

The AI Bootcamp

6

**Let it think harder
when the question
deserves it**

In this chapter you'll learn

What the thinking setting really trades, where it hides on each assistant, and the 4 questions that decide when to spend the time

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 A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes.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 will know what the thinking setting on your assistant actually trades, where to find it whatever it has been renamed to since this went to print, and a 4-question test for when it is worth the wait, so you stop reaching for it out of superstition.

Chapter 2 told you this control existed and told you not to touch it yet. This is where that promise comes due. Most people land in one of two habits: never pressing it, and grinding through a genuinely hard question on a setting built for easy ones, or pressing it for everything and spending the week waiting. Both cost you. The difference between them is one question you can ask yourself in 2 seconds.

CONTENTS	6.1	Not a smarter setting, a slower one
	6.2	Every assistant renames this control
	6.3	4 questions that tell you it is worth the wait
	6.4	Most of what you ask does not need it

TOPICS REQUIRED

Chapter 5. You need your assistant open and one clearly written instruction from that chapter, something with context, a task and a format, so that when an answer disappoints you it is the setting at fault and not the ask.

AFTER THIS CHAPTER

Chapter 7 assumes you can tell a question that needs the slower setting from one that does not, and that you have stopped reaching for it by reflex.

6.1 Not a smarter setting, a slower one

Every assistant now has some way of telling it to take longer before it answers. What you are buying with that wait is care, not a better brain.

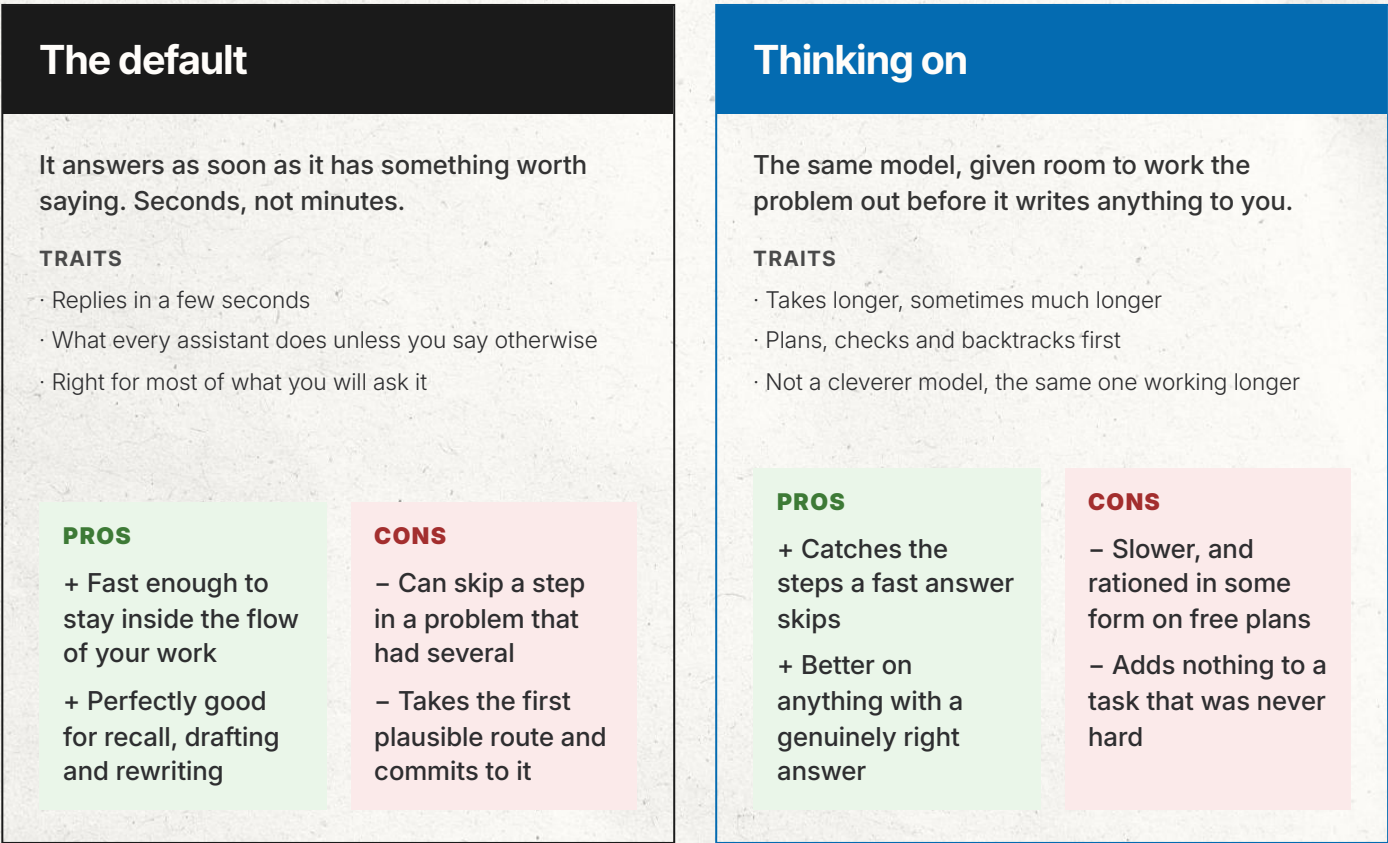


Fig 6.1 The same model sits on both sides of this figure. What the slower setting buys is care, and care only helps a question that had somewhere to go wrong.

6.1

Same model. On most tools the slower setting is the assistant you already had, given more time. It is not an upgrade.

The cost is time. Seconds on the default. Ten seconds to several minutes with thinking on, depending how far the setting goes.

What it will not fix. It cannot come to know something it never knew. Checking the answer is still your job, and Chapter 8 is about that.

Here is what that button actually does, because the name on it promises something it does not deliver. Turning on thinking does not hand you a cleverer assistant. On most tools it hands you the same assistant you already had, given room to work the problem out before it starts writing to you.

Left on the default, it begins writing almost as soon as you press send. It takes the first plausible route through your question and commits to it. For the large majority of what you will ask, that is exactly right, because most questions do not contain a route that can go wrong.

With the setting on, it does something closer to working on paper first. It lays out the steps, tries one, sees that it does not work, goes back. You will not usually see that work happening. What you see is the wait, and then an answer that has already been checked against itself once before it reached you.

The cost is time, and it is not trivial. A default answer arrives in a few seconds. A slower one runs from about 10 seconds on the light settings to several minutes on the heavy ones. Free plans also ration it, in one way or another: a weekly allowance, a queue behind paying customers, or a lighter model doing the thinking. You have a budget here, even if nobody shows you the meter.

One thing the wait does not buy you. It does not make the assistant know something it never knew. Extra time helps it work things out, not look things up, and a carefully reasoned answer can be confidently wrong in exactly the way a quick one can. Chapter 8 is entirely about catching that, and nothing in this chapter replaces it.

6.2 Every assistant renames this control

It will not be called the same thing on your screen as on your colleague's, and it will have moved again by the time you read this. The shape underneath does not change.




Fig 6.2 Four shapes, one question underneath all of them: how long should it work before it answers. Bar heights here are visual rhythm, not a ranking. The words on your screen will differ from these, and that is the point of learning the shapes.

6.2

ChatGPT Free accounts get a Think button that applies to the one message you are typing. Paid accounts get a sliding effort setting instead.

Claude Mostly decides for itself: it reads the question and thinks longer when the question warrants it. A manual toggle appears and disappears.

Gemini A picker above the chat, Fast, Thinking and Pro, with a separate thinking level under it on some accounts.

Copilot Auto by default, with Quick response and Think deeper as the manual choices. Think deeper is capped at around 10 seconds.

The control exists on all four assistants. What it is called, where it sits, and whether you get a say in it at all is the part that keeps moving, and any book naming the exact label is out of date within months. So learn the 4 shapes instead, and read whatever is on your screen as one of them.

A switch is a single button beside the box you type into, on or off, applying to this message and nothing after it. ChatGPT works this way on the free tier: press it, and the same model spends longer on that one question. Nothing carries over to your next one, which is either a nuisance or a useful discipline, depending on your temperament.

A dial is the paid version of the same idea: an effort setting running from low to high that stays where you put it. Higher settings buy more internal work rather than a longer reply, which catches people out the first time. Turning it up is not how you ask for more detail. That is a wording problem, and Chapter 5 already solved it.

A picker is a different thing in similar clothing, and it is worth getting straight. Gemini's Fast, Thinking and Pro menu is not one model at three speeds. Those are genuinely different models, and the slower ones are also the more capable ones. Choosing there changes what you are talking to, not only how long it takes to answer.

Auto is the fourth shape and the direction all of them are drifting. Copilot defaults to it. Claude leans on it hardest: rather than give you a toggle, it reads each question and decides for itself whether to think and how much, which is why many of its users never see a dial at all. Auto is right more often than people expect, and it removes a decision you would mostly get wrong anyway.

Go and find yours now. It sits beside the box you type into, or just above it. Then leave it alone. Knowing where it is matters. Reaching for it every time does not.

6.3 4 questions that tell you it is worth the wait

Here is the test, and it replaces every rule of thumb you will read online. Ask whether the question has somewhere to go wrong in the middle, because that is the only thing extra thinking protects you from.

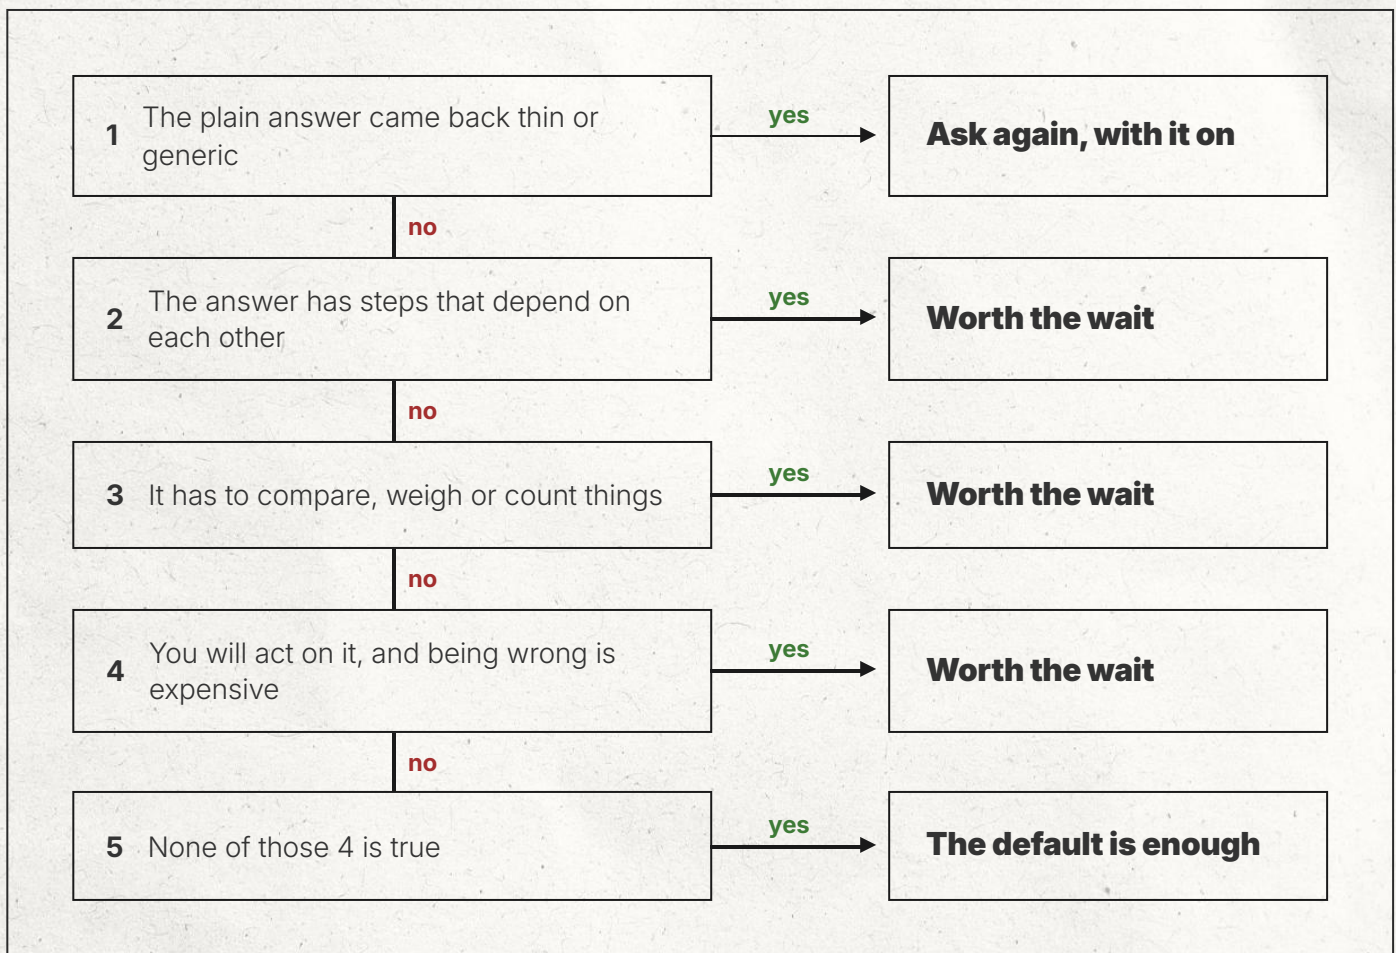


Fig 6.3 Read down and stop at the first line that is true. Three of them land on the same answer on purpose: any one test on its own would miss the cases the other two catch.

6.3

The test. Does the question have somewhere to go wrong in the middle? If it does not, the slower setting has nothing to catch.

Stakes count. The same question earns the wait if you are going to act on the answer, and does not if you are only going to read it.

Decide afterwards. Ask on the default first. Re-ask with thinking on when what comes back is thin. That is the cheapest test there is.

Here is the test worth memorising: does this question have somewhere to go wrong in the middle? Extra thinking protects you from mistakes made along the way to an answer. Where there is no along the way, it protects you from nothing at all, and you have paid a minute for the privilege.

Steps that depend on each other are the clearest case. Work out the cost per unit, then find which discount tier that cost falls into, then compare the result against last year. Miss one step and everything after it is wrong, quietly and plausibly. That is the precise shape of failure the slower setting exists to catch, and it is the shape a fast answer is worst at.

Comparison and counting is the second case: weighing 5 options against 4 criteria, reconciling two lists, working out which of these numbers is genuinely the largest. Arithmetic deserves particular suspicion, because a fast answer will hand you a figure that looks entirely right and is not. Chapter 19 goes into exactly where it will lie to you about numbers.

Stakes are the third case, and this one is a judgement about you rather than a property of the question. If you are going to forward the answer, act on it, or put money behind it, an extra minute is cheap insurance. If you are going to read it, nod, and move on, it is not.

The fourth case is the one you will actually use most, and it is the one nobody teaches: you do not have to decide in advance. Ask on the default. If what comes back is thin, generic, or has visibly skipped something, ask the same question again with the setting on. Try that on the next answer that disappoints you rather than deciding anything now, and within a week you will know which of your questions need it.

6.4 Most of what you ask does not need it

This is the part that will save you the most time, and the reason this chapter is not an advert for the button. The default is the correct setting for the large majority of your working day.




Fig 6.4 4 asks the slower setting does not improve. Thinking longer helps a model work something out, and none of these 4 has anything to work out. Bar heights are rhythm only, not a measure of anything.

6.4

The default is fine. For most of a working day the fast answer is also the correct answer. Treating the slow one as the serious option costs you hours.

Longer is not righter. Past a point, more reasoning has been shown to cost accuracy rather than add it. This is measured, not folklore.

Spend the allowance.

Whatever your plan ratios, spend it on the questions that pass the test in 6.3, not on tone edits.

Now the honest part, and the reason this chapter is not an advert for the button: most of what you ask all day does not need it. The default is the right setting for the large majority of your work, and treating the slower one as the serious, professional option is how people end up waiting 3 minutes for a subject line.

Recall is the weakest case for it. Asking what a term means, or who runs a company, does not send it off to look anything up. It is the same stored knowledge, examined for longer. Longer examination does catch some of its own contradictions, which is worth a little, but it is not the same as being right, and it does not excuse you from checking.

Drafting gains almost nothing, because a first draft is not a problem with a correct answer. You were going to rewrite it regardless. The same goes for rewriting and for tone: shorter, plainer, warmer, less corporate. There is nothing in any of those to work out, only something to do, and doing it fast is strictly better.

There is a subtler reason not to overuse it. Researchers studying these models have found the relationship between thinking length and accuracy is not a straight line. Past a certain point, longer reasoning starts to hurt, and there are documented cases of a model reaching the right answer early, carrying on thinking, and talking itself out of it. More effort is not more correct, which is worth remembering the next time you are tempted to turn the dial to maximum on principle.

Then the practical argument. Whatever your plan gives you of the slow setting, you want it spent on section 6.3 questions rather than on tidying an email. Pick one task this week that genuinely passes the test, give that one the wait, and leave everything else on the default. Notice how rarely you miss it.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 7.

Do these before Chapter 7

3 exercises. About 15 minutes

Find the control on your own assistant. Write down exactly what it is called on your screen, where it sits, and which of the 4 shapes from 6.2 it is.

Take one real question from your work that has steps depending on each other. Ask it twice: once on the default, once with thinking on. Note what actually changed between the two.

DEFAULT ANSWER

THINKING ANSWER

Name one recurring task of yours that passes the test in 6.3, and one that clearly fails it. The second one matters as much as the first.

Chapter 6 on one page

WHAT THE SETTING DOES

Trades time for care. Usually the same model, given room to work the problem out before it writes to you

THE 4 SHAPES

A switch (one message) · a dial (effort, stays set) · a picker (a different model, not just more time) · Auto (it decides)

WHAT IT IS CALLED

ChatGPT: a Think button, or an effort slider if you pay · Claude: mostly automatic · Gemini: Fast, Thinking, Pro, plus a thinking level · Copilot: Auto, Quick response, Think deeper

THE TEST

Turn it on when the steps depend on each other, when it has to compare or count, or when being wrong is expensive. Otherwise the default is enough

THE RETRY MOVE

Do not decide in advance. Ask on the default, and re-ask with thinking on only when the answer comes back thin

BRING TO CHAPTER 7

The name of the control as it appears on your screen, and one question you have now run both ways. Chapter 7 changes what your assistant is working from, not how hard it works.

BOTTOM LINE

You know what the thinking setting trades, where it lives on your assistant, and the 4 questions that decide when to spend the time.

NEXT CHAPTER

Give it your own material to work from

Chapter 7 stops asking your assistant to work from whatever it happens to know. You give it your own material instead: the document, the deck, the spreadsheet, the screenshot, and ask your question against that.

BRING WITH YOU

- ☐ The control you found in this chapter, by the name it has on your screen
- ☐ One document, deck or spreadsheet you had to read closely this week

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

