

The AI Bootcamp

9

What You Can and Can't Put Into AI at Work

In this chapter you'll learn

The account that decides everything, 4 words for what never goes in, and who owns what comes out

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 A LinkedIn carousel here, a YouTube tutorial there, a \$2,000 cohort course that goes out of date before the refund window closes.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 your organisation already provides and what its rules actually say, you will have looked at your own data settings instead of assuming them, you will carry a 4-word test for what never goes in the box, and you will be able to say plainly what you do and do not own when you use the output.

Chapter 7 taught you to give it your own material, then stopped deliberately short of saying which material. This is where that gets answered. Almost nobody who gets into trouble here was careless. They were doing company work in a personal account, because nobody told them there was another option and nobody told them the rules.

CONTENTS	9.1	It depends which account you use
	9.2	What happens after you press send
	9.3	Four words for what never goes in the box
	9.4	You own the output, not a copyright

TOPICS REQUIRED

Chapter 8. You are carrying an assistant you already use on real work, your own documents from Chapter 7, and the verification habit from Chapter 8. What this chapter adds is the account all of that should be running on.

AFTER THIS CHAPTER

Chapter 10 assumes you can decide, without stopping to think, which pieces of your own writing you are allowed to paste in as samples. That decision is the whole method it teaches.

9.1 It depends which account you use

The same screen, the same words typed into it, and two completely different sets of consequences depending on whose contract you are sitting under. Almost everything else in this chapter follows from that one fact.

Your personal account	A company seat
A contract between you and a technology company. Your employer is not a party to it. TRAITS · You accepted the terms, on your own · Training and retention are settings you must find · Nothing your company negotiated applies PROS + Instant, free, already open + Fine for anything you would say in public CONS – Company material goes to a party you have no agreement with – No processing agreement, no indemnity, no admin oversight	The same product under a business contract your organisation signed and can enforce. TRAITS · Not used to train the vendor's models, by default · Covered by a data processing agreement · Conversations are logged and discoverable PROS + Removes most of the problem without you doing anything + Business tiers carry copyright indemnity CONS – Your employer can retrieve what you typed – Someone still has to tell you it exists

Fig 9.1 The same interface, sold twice. The difference is not a feature list, it is who signed the agreement and what it says about your content.

9.1

Personal account. A contract between you and a vendor. Your employer is not a party to it, so nothing it negotiated protects you.

Company seat. Same interface, different terms. All four business tiers say your content is not used to train their models.

The two questions. What are we licensed to use, and do we have a written policy. Both fit in one message to IT.

Here is the question that decides almost everything else in this chapter, and it is not the one people ask. The question is not whether AI is safe to use at work. It is which account you were signed in to when you typed.

A personal account is a contract between you and a technology company, agreed by you, on terms you clicked past. Your employer is not a party to it and never saw it. So nothing you paste into that box is covered by anything your company negotiated, because your company negotiated nothing. In plain terms, you have handed material to a third party you have no agreement with.

A company seat is a different product wearing the same interface. All four vendors sell business tiers, and on all four the contract says close to the opposite of the consumer one: your content is not used to train their models, by default, with no setting for you to hunt down. ChatGPT Business, Enterprise and Edu. Claude for Work. Google Workspace with Gemini. Microsoft 365 Copilot when you are signed in with your work account rather than a personal one. Same box, same skills, different terms behind it.

Which is why the most useful thing you can do this week is not to memorise a policy. It is to find out what your organisation already provides. A great many companies bought seats across 2024 and 2025 and never announced it clearly, so a meaningful number of people are pasting work into a free personal account in the same building as a paid licence nobody mentioned.

Ask two things while you are asking. What are we licensed to use, and what does our policy actually say. If the answer to the second is that no written policy exists, that is worth knowing too, and if you are senior enough to be reading this book you are probably senior enough to be the person who fixes it. In Europe there is a further nudge: since February 2025 the AI Act has required organisations deploying AI to make sure the people using it on their behalf have adequate AI literacy, with the enforcement machinery due from August 2026.

Send that message before you read the next section. The rest of this chapter gets a great deal shorter if the answer comes back "we already have Copilot, it is on your laptop."

9.2 What happens after you press send

Three separate things can happen to a message: it gets stored, it gets used to improve the product, and a person may read it. They are not the same thing, and they are not controlled by the same switch.

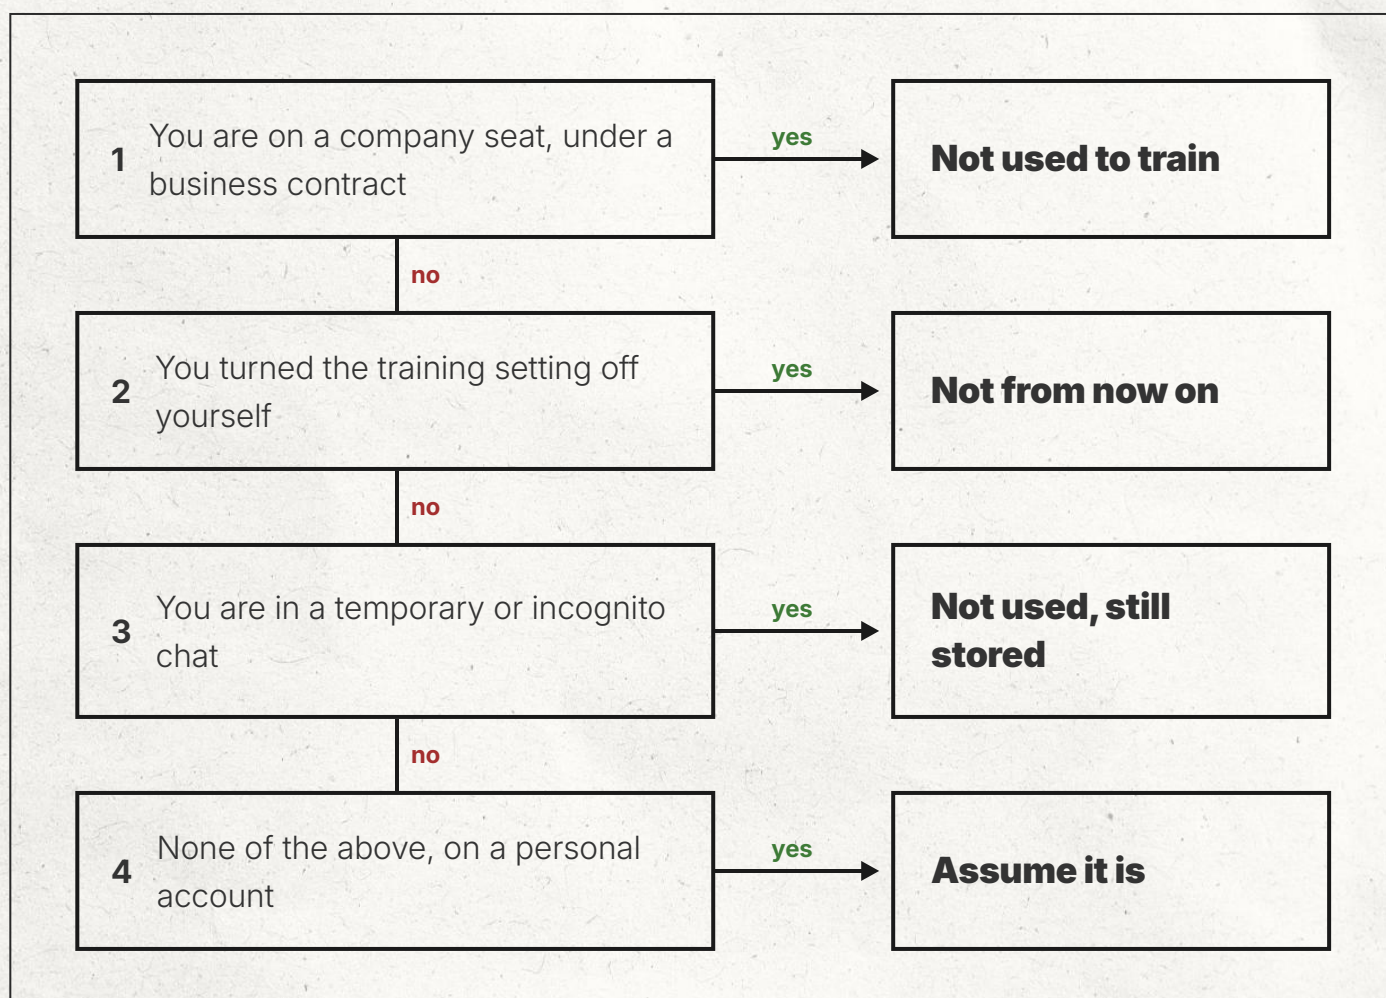


Fig 9.2 Read down until a line is true for you. The thing to notice is line 4: on a personal account the permissive answer is the default one, so saying nothing is a decision.

9.2

ChatGPT Settings, then Data Controls. The switch is called Improve the model for everyone, and on personal plans it starts on.

Claude Privacy settings. Allowing training extends retention to 5 years. Declining keeps it at 30 days.

Gemini Gemini Apps Activity. Kept 18 months by default, and a sample of chats is read by human reviewers.

Copilot Privacy, then training on conversation activity. Signed in with a work account it is a different product entirely.

The honest answer to "does it learn from what I type" is that it depends on which of those two accounts you are in, and the gap between them is not subtle.

On personal accounts the four vendors differ in the details and agree on the shape. ChatGPT keeps a switch called Improve the model for everyone in Settings, under Data Controls, and on personal Free, Plus and Pro accounts it starts on. Claude asked every consumer user to make an explicit choice in late 2025: allow it and conversations can be kept for 5 years, decline and retention stays at 30 days. Gemini keeps chats for 18 months unless you change it to 3, or 36, or off. Microsoft Copilot has its own toggle for training on conversation activity. You almost certainly answered one of these questions once, quickly, and have not looked since.

Google is the bluntest of the four about the consequence, and it is worth quoting because it sits on their own help page rather than in a critic's blog post: do not enter confidential information that you would not want a reviewer to see, or Google to use to improve its services. That is not a warning about hackers. It is a warning that a subset of conversations is read by people, deliberately, as part of how the product gets better.

Turning the setting off is worth doing and does less than people assume. It applies going forward, not backward. It does not stop the conversation being stored, and stored is its own category of risk: in 2025 a court order in a copyright case required OpenAI to preserve consumer chat logs it would otherwise have deleted. That order was lifted in October 2025, and the lesson outlives it. What a company promises to delete and what a court can later require it to keep are two different questions. Temporary and incognito chats sit in the same gap: they stay out of your history and out of training, and they are still retained for a period, typically around 30 days.

A company seat closes most of that, and opens one thing worth saying plainly because no vendor will volunteer it. On business tiers your conversations are logged, exportable and discoverable. ChatGPT Enterprise ships a compliance interface built for exactly that. Microsoft 365 Copilot activity lands inside your own tenant's audit and eDiscovery. Claude's Team and Enterprise plans expose the same, incognito chats included. The seat protects the company from the vendor. It does not make you private from your employer, and it was never sold as doing so.

Go and look at your own setting now, before the next page. It takes under a minute, and you cannot make a judgement about line 4 of that figure without knowing which line you are actually standing on.

9.3 Four words for what never goes in the box

You will never read a confidentiality policy in the 4 seconds between copying something and pasting it. So you need a test short enough to run from memory, with your hand still on the mouse.




Fig 9.3 People, promises, passwords, plans. The bar heights are visual rhythm only, not a ranking: any one of the four is enough to stop you.

9.3

The four Ps. People, promises, passwords, plans. If what you are about to paste contains one, it does not go in.

The 4-second test. You will not read the policy at the moment you paste. You will run 4 words, or nothing at all.

A seat changes the answer. It does not delete the list. It moves some categories from never to check what our policy allows.

Every organisation's confidentiality policy is longer than anyone reads, and none of it is available in the moment that actually matters. So the rule has to be four words long.

People, promises, passwords, plans. If what you are about to paste contains any of those, it does not go into a personal account at all, and on a company seat it goes in only where your own policy says that category is allowed.

People means anything that identifies a named individual: a customer list, a CV, a patient record, a performance note, a call transcript with real names in it. In Europe this is the sharpest of the four, because putting a client's personal data into an account your employer has no processing agreement with stops being a matter of taste and becomes your employer's exposure under GDPR. North American privacy law is more fragmented and the direction of travel is the same.

Promises means anything confidential because you agreed it would be: material under an NDA, a client's documents, a partner's roadmap, code you do not own. You cannot promise confidentiality to one party and then decide on your own to share it with a second. Passwords means credentials, keys and tokens, which ought to be obvious and still get pasted regularly, because they happened to be sitting in the middle of the log file someone asked you to explain. Plans means anything true but not yet public: unreleased results, a deal in progress, a restructuring, pricing that has not shipped. For a listed company that last category has another name and a regulator attached to it.

The well-known version of this is Samsung in 2023. Engineers pasted source code and a meeting transcript into ChatGPT to get help with their actual jobs, and within weeks the company had banned the tools on its own devices. Nobody in that story was acting badly, and that is the part to take from it.

In the largest global study on the subject, 48,000 people across 47 countries, almost half of employees admitted using AI in ways that break their own company's rules. More rules were not the missing ingredient. A rule you can apply in 4 seconds was. Write your own four on the exercises page, in your company's nouns.

9.4 You own the output, not a copyright

Two different questions get answered with the same word, which is how careful people end up surprised. What the vendor gives you and what the law protects are separate things.

What the terms give you

Every major provider assigns you whatever rights it holds in the output it produced for you.

TRAITS

- You own the output as against the vendor
- No fee, no credit, no licence to renew
- Business tiers add copyright indemnity

PROS

- + You can use it commercially
- + Nothing to negotiate, it is already in the terms

CONS

- A contract promise, not a property right
- Consumer plans generally carry no indemnity

What copyright gives you

Protection for the part a human actually authored, and only that part.

TRAITS

- Prompting alone is not authorship
- Your selection, arrangement and edits can be
- Settled in the US, tested up to March 2026

PROS

- + The more you rework it, the more is yours
- + Irrelevant for internal work, so do not overthink it

CONS

- An unedited draft may be protectable by nobody
- A similar answer can be produced for someone else

Fig 9.4 Two questions, one word. The contract settles whether you may use it. Copyright settles whether anyone else can be stopped from copying it.

9.4

The terms. All four assign you the output. You are not borrowing the words and you owe nobody a credit.

The law. Prompting is not authorship. What you select, arrange and change is what can be protected.

Your employer. Work produced in the course of your job belongs to them, exactly as it did before any of this.

Owning the output and being able to copyright the output are two different questions, and they get answered with the same word.

The contract half is settled and generous. All four providers hand you whatever rights they have in what the model produced for you. OpenAI's terms assign it to you outright; Anthropic's say the customer owns its outputs. You owe no fee, no credit and no attribution, and you are not borrowing the words on licence.

The copyright half is narrower, and it has been tested rather than assumed. The US Copyright Office concluded in January 2025 that prompting alone, however long and careful the prompt, does not make you the author of what comes back. What can be protected is the human contribution around it: your selection, your arrangement, your edits. In March 2026 the Supreme Court declined to take up the case that would have challenged the human authorship requirement, so this is now the settled position rather than a live argument.

In practice that means the more you cut, rewrite and shape, the more of the result is genuinely yours. A draft you accepted unchanged is a draft nobody owns. That matters a great deal for anything you intend to license, defend or sell, and not at all for an internal memo. Know which one you are producing before you worry about it.

Two more things worth knowing. An output is not exclusive: a similar question from someone else can produce a very similar answer, so treat what comes back as a draft rather than as an asset with a fence around it. And if you produced it in the course of your job, your employer almost certainly owns it whatever the vendor's terms say, exactly as with everything else you write at work. AI does not change that chain of ownership, it only shortens it.

The one real asymmetry is worth carrying into your next budget conversation: the major vendors offer copyright indemnity to business customers and not to consumers. That is one more line in the argument for using the seat your company pays for. Whether it covers your particular situation is a question for whoever signed the agreement, which is the honest answer to a lot of this chapter. Read your own contract, or go and ask the person who has.

READING ENDS HERE

Everything before was reading. Everything after is doing.

3 exercises, About 15 minutes and one real decision that opens Chapter 10.

Do these before Chapter 10

3 exercises. About 15 minutes

Send one message to IT, or to whoever owns your Microsoft or Google relationship, asking two things: what AI tools are we licensed to use, and do we have a written policy. Write the answer here when it comes back.

Open your assistant's settings and find these two things. Write what yours are actually set to, not what you assume they are set to.

TRAINING SETTING

HOW LONG CHATS ARE KEPT

Write your own four Ps in your company's own nouns. Name the actual documents, systems and categories of person you handle in a normal week.

Chapter 9 on one page

THE ACCOUNT DECIDES

A personal account is a contract you agreed alone. A company seat carries terms saying your content is not used for training. Find out which one you are working in

THE FOUR PS

People, promises, passwords, plans. If what you are pasting contains one of them, it does not go into a personal account

WHERE THE SETTING LIVES

ChatGPT: Settings, Data Controls · Claude: Privacy settings · Gemini: Gemini Apps Activity · Copilot: Privacy, training on conversation activity

OFF IS NOT DELETED

Turning training off applies going forward only, and does not stop the chat being stored. Temporary and incognito chats are still retained, typically around 30 days

THE OUTPUT

The terms give it to you. Copyright protects only what a human added. Business tiers add indemnity, consumer plans do not

BRING TO CHAPTER 10

The answer from IT about what your company provides, and the setting you found in your own data controls. Chapter 10 has you feeding it your own writing, so it matters that you know where that writing goes.

BOTTOM LINE

You know which account to work in, what never goes in it, and what you actually own when something comes back out.

NEXT CHAPTER

Make AI Sound Like You

Chapter 10 fixes the thing everybody notices and nobody can name: the flat, over-polished register that gives AI writing away. You give it samples of your own work, and it stops sounding like a press release.

BRING WITH YOU


- ☐ The account and the settings you sorted out in this chapter
- ☐ 3 pieces of your own writing you would happily show a colleague

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



We turn attention into pipeline.

visit thecentral.ai

➤