

The AI Bootcamp

7

How to Get Answers From Your Own Files

In this chapter you'll learn

Paste it or attach it, what each format really gives you, and what it quietly failed to read

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 A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes.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 get your own documents, spreadsheets and screenshots in front of your assistant, choose between pasting and attaching without thinking about it, and tell the difference between an answer built from your material and one built from general knowledge.

Up to now your assistant has answered every question from general knowledge. It has read an enormous amount about how contracts, board papers and reorganisations usually work, and nothing whatsoever about yours. That gap is why so many answers come back sounding competent and useless at the same time. Closing it takes one click, and it is the biggest single jump in usefulness in this track.

CONTENTS	7.1	Give it the document, it stops guessing
	7.2	Paste it if it is short, attach it if it is a file
	7.3	Documents, sheets and screenshots go in
	7.4	It will not tell you what it did not read

TOPICS REQUIRED

Chapters 4 to 6. You need your open account, the habit of writing a clear instruction from Chapter 5, and one real document from your own work sitting somewhere you can find it: a contract, a report, a deck, a spreadsheet, or the screenshot you took last week.

AFTER THIS CHAPTER

Chapter 8 assumes you have watched it work from a source you can check. Verifying an answer is a different exercise when the original is sitting right there.

7.1 Give it the document, it stops guessing

The same request, asked twice, produces two completely different answers depending on whether the thing you are asking about is attached. The gap is larger than almost anyone expects.

Without your material	With your material
It answers from general knowledge: how documents like yours usually work. TRAITS · Describes the category, not your case · Hedges, because it cannot see the specifics · Reads as a competent checklist PROS + Instant, no setup, no file to find + Genuinely useful when you want the general shape of a thing CONS – Tells you what you already knew – Any specific it offers is a guess dressed as a fact	It answers from the document you gave it, and quotes the parts it used. TRAITS · Names clauses, figures and dates from your file · Contradicts you when the file does · You can check every claim against the source PROS + Answers only your case, not the category + Wrong answers become findable, because the source is right there CONS – You have to have the file, and be allowed to upload it – It may read only part of a long one, see 7.4

Fig 7.1 The same question, asked twice. The right-hand answer is not smarter. It just had access to the thing you were asking about.

7.1

The tell. A generic answer names categories. An answer from your material names things only your document contains.

Not smarter. Attaching a file changes what it can see, not how well it thinks.

One rule. If the answer depends on a document, the document goes in the chat.

Here is what has been missing from every conversation you have had so far. Your assistant has answered from general knowledge, which is a very large amount of reading about how contracts, board papers and supplier renewals usually work, and nothing at all about yours.

Take a supplier renewal that lands on your desk. Ask, with nothing attached, what you should push back on. You get a good answer: check the auto-renewal, look at the notice period, watch the price uplift, confirm the support terms. Every line is true. Not one line is about your contract. You could have written it yourself, which is the polite way of saying it was worth nothing.

Now attach the contract and ask the identical question. The reply tells you the notice window runs to 90 days and you are 11 days inside it, that the uplift is written as inflation plus 4% rather than the flat 3% you remembered, and that a clause halfway down quietly moves support to business hours only. None of that is general knowledge. It is your document, read carefully and handed back to you.

The assistant did not get cleverer between those two answers. It stopped guessing. A generic reply is what any well-read stranger gives you when they cannot see the thing you are describing, and no amount of better wording in your question fixes it, because the problem was never the wording.

This is the cheapest improvement available anywhere in this course. It costs one click, it works on every one of the four assistants, and it works on the free tier of all of them.

One thing to settle before you start uploading work documents: what is and is not safe to put in the box at your company is a real question with a real answer, and Chapter 9 gives it. For this chapter, practise on something you would be comfortable emailing to a colleague.

7.2 Paste it if it is short, attach it if it is a file

There are two ways to get material in and beginners agonise over the choice. The rule is simpler than the agonising: paste what you can read on a screen, attach what already exists as a file.

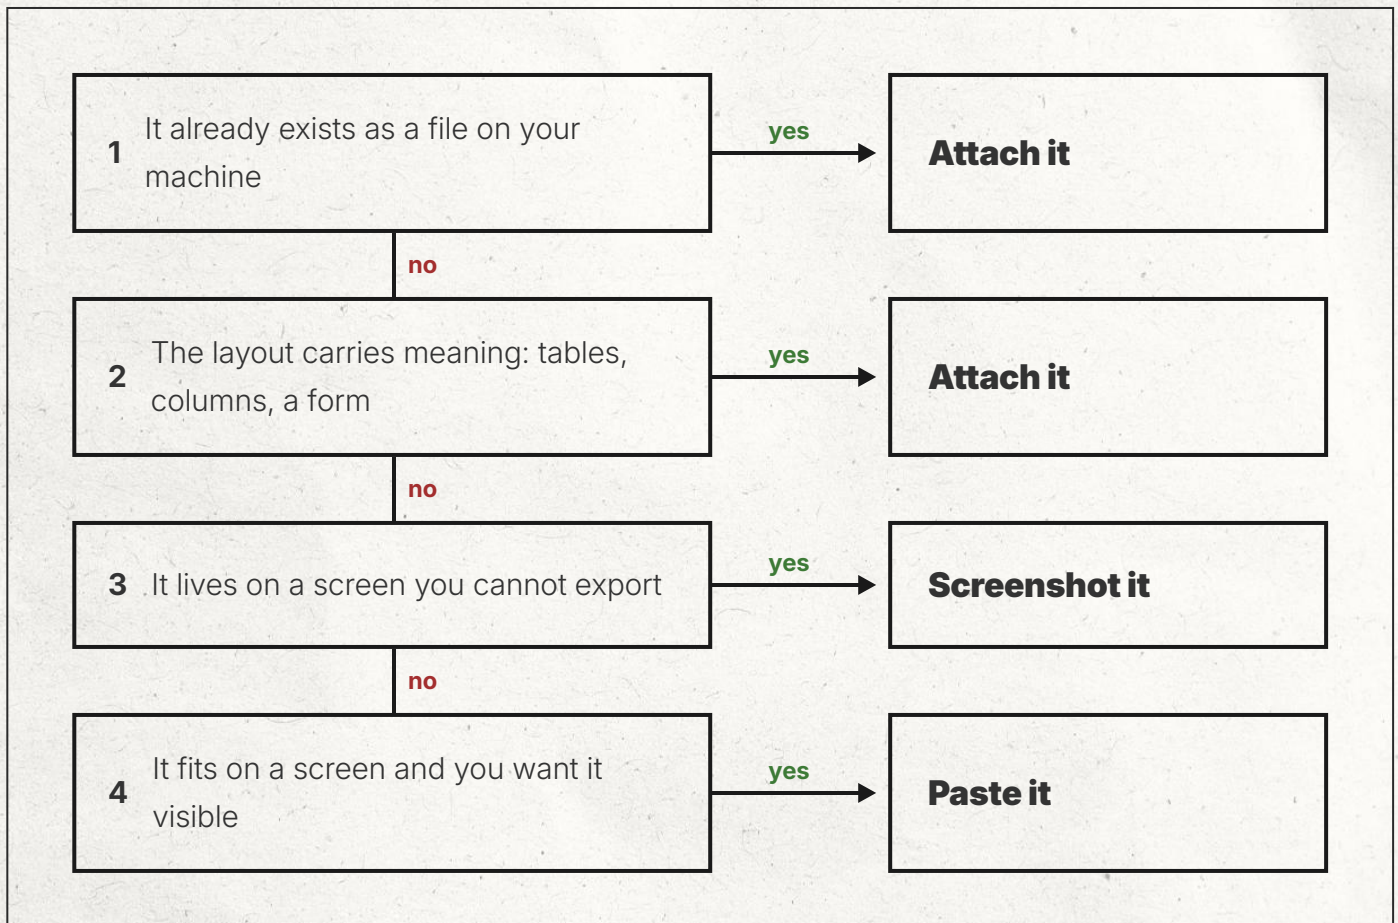


Fig 7.2 Read down until a line is true of your material. Most people default to whichever they did last time, which is how a 40-page report ends up pasted into a message box.

7.2

Paste. Short and visible in the thread. An email, a paragraph, a job advert, three bullet points from a call.

Attach. Long, structured, or already a file. Nothing is lost in transit and the original stays available.

The paperclip. All four put it at the left edge of the message box. Drag and drop onto the box works too.

Pasting is what you already do everywhere else: copy the text, click in the message box, paste, then type your question above or below it. It works fine, and it has one real advantage, which is that both you and the assistant can see exactly what is being discussed, right there in the thread.

Attaching means clicking the paperclip at the left of the message box, or dragging the file onto the box, and then typing your question as normal. Use it when the material is long, when the layout matters, or when the thing is already a file and turning it into text would lose something.

Watch what happens when you paste something very long into ChatGPT: past roughly ten pages of text it will turn your paste into an attachment on its own, and show it as a small file above your message. That is deliberate behaviour, not an error, and it has applied to free accounts as well as paid since mid 2026. It is also a decent hint about where the line sits.

Pasting out of a PDF or a formatted document drags stray line breaks, page numbers and running headers along with it. Usually harmless. If the reply seems confused about the structure of what you sent, that is your cue: attach the original instead of arguing with it.

Do not overthink which one you picked. Both routes end in the same place, with your material in front of the assistant instead of behind it. Try the next email thread you were about to describe in your own words: paste the thread instead, and notice how much less explaining you have to do.

7.3 Documents, sheets and screenshots go in

Four families of material cover almost everything a working professional needs to hand over. They are not read equally well, and the differences are worth knowing before you rely on one.




Fig 7.3 Ranked by how reliably the content survives the trip, most dependable on the left. This is an ordering, not a measurement: the bar heights carry no number.

7.3

ChatGPT Documents, spreadsheets and images. Free accounts get a small number of uploads a day, paid accounts get a much larger allowance per 3 hours.

Claude Up to 20 files in one chat. PDFs to 1,000 pages, but it only reads charts and layout in PDFs of 100 pages or fewer. Excel files need file creation switched on in settings.

Gemini Up to 10 files per message, 100MB each. Video and audio go in too, with free accounts capped at short clips and paid accounts at much longer ones.

Copilot Up to 20 files per conversation, 50MB each. Word, Excel, PowerPoint, PDF, images, and plain text or CSV.

Four families cover almost everything you will actually want to hand over, and they are not equally reliable, so it is worth knowing where each one sits.

Documents are the easy case. A Word file, a plain text file, or a PDF that was generated from a document all carry a real text layer, and every assistant reads that layer directly and accurately. If you have a choice of formats, this is the one to send.

Slides go in as PowerPoint files or as exported PDFs. The words come through cleanly. What comes through less well is everything the deck communicated visually: which box pointed at which, what the shading meant, what the tiny footnote under the chart said. Send the underlying document too if there is one.

Spreadsheets are the family with a catch. All four accept Excel and CSV, and all four can see your rows and columns and describe what is in them. Whether they genuinely calculate over those numbers, or simply read them as text and reason about them, depends on whether the assistant is running actual code underneath. It sometimes is. Chapter 19 takes this apart properly, including exactly where it will get arithmetic wrong. For now: trust it to find things in a spreadsheet, and check anything it adds up.

Images are the family that surprises people. Every one of the four reads pictures, which means a screenshot is a legitimate way to get material in. A dashboard you cannot export, an error message, a page of a printed report, a photograph of a whiteboard after a meeting. When the information you need is stuck inside a system that will not give it to you, a screenshot is the way out.

Go and attach one document today, using the file you were told to bring. Pick something you already know well, so you can judge the answer instead of trusting it.

7.4 It will not tell you what it did not read

The failure mode with attachments is not rejection. It is a confident answer built on the part of your file that happened to get read.

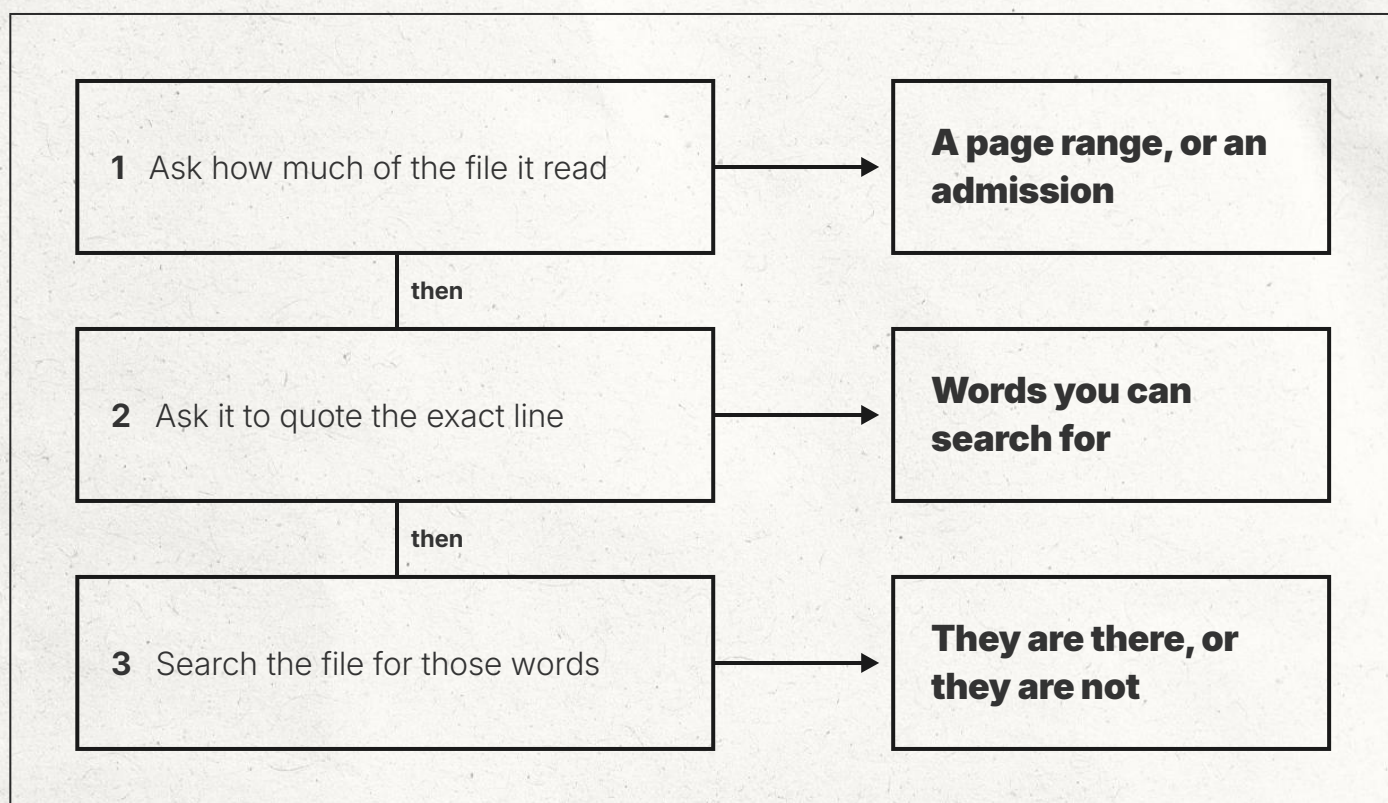


Fig 7.4 A 30-second check, worth running on anything that matters. Step 2 does the real work: a quote is checkable against the file, and a summary never is.

7.4

Long files. Past a certain size the assistant searches your document for relevant passages rather than reading all of it. It answers either way.

Scans. Clean printed pages read well. Faded, skewed, handwritten or densely tabulated ones read badly, and confidently.

Ask for quotes. The single cheapest habit in this chapter. Quoted text can be found in the file or it cannot.

Now the part nobody puts in the marketing. When you attach something very long, the assistant frequently does not read all of it. It indexes the file, searches it for passages that look relevant to your question, and answers from what it found.

That is a sensible design and it works well most of the time. Its cost is the quiet failure: your file uploads without complaint, the answer arrives sounding entirely normal, and nothing in it flags that a third of the document was never opened. Ask a 200-page report a question about page 180 and you may get a fluent answer drawn from page 12.

Scans behave the same way. A PDF that came out of a scanner or a photograph of a page has no text layer, so the assistant reads the picture instead. On clean printed text that works surprisingly well now. On a faded fax, a skewed photograph, handwriting in a margin, or a dense financial table it misreads characters, and it does not mark its own uncertainty. A wrong digit arrives looking exactly like a right one.

The fix costs half a minute and it is the same fix every time: make it quote. Ask which pages it used and ask it to give you the exact sentence its answer rests on, then search your file for that sentence. If the words are there, the answer has a foundation. If they are not, you have learned something far more useful than the answer.

That habit is the front door to Chapter 8, which is entirely about catching it when it is wrong, and it works better here than anywhere else in the book, because for once the source of truth is sitting in the chat with you. Reading long documents properly, extracting and summarising the way an analyst would, belongs to Chapter 18.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 8.

Do these before Chapter 8

3 exercises. About 15 minutes

Run the same request twice. Pick a document you know well and ask a question about it with nothing attached, then attach it and ask the identical question. Note what each answer gave you.

ANSWER WITH NOTHING ATTACHED

ANSWER WITH THE FILE ATTACHED

Take a screenshot of something you cannot easily copy out: a dashboard, a form, an error message, a page of a printed report. Send it and ask what it says. Write down the one thing it got wrong or missed.

Run the quote check on the answer from exercise 1. Ask it to give you the exact sentence its answer rests on, then search your file for that sentence. Write down whether you found it.

Chapter 7 on one page

THE CORE MOVE

If the answer depends on a document, the document goes in the chat. Everything else in this chapter is detail

PASTE OR ATTACH

Paste what fits on a screen. Attach what is already a file, or where layout carries meaning. Screenshot what you cannot export

ROUGH LIMITS

ChatGPT: few uploads a day on free, many on paid · Claude: 20 files a chat, charts read only in PDFs under 100 pages · Gemini: 10 files a message, 100MB each · Copilot: 20 files a conversation, 50MB each

THE QUIET FAILURE

A long file gets searched, not read, and nothing tells you which parts were skipped. Scanned and handwritten pages are misread confidently

THE 30-SECOND CHECK

Ask how much it read, ask it to quote the exact line, then search your file for that line

BRING TO CHAPTER 8

The document you attached and the two answers you compared. Chapter 8 uses exactly that pairing, because a claim you can check against a source in front of you is the easiest kind of error to learn to spot.

BOTTOM LINE

You can put your own documents, spreadsheets and screenshots in front of your assistant, and you know how to tell what it actually read.

NEXT CHAPTER

4 Ways AI Gets Things Wrong

Chapter 8 turns the quote check you just ran into a general habit. Why an assistant invents things in the first place, the tells that give it away before you have verified anything, and a verification routine short enough that you will actually keep doing it.

BRING WITH YOU


- ☐ The document you attached in this chapter, and the answer it produced
- ☐ One claim from that answer you have not yet checked

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

