction is yours

Chapter 6 handed you the switches. This chapter is about the words, which work the same whether the switch is on, off, or absent from your assistant entirely.

The setting

A control the vendor gives you. It buys depth: how long the assistant works before it speaks.

TRAITS

· The Think button, the effort level, the model picker
· Chapter 6 covered all of it
· Differs by assistant, and keeps changing

PROS

+ One click, no wording to get right
+ Genuinely helps on hard problems

CONS

– Buys time, not direction
– Nothing about the answer you get back changes shape

The instruction

Words you type. It buys direction and shape: what the answer must consider, and what it must show you.

TRAITS

· Works on every assistant, on any setting
· Survives a model change next month
· Costs you one extra sentence

PROS

+ Makes the answer checkable instead of take-it-or-leave-it
+ You control what gets exposed

CONS

– You have to write it, every time, until Chapter 17
– Adds length and a little latency

Fig 15.1 Two different purchases. More thinking time spent on the wrong question produces a more thorough wrong answer, which is why the sentence matters more than the setting.

15.1

Chapter 6's job. The switch. Modes, effort settings and model choice: how long it works.

This chapter's job. The instruction. What it must weigh and what it must show: how it works.

Not Chapter 16. Splitting one big job across several messages is the next chapter. This one is about the quality of a single answer.

Chapter 6 gave you the switches: the Think button, the effort setting, the model picker. Those belong to the vendor, they differ by assistant, and they all buy the same thing, which is time. This chapter is about the one control the vendor does not own, which is what you write in the box.

The distinction is worth 10 seconds of your attention because it decides what to reach for. A setting buys depth. An instruction buys direction. An assistant given twice as long to think about a question you asked badly will produce a more thorough answer to the wrong question, and it will sound better while doing it.

Here is the situation this chapter exists to fix. You ask something that genuinely matters, a pricing call, a hiring choice, whether a contract clause is a problem. Back comes a confident recommendation, well written, with nothing in it you can test. You cannot see what it assumed about your business, what else it weighed, or which single fact the whole thing rests on. You either take it on trust or you throw it away, and neither of those is using the tool properly.

What changes that is not a harder setting. It is requiring the answer to arrive in a shape that exposes its own logic: what it assumed, what it considered and rejected, and which step is load-bearing. That is an instruction, and it works identically on a free ChatGPT account, a Claude window, a Gemini tab and a Copilot pane.

One boundary before you start. This chapter is about the quality of thinking inside a single answer. When the job is genuinely too big for one answer, no instruction rescues it and the fix is sequencing, which is Chapter 16.

15.2 The advice you have heard 3 times has quietly expired

Adding "think step by step" used to make answers measurably better. On a current assistant it mostly buys you latency, and all 3 vendors who publish guidance on this now say so themselves.

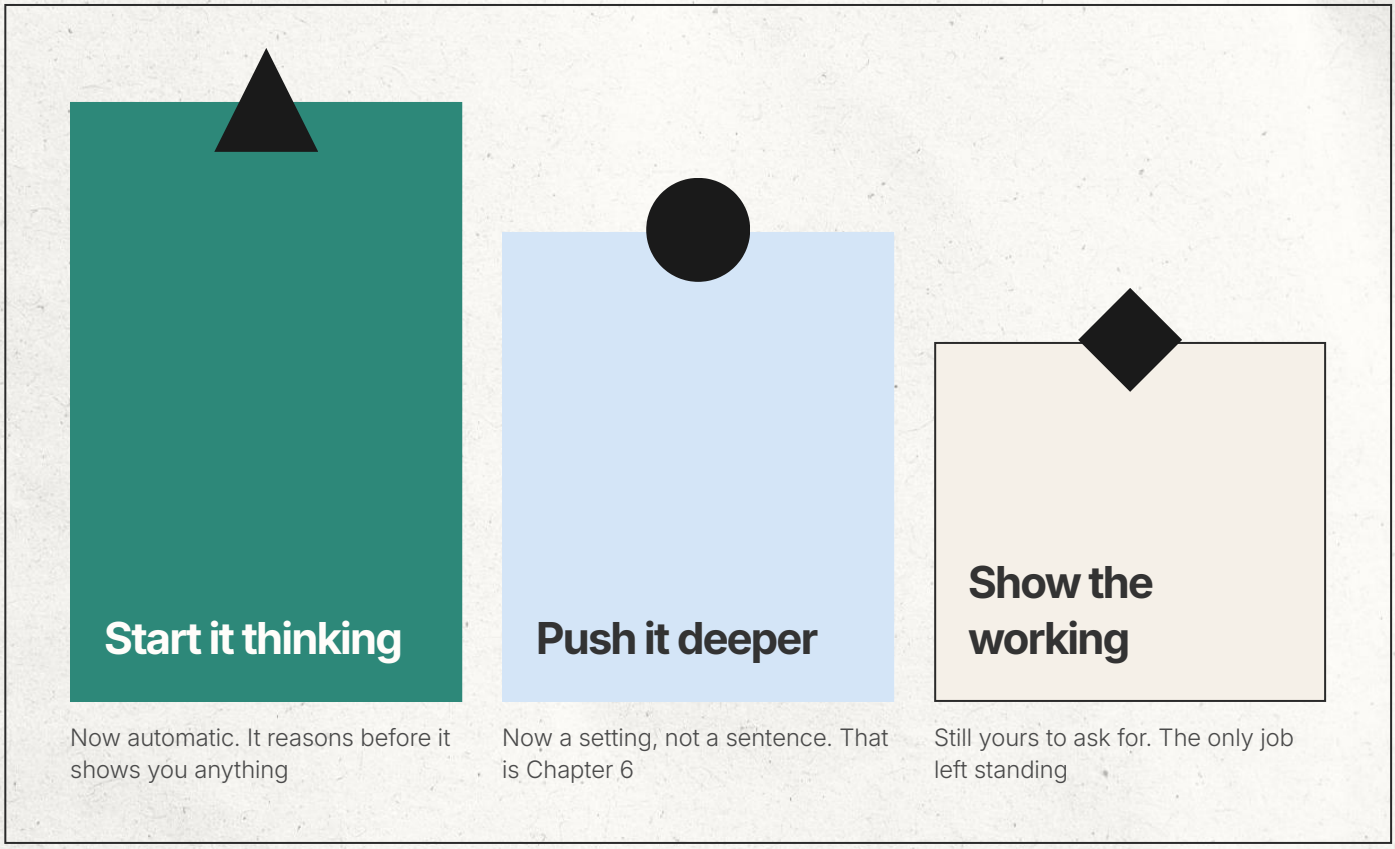


Fig 15.2 Three jobs one phrase used to do. Two of them got absorbed into the product and are no longer yours to ask for. The third never was, and it is the whole of section 15.3. Bar heights are visual rhythm, not a measurement.

15.2

ChatGPT OpenAI's reasoning guidance says to avoid chain-of-thought prompts outright: the model reasons internally, so asking it to think step by step is unnecessary.

Claude Anthropic advises general instructions over prescriptive steps, on the grounds that the model's own reasoning frequently exceeds a hand-written plan.

Gemini Google says it is generally not necessary to have the model outline, plan or detail its reasoning steps in the answer itself.

Copilot Microsoft publishes no equivalent warning. Its guidance stays general: be clear and specific, give context, name the format you want.

You have heard it at three conferences and read it in a dozen posts: add "think step by step" and the answers get better. That was true, it was well evidenced, and it is now largely obsolete. The interesting part is that the people telling you so are the vendors themselves.

OpenAI's own guidance for its reasoning models tells you to avoid chain-of-thought prompts, because the model performs that reasoning internally and asking for it is unnecessary. Google's Gemini guidance says it is generally not necessary to have the model outline, plan or detail its reasoning steps in the response at all. Anthropic goes furthest and says to prefer general instructions over prescriptive steps, because the model's own reasoning frequently exceeds what a person would have written out for it. Microsoft is the exception: its Copilot advice stays at the level of be clear, give context, name your format.

The measurements agree. A Wharton research group ran the test properly on 198 PhD-level science questions, 25 trials per condition. Older models with no reasoning mode of their own gained a lot from the phrase, in one case over 13 points. Reasoning-capable models gained roughly 3 points, one of them got slightly worse, and every one of them took between 20% and 80% longer to answer. Their conclusion was that chain-of-thought prompting is not universally optimal, which is an academic's way of saying the ritual no longer pays for itself.

Read that with the caveat it deserves. Multiple-choice science questions are not your board memo, and a benchmark is a proxy, not your work. But the direction is consistent across independent measurement and 3 separate vendors, which is about as much agreement as this field ever produces.

The reason it broke is worth understanding, because it tells you what to do instead. The phrase was scaffolding for a model that would otherwise commit to the first plausible answer and then defend it. Current assistants do that deliberation before they show you a word of the reply. Telling one to think step by step is telling it to do a thing it already did, and on some setups it makes it do the whole performance again, out loud, in the answer, where it costs you reading time and buys you nothing.

You are not being scolded for having used it. It cost a little time and it rarely did harm. But if it is the only move you have for a hard question, you are a technique behind, and the next section is the one that replaced it.

15.3 3 asks that still change the answer

The move that survived is not a phrase about thinking. It is naming a piece of work you want done, and never the method for doing it.

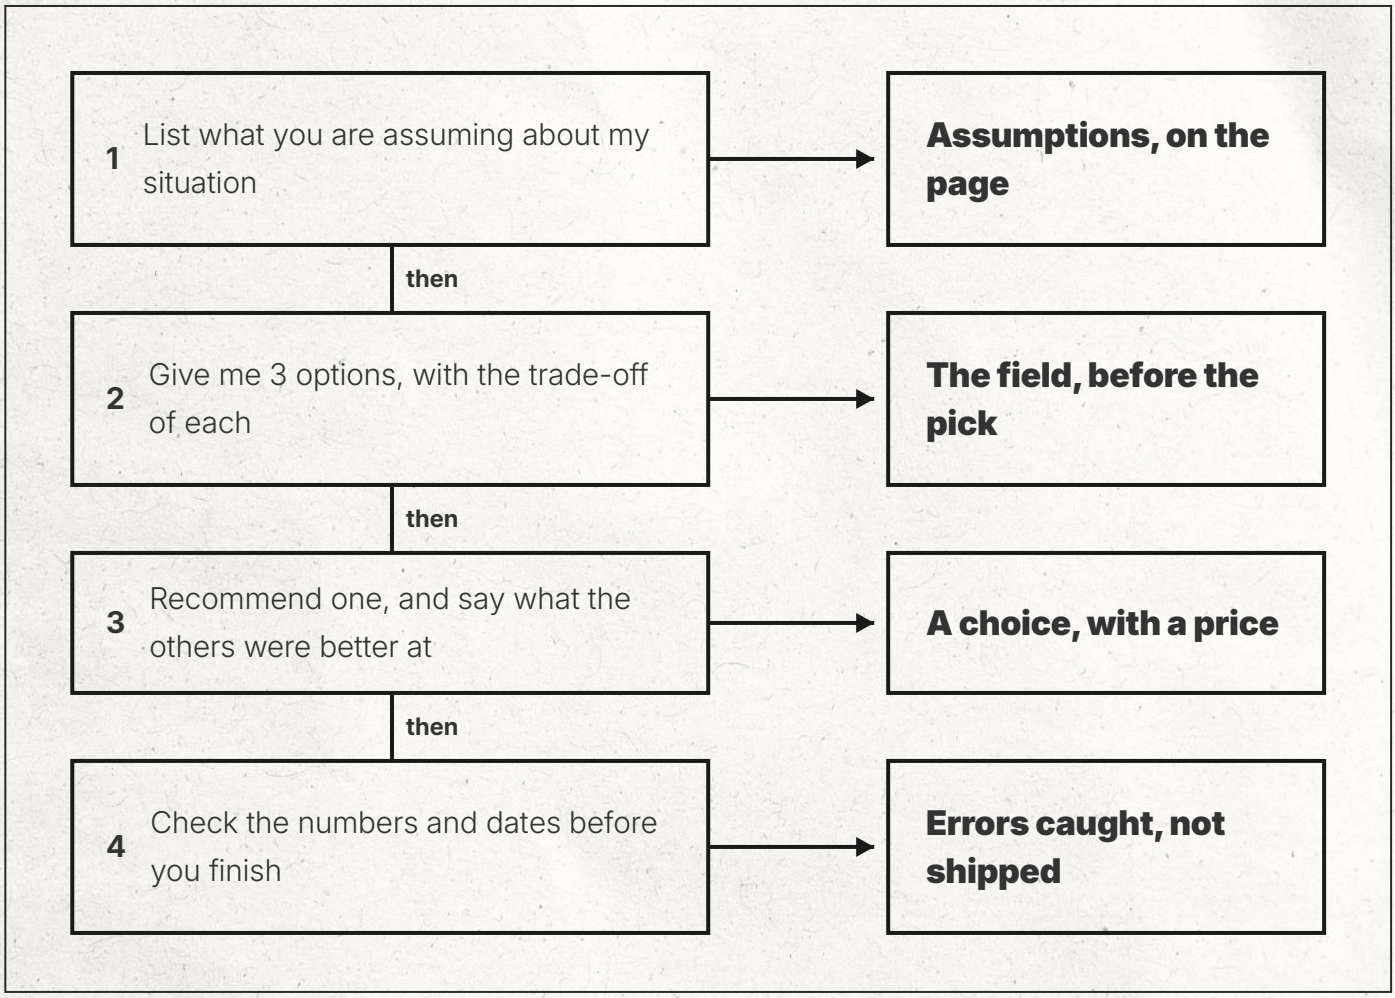


Fig 15.3 One instruction, 4 clauses. Every clause names an output you can read and argue with. Not one of them tells the assistant how to reason.

15.3

Assumptions. Your question was underspecified. It answered anyway. Ask to see what it quietly filled in.

Options. One recommendation is a conclusion. Three options and a pick is a decision you can audit.

Working. Do not ask it to explain itself. Name the specific item you intend to check.

Here is what to type instead, and it is a shape rather than a phrase: name the work you want done, and say nothing about how to do it.

The first ask is assumptions, and for a senior reader it is the highest-yield sentence in this book so far. "Before you answer, list the assumptions you are making about my situation, and mark the ones you are least sure of." Your question was underspecified, because every real question is. It answered anyway, which means it filled the gaps silently and you never saw the filling. This clause drags it into the open, and roughly half the time you read the list and realise the problem was your question, not its answer.

The second ask is options before the pick. "Give me 3 options with the main trade-off of each, then recommend one and say what the other two would have been better at." The failure this fixes is specific and common: an assistant settles on the first workable answer and spends the rest of the reply building a case for it. Requiring the field before the choice stops a reply from being an argument for a conclusion it reached in its opening sentence. Making it argue against a position, rather than merely lay out alternatives, is Chapter 21.

The third ask is the working, and this is where most people go wrong by asking for too much. "Show your reasoning" gets you narration: a pleasant account of how a sensible person might have approached it. Name the load-bearing item instead. The numbers you used and where each one came from. The criterion you applied. The single assumption this recommendation would collapse without. Ask for what you would need in order to prove it wrong.

There is a fourth clause worth adding when the answer contains anything checkable. Google's guidance is explicit that you should include verification steps in the instruction, and Anthropic's is that appending a line asking it to verify the answer against your criteria catches errors reliably. One honest caveat: the newest models increasingly do this unasked, and Anthropic now tells developers moving to its latest model to delete inherited verification instructions rather than rewrite them, because piling them on produces a model that checks, re-checks, and then reports at length on having checked. Add the clause when there is arithmetic, a date, a quoted line or a rule with conditions. Leave it off otherwise.

Take the last question you asked an assistant and re-ask it with the assumptions clause on the front. Nothing else changed. Read what it filled in.

15.4 Read the working, do not admire it

What comes back when you ask for reasoning is a written account of an answer. It is not a recording of how that answer was produced, and the difference decides what you can safely do with it.

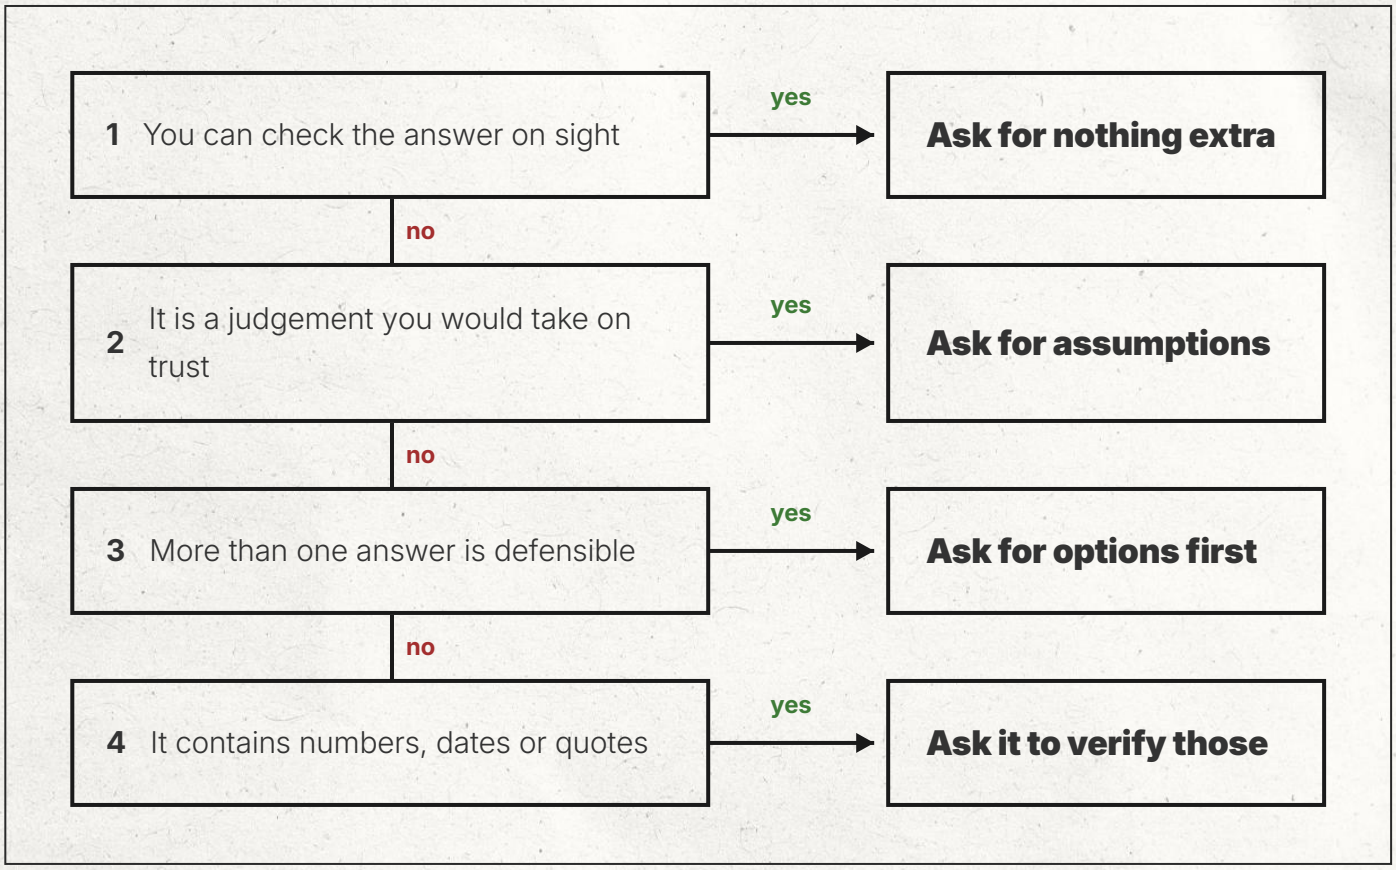


Fig 15.4 How much to ask for. Read from the top and stop at the first line that is true: every clause below it costs you length, latency and a longer thing to read.

15.4

It is an account. The written reasoning explains the answer. It does not record how the answer was reached.

Check claims, not the mind. A wrong assumption is wrong whatever actually drove the conclusion.

Length is not rigour. Asking for the working buys you words. Only some of them are checkable.

Now the part that keeps you honest. When an assistant shows you its reasoning, you are reading a written account of an answer, not a transcript of how that answer came about, and treating the two as the same thing is how a careful person gets fooled by a thorough-looking reply.

Anthropic's own researchers tested this directly. They planted a hint in a question, confirmed the model had used it, then checked whether the visible reasoning admitted to using it. Across the reasoning models they tried, it mentioned the hint that had actually changed its answer well under a fifth of the time. Their concern was safety monitoring rather than your Tuesday, but the narrow lesson transfers cleanly: the trace tells you what the answer claims, not why the answer came out that way.

Which is still worth having, for a narrower reason than most people assume. An assumption you disagree with is wrong whether or not it drove the conclusion. A number is right or wrong on its own terms. An option dismissed on a ground you know to be false is a bad dismissal, whatever else was going on. You are checking the claims in front of you, which is a normal editorial act. You are not auditing a mind, and you cannot.

The costs are real and worth naming. Asking for the working makes every answer longer, slower and more confident-sounding, and length reads as rigour when it is nothing of the kind. On the newest models, inherited reasoning and verification instructions produce visible padding rather than better answers. That is why the figure above is a rule you read from the top and stop, and not a checklist you apply to everything.

So leave it off for short factual questions, for drafting, for anything you can eyeball. Spend the extra sentence when the answer is a judgement you would otherwise accept on trust, which is exactly the category where an assistant is most useful to you and least checkable.

One thing to do before Chapter 16. Find a recommendation you accepted from an assistant this month and ask it, in the same thread, what it was assuming about your situation when it gave you that. Read the assumptions before you read anything else, and see whether you would still have accepted it.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 16.

Do these before Chapter 16

3 exercises. About 15 minutes

Write down the last recommendation an assistant gave you that you accepted without checking. Then write one sentence on what it must have assumed about your situation in order to say that.

Re-ask a real question from this week with the assumptions clause on the front: "Before you answer, list the assumptions you are making about my situation." Fill in both boxes from what comes back.

WHAT IT ASSUMED	WHAT IS ACTUALLY TRUE

Take one prompt you saved in Chapter 13 or 14 and rewrite its closing line so it asks for 3 options and a recommendation instead of an answer. Write the new closing line here.

Chapter 15 on one page

THE EXPIRED PHRASE	"Think step by step" was scaffolding for models that would not reason unprompted. OpenAI, Google and Anthropic now all advise against scripting the method; Microsoft publishes no equivalent guidance
THE ASSUMPTIONS CLAUSE	"Before you answer, list the assumptions you are making about my situation, and mark the ones you are least sure of." The single highest-yield sentence in this chapter
OPTIONS BEFORE THE PICK	"Give me 3 options with the main trade-off of each, then recommend one and say what the other two would have been better at"
NAME WHAT YOU WOULD CHECK	Never ask it to explain its reasoning. Ask for the numbers and their sources, the criterion applied, or the one assumption the recommendation rests on
THE TRACE IS AN ACCOUNT	Written reasoning explains an answer, it does not record how the answer was produced. Check the claims in it, not the mind behind it

BRING TO CHAPTER 16

The rewritten prompt from exercise 3, with the assumptions clause on the front. Bring one more thing too: a job you already tried to do in a single prompt and watched go thin.

BOTTOM LINE

You can now ask for reasoning in the one way that still changes the answer, and read what comes back for the parts you can actually check.

NEXT CHAPTER

5 Steps for Jobs Too Big for One Prompt

Chapter 16 asks the question this one deliberately left alone: what to do when the job is too large for a single answer. One giant prompt fails in a predictable way, and no amount of better wording rescues it. The fix is sequencing.

BRING WITH YOU

- ☐ Your rewritten prompt from exercise 3
- ☐ One job you tried to do in a single prompt and watched thin out

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 ↗