AI Accelerator

THE GUIDE From inside real programmes

AI Adoption That Lasts

25 lessons learned the hard way

For the person leading AI adoption: one of the hardest jobs in any organisation right now.



Running the AI Primer

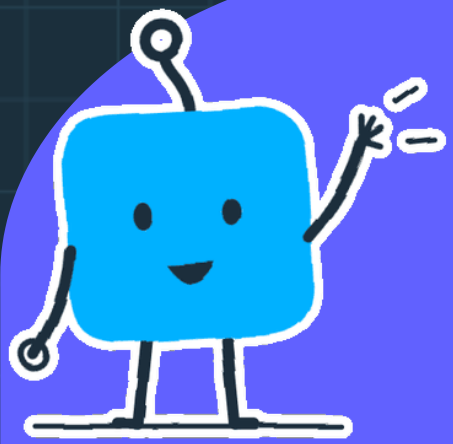
INSIDE, IN FOUR PARTS

01 What turns the dial

02 What to set up

03 Don't do this

04 Don't worry about this



Written by the practitioners at the AI Adoption Lab, Greenwich

Before you start

This guide is for the person who has been handed AI adoption. You might run L&D, lead transformation, carry the title Head of AI, or own the rollout of your AI tools. Either way, the licences have been bought and the question is now yours: how do you get a whole organisation using this well, safely, and in a way that lasts?



A live session at the AI Adoption Lab, Greenwich

START HERE

Start from the right mindset: this is behaviour change, not a technology rollout. Nothing changes until people do their own work differently, week after week, and that is hard. It's harder still because the tools keep changing while people are learning them. Your executive team and your programme team both need to start from there. Run it like a system going live and you'll get a launch, a spike in logins and a quiet drift back to the old way.

Everything here comes from our own enablement work: programmes inside construction, insurance, consumer brands and professional services, plus everything we've learned running our own business on AI. We haven't named the organisations. We have kept the lessons, including the ones where we got it wrong.

It isn't a methodology. We've built a few of those and watched them go out of date (lesson 20 explains). It's a set of lessons in four parts:

1

What turns the dial

The things that made the biggest difference.

2

What to set up

The activities and structures worth putting in place early.

3

Don't do this

Mistakes we've seen, three of them our own.

4

Don't worry about this

Things that take up time and attention without moving much.

After that: how to run adoption the new way, using the tools themselves; where this is all heading; the five common ways organisations go about adoption and where each falls short; how we work and why; and where to start.

Every lesson here is one we're happy to justify. If you want the detail behind any of them, ask us.

The ten that matter most

If you only have five minutes, or need something to forward to your sponsor:

TURNS THE DIAL | SET UP | DON'T DO

01 **Leaders go first, and they have to actually use it.**
Your people can tell when a leader isn't all in.

02 **Teach how AI works, not where the buttons are.**
The buttons move. The mental model doesn't.

03 **Start with what's in it for each person.**
Business value follows personal gain, not the other way round.

04 **Short, frequent and in a fixed slot.**
A fixed weekly slot will outlast a training day.

05 **Put practitioners, not trainers, in front of people.**
The questions unlock the value, and practitioners have the answers.

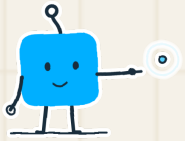
06 **Write a short manifesto before anything else.**
AI isn't cheating, you own the work, and here are the edges.

07 **Start governance light, then build it out.**
Tier the use case, not the tool, and add checks as AI starts to act on its own.

08 **Train and equip your champions, don't just appoint them.**
Training, kit and protected time. Judge them at month three.

09 **Measure adoption with five metrics, not logins.**
Four show the habit forming, and one business result is agreed with finance.

10 **Don't build your own AI learning platform.**
Home-built ones end up with too little content, and the wrong content. Subscribe to the signal instead.



PART 1

What turns the dial

01 Leaders go first, and they have to actually use it

02 Make the first session mandatory, and invite people by name

03 Start with the tools people already have

04 Teach how AI works, not where the buttons are

05 Start with what's in it for each person

06 Short, frequent and in a fixed slot

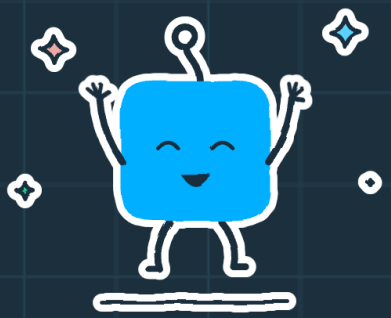
07 Train on people's real work, function by function

08 Build skills, not one-off prompts or agents

09 Put practitioners, not trainers, in front of people



Leadership first, then everyone else



01 Leaders go first, and they have to actually use it



Every programme we've seen move quickly started with the leadership team. Most organisations know this and book a leadership briefing. The bit they miss is that going first isn't the same as using it.

WHAT WE'VE SEEN

In one programme, some leaders were visibly using AI in their own work. Others asked their teams why AI couldn't do something, without ever having tried it themselves. Staff noticed the difference within weeks. A leader who is privately anti-AI, or simply doesn't get it, will undermine a programme without saying a word: people read what leaders do, not what the launch email says.

WHAT TO DO

- Run a session for leaders before anyone else, and make it hands-on. Show live, real-time demonstrations of what's possible now, using the leaders' own kind of work. When people see how easy something is once you understand the model, the conversation changes.
- Ask each leader to do one real piece of work with AI before the wider launch, and to talk about it.
- Have leaders give explicit permission to experiment. Curiosity is the most useful trait a team can have with AI, because these systems keep doing things nobody planned for. People need to hear that trying things is the job, not a distraction from it.

WHERE WE COME IN

Our leadership sessions are run by people who use these tools all day and can show what's possible on the spot, then answer the hard questions that follow.



Leaders first: show what's possible, then take the hard questions

02 Make the first session mandatory, and invite people by name

Voluntary training fills with the people who were going to use AI anyway. Everyone else waits to see if it matters.

WHAT WE'VE SEEN

At one client, attendance at the foundation session stalled. When the session was made mandatory and invitations went to people individually, by name, the cohort filled within days.

WHAT TO DO

- Make the foundation session mandatory for everyone who needs the basics, with a faster track for experienced users (lesson 22). Keep it short (see lesson 6).
- Send named invitations from someone people report to, not a generic all-staff email.
- Make everything after that voluntary. Once people have the basics, the pull takes over.

03 Start with the tools people already have

It's tempting to design the programme around the most powerful licence. Most of your people won't have it.

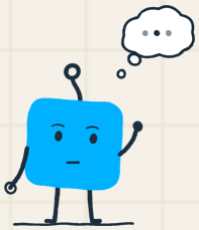
WHAT WE'VE SEEN

In a Copilot rollout, far more people picked up the basic chat tool in the first month than the plan had assumed, several times more. The programme was designed so that everything started with what works on any licence, and the paid licence was there for the people who needed it.

WHAT TO DO

- Design the foundation for the licence most people have.
- Check every exercise and demo against what your audience can actually access. Content built for features people don't have is wasted, and it frustrates them.
- Treat paid licences as a step people earn through use, not a starting point.

04 Teach how AI works, not where the buttons are



Tool training feels productive and goes out of date fast. The menu moves, the feature gets renamed, a new model behaves differently, and people who learned the clicks are lost again.

WHAT WE'VE SEEN

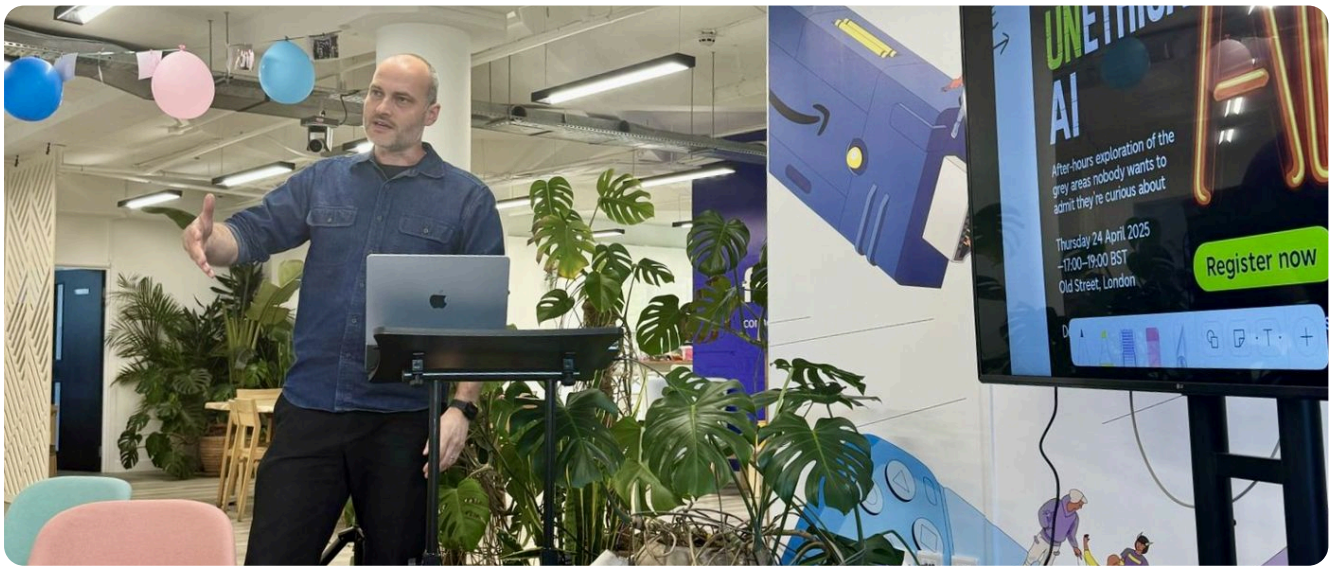
The people who get furthest aren't the ones who know a particular tool best. They're the ones who understand what an AI model is actually doing: why it gets things wrong, what it needs from you, why the same request gives a different answer tomorrow. With that understanding they can pick up any tool, and they know why the button moved. Mindset beats tool knowledge every time.

WHAT TO DO

- Start the learning with how AI works, once people know what's allowed (lesson 10): what a large language model is, what it's good and bad at, why context matters, where the risks are.
- Teach techniques that work in any tool, then show them in your tool.
- Keep a short, plain explanation of the risks in front of people from day one.

WHERE WE COME IN

Our AI Primer starts here, and our platform keeps the mental model current as the tools change underneath it.



Teach the model, live: what AI is good at, and where it trips

05 Start with what's in it for each person

Organisations want productivity. People want their Tuesday afternoon back. If you only sell the first, you won't get the second.

WHAT WE'VE SEEN

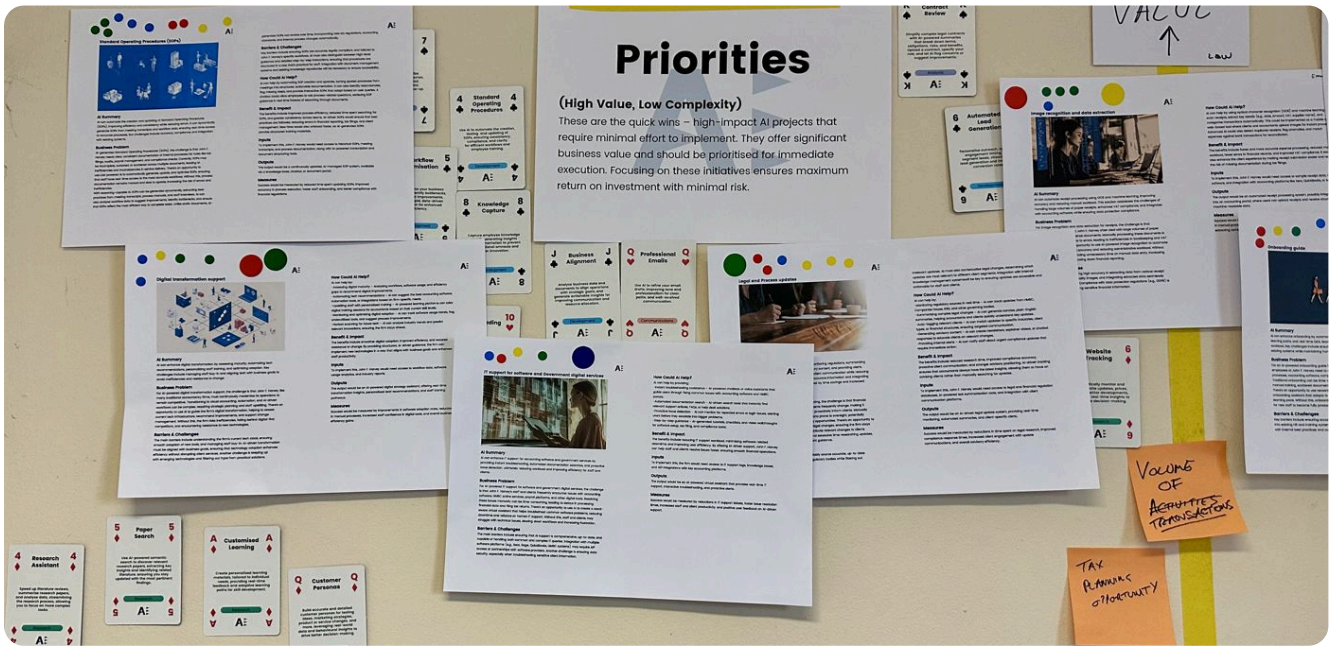
The moment adoption takes off for an individual is when they find one thing that saves them real time in their own job. After that they go looking for more. Programmes that lead with organisational ROI get polite attention. Programmes that help each person find their own use case get people coming back.

WHAT TO DO

- Build use-case discovery into the first weeks: help each person find two or three uses in their own work, not generic examples.
- Ask "what do you do every week that you'd happily never do again?" It's a better starting point than any list of use cases.
- Let the business case build up from the bottom. Personal wins add up to team wins.

WHERE WE COME IN

We run use-case discovery as a structured activity, and we have a library of real use cases by role to start the conversation.



Use-case discovery: people's own work, sorted by value

06 Short, frequent and in a fixed slot



Long training days feel thorough, and a launch event gets attention for a day. Most of both is forgotten by the following week. Small pieces, in the same slot every week, build a habit.

WHAT WE'VE SEEN

A leadership session briefed at four hours met a chief executive who expected two. We cut it to three, with a break every hour and the most useful material first, and we've designed in small pieces ever since. At one client, a voluntary webinar in the same slot every week drew around a quarter of the organisation, week after week. We've also run single broadcasts to thousands of people verify: ~4,500, Mike's figure and regular sessions of 100 to 200 verify. The format that works best is about 15 minutes of content, then 15 to 45 minutes of questions answered by a practitioner.

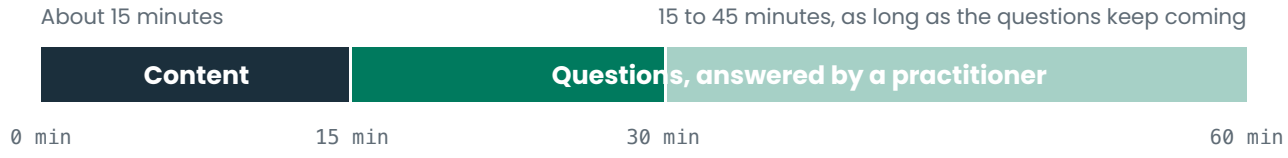
WHAT TO DO

- Break learning into units people can finish in a few minutes, and give them something to come back to. We use a deck of cards, each linking to a short lesson.
- Pick a weekly slot and keep it. Same day, same time.
- Keep the content to around 15 minutes and spend the rest on questions. The questions tell you what to teach next.
- Make AI a standing item on every team's agenda, and ask the chief executive to open each all-hands with an AI use case.
- Produce it properly: a studio set-up, with someone producing alongside the presenter.

WHERE WE COME IN

Our platform is built around short learning cards, and we run studio-produced webinars with practitioners taking questions live.

The shape of a good webinar



Illustrative. The questions are where most of the value is.

07 Train on people's real work, function by function

Generic training gives everyone the same examples. Nobody's job looks like the examples.

WHAT WE'VE SEEN

With one insurer's marketing team we ran two half-days. The first mapped their real work and the ideas worth trying. Between sessions we sent a written summary and a small piece of homework. In the second, everyone loaded and used an AI skill on the day, four small teams each built a skill for their own work, and another team tested each one.

WHAT TO DO

- [x] Run sessions by function or team, not just company-wide.
- [x] Split it in two: discover first, build second, with time to think in between.
- [x] Have people work on their own material, not exercises.
- [x] Swap-test whatever gets built. If another team can't use it, it isn't finished.

WHERE WE COME IN

Our functional workshops follow this pattern and leave teams with something working, not just notes.



Discover first: a team sorting its own work

08 Build skills, not one-off prompts or agents

A good prompt helps one person once. A skill (a reusable, shareable set of instructions an AI follows for a particular job) helps everyone who does that job, every time, and gets better as you improve it.

WHAT WE'VE SEEN

When champions at one client were encouraged to build agents, they built quite a few, and most of them had exactly one user: the person who built them. Meanwhile the teams that built shared skills got repeatable results, and the major vendors, Microsoft included, are now building skills into their own products. Whatever tool you're on, this is where it's going.

WHAT TO DO

- [x] Teach people to turn a task they repeat into a skill.
- [x] When the output is wrong, fix the skill, not just the output. Otherwise you fix the same thing again next week.
- [x] Keep look and task separate: one skill for your house style, another for the job.
- [x] Give every shared skill a named owner, so it improves in one place.

WHERE WE COME IN

We teach skill-building from the foundation level up, and our own business runs on them.

house-documents/SKILL.md — A SKILL

```
---
name: house-documents
description: Turns rough drafts into finished, on-brand documents.
---

# House documents
1. Start from the matching template.
2. Follow references/style-guide.md.
3. UK English. Never invent a fact.
```

09 Put practitioners, not trainers, in front of people



The best moment in any session is when someone asks the question that isn't on the slides. Whether that moment helps or frustrates depends on who's at the front.

WHAT WE'VE SEEN

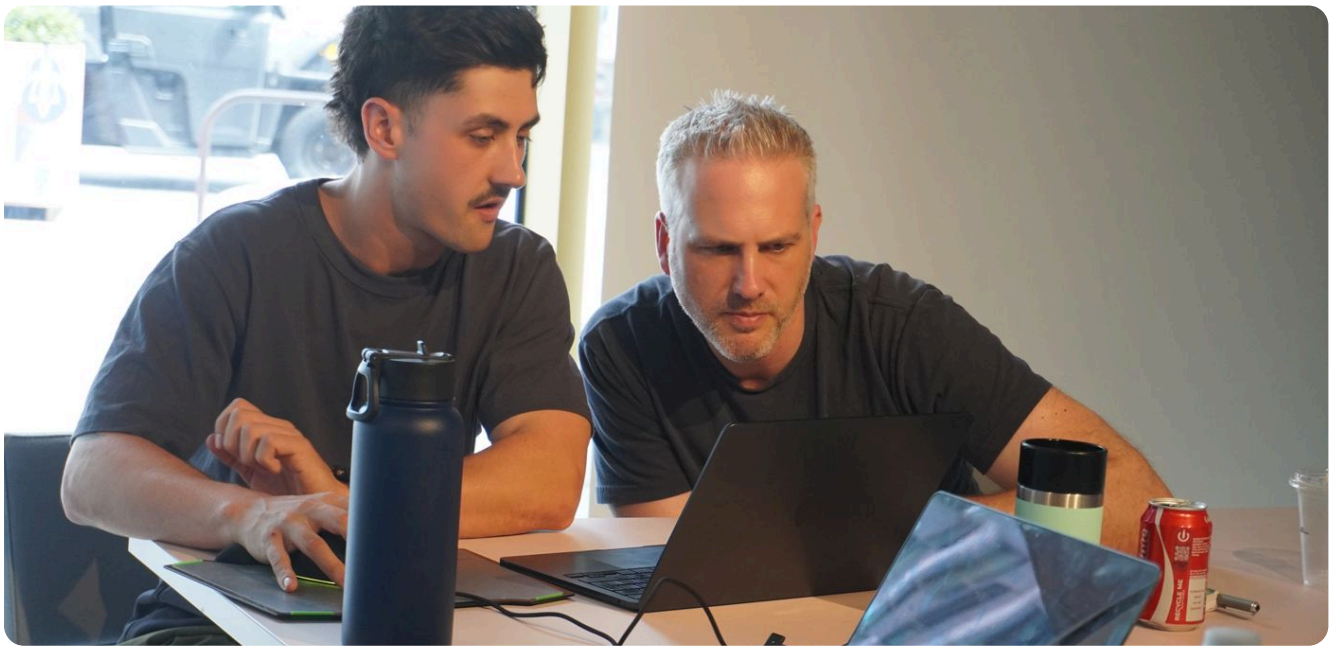
In one programme's sessions, more than a hundred questions came from dozens of different people, and most of them were specific: about their tools, their data, their job. Rooms are mixed: some people have never used AI, others have been using it at home for years. The advanced users rated generic sessions lowest. What all of them valued was someone who could answer from experience.

WHAT TO DO

- [x] Choose people who use AI daily in their own work to run sessions, not people who learned the material last month.
- [x] Leave long Q&A time. If your presenter can't fill it, that tells you something.
- [x] Capture every question. They become your FAQ, your next session and your measure of what's landing.

WHERE WE COME IN

Everyone who presents for us builds with these tools every day. That's why we can keep the Q&A open.



Practitioners answer the next question, one to one if needed

PART 2

What to set up

10 A short manifesto, before anything else

11 Start governance light and build it out, with the people who can say no

12 A baseline on day one, and the person who owns your usage data

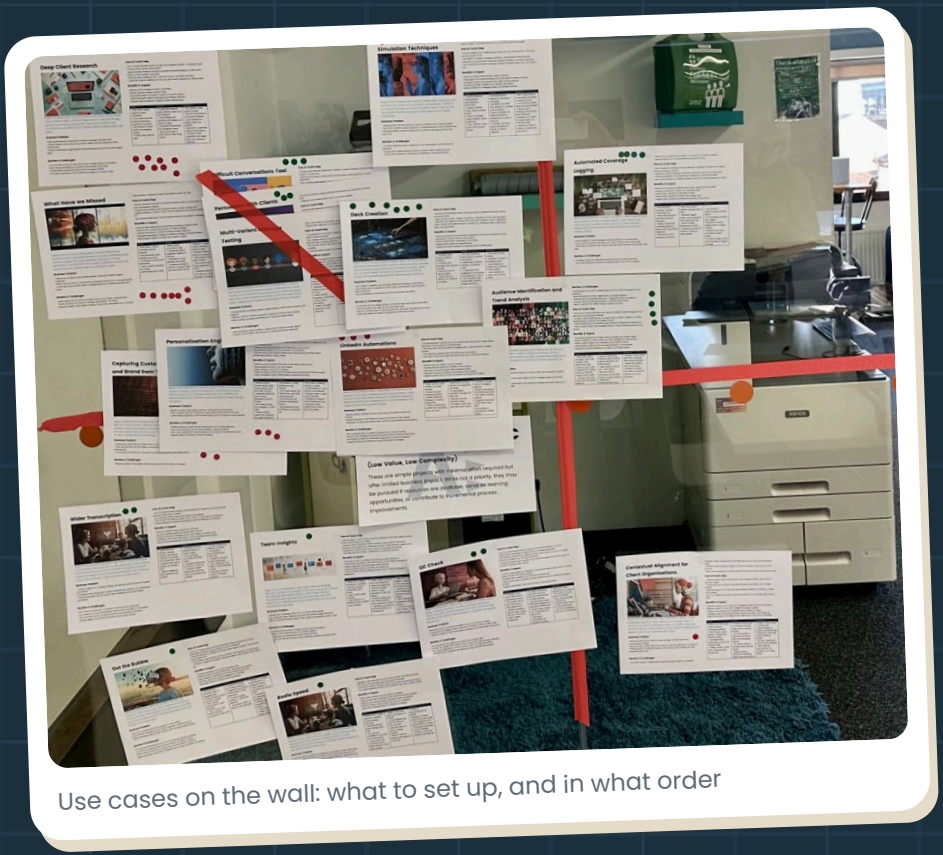
13 Have the answers ready before the questions arrive

14 Train and equip your champions, don't just appoint them

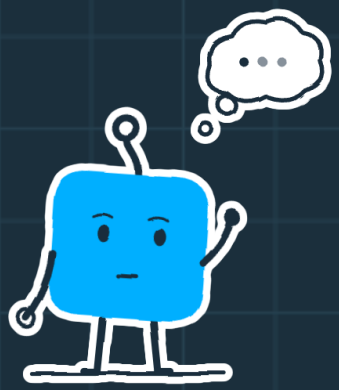
15 Somewhere to go that stays current, and a way to hear what people need

16 Measure adoption with five metrics, and set a target for each

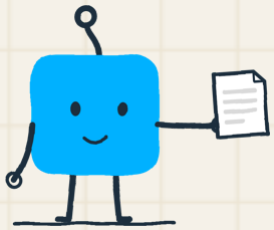
17 Plan with what you know, and review it every quarter



Use cases on the wall: what to set up, and in what order



10 A short manifesto, before anything else



Before people learn anything, they need to know what's allowed. Without it, the cautious don't start and the keen go off on their own.

WHAT WE'VE SEEN

The question people most often ask in their first session isn't how to do something. It's "am I allowed to?" Organisations that answered that up front moved faster than those still drafting a policy.

WHAT TO DO

Write one page that says:

- [x] Using AI isn't cheating. It's expected.
- [x] You own the work. Whatever comes out of AI, you're accountable for it, so check it.
- [x] Here are the edges: what data never goes into which tools, and where to ask if you're unsure.
- [x] Here's where to go to learn more and to share what you've found.

Make sure everyone has the same understanding of the risks, so they know where the edges are, and build the security rules into the training itself, not a separate document nobody reads.

WHERE WE COME IN

We help leadership teams write this in a single session, and we build it into every piece of training that follows.

Our AI manifesto

ILLUSTRATIVE · ONE PAGE

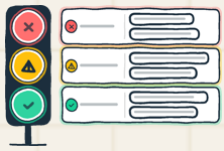
01
AI isn't cheating.
It's expected.

02
You own the work.
Whatever AI produces, you check it and answer for it.

03
Here are the edges.
What never goes into which tool, and who to ask.

04
Here's where to learn.
And where to share what you find.

11 Start governance light and build it out, with the people who can say no



Governance programmes that try to approve every tool in advance either never finish or block everything. The risk isn't in the tool. It's in how it's used. And every function that can stop the programme (legal, information security, data protection, HR, procurement) will stop it if they first meet it when something goes wrong.

WHAT WE'VE SEEN

At one client the most useful idea was a short pre-use check: a handful of questions about accuracy, accountability, data and unintended consequences, which sorts each use case into go ahead, check first, or don't. Where organisations set out to design a full governance framework before starting, it absorbed months. The steering group that worked best had technology, information security, people, legal, data protection and procurement in it from the start, and met regularly. Where the group lapsed, the programme drifted.

WHAT TO DO

- [x] Start with the manifesto (lesson 10) and a short pre-use checklist.
- [x] Tier by use case: low-risk work moves fast, scrutiny goes where the stakes are real.
- [x] Put the "no" functions on the steering group from the start. Meet monthly, and give the group one job: clearing blockers fast.
- [x] Decide early who answers the awkward questions: whether to disclose AI use in client work, how long prompts and outputs are kept, how meeting transcripts can be used.
- [x] Grow the governance as use grows. It doesn't need to be finished on day one. Agents that send, buy or change things will need more checks than a chat window.

WHERE WE COME IN

We bring the questions organisations usually only discover later, and we weave the answers into training so people know where the green lane is.

Tier the use case, not the tool

Go ahead
Low-risk work moves fast

Check first
Run the pre-use questions

Don't
Out of bounds, and people know why

12 A baseline on day one, and the person who owns your usage data

If you don't know where people started, you can't show how far they've come. You'll also want usage numbers, and they're harder to get than anyone expects. Most organisations remember both three months in.

WHAT WE'VE SEEN

A short form at the start and end of each session is enough. Across our own live sessions, participants report more than twice as many people feeling confident using AI after working with us.* On a Microsoft rollout, the usage data was weekly, organisation-wide and partly exported by hand verify: current Microsoft reporting position. One call with the person who looked after the admin centre agreed a weekly export and a six-month baseline, after months without usable numbers.

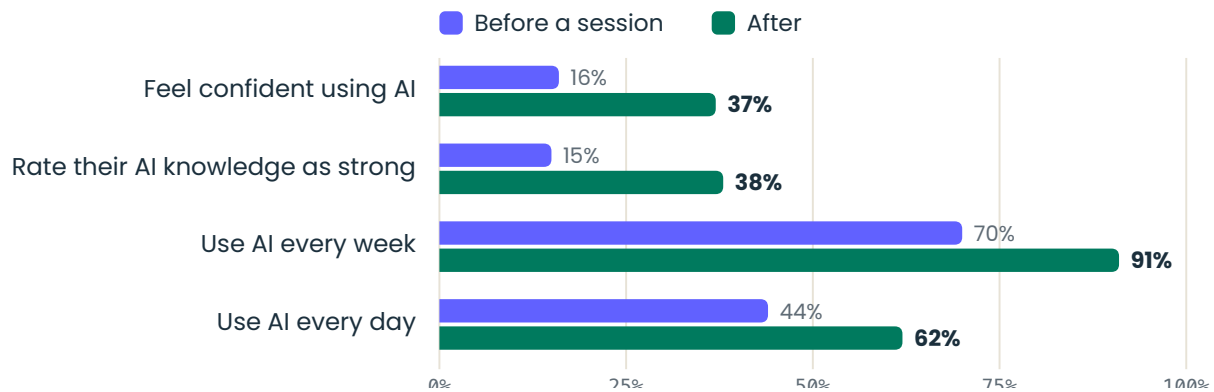
WHAT TO DO

- [x] Ask three things before people start: how confident they feel, how much they know, how often they use AI. Ask the same three afterwards, and keep it to two minutes.
- [x] Find whoever administers your AI tools and their reporting, and get them on side in week one.
- [x] Agree what data you can get, how often and in what format before you promise anyone a dashboard, and take a historical baseline before the programme starts.

WHERE WE COME IN

Our platform runs welcome and exit forms for you and shows programme leads the change.

What changes after a live session with us



Self-reported by participants on our welcome and exit forms, 504 and 120 responses, January 2025 to October 2026. Confidence and knowledge: 5 or more out of 7.

13 Have the answers ready before the questions arrive

Once people understand how AI works, the same questions come in every organisation: Is it safe? Does it learn from my data? Can I use it with clients? What happens to my transcripts? If every answer has to be worked out from scratch, people wait, or they guess.

WHAT WE'VE SEEN

The questions are remarkably consistent. We now issue a library of well over a hundred standard questions with model answers at the start of a programme, to be adapted to the organisation's own approach.

WHAT TO DO

- [x] Prepare answers to the common questions before launch.
- [x] Adapt them to your policies and tools. A generic answer to "can I use this with client data?" isn't good enough.
- [x] Publish them somewhere people can find, and add to them every week.

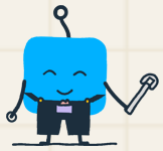
WHERE WE COME IN

We start every programme with our FAQ library and keep it current.



Questions asked live become your FAQ library

14 Train and equip your champions, don't just appoint them



Champions are one of the best ideas in adoption and one of the most often wasted. Most organisations appoint them, thank them, and leave them to it.

WHAT WE'VE SEEN

Champion roles are usually designed at around four hours a week. In practice we've seen one to two, because the day job wins. At one client, champions were asked to come up with agent ideas, and only a small group had the skills to. The real value started to show at around three months, not one.

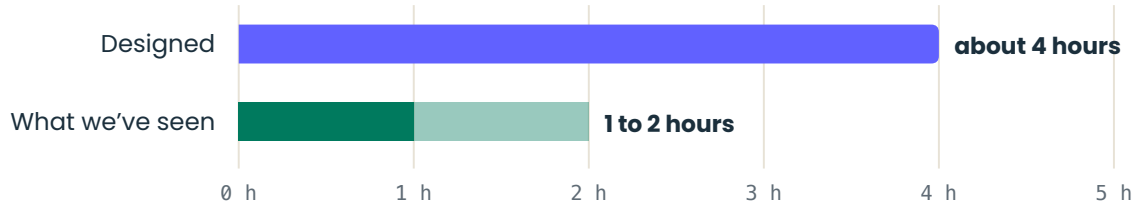
WHAT TO DO

- Give champions the kit, not just the title: ready-made lessons and exercises they can run, how to answer common questions, how to handle objections, how to deal with ethical and security concerns, how to encourage people, and what meetings to set up.
- Train them first, for the skills and a common approach, then train them to train others.
- Protect their time explicitly, with their line manager's agreement.
- Make the ask narrow and specific: "run one 20-minute session a fortnight" beats "drive adoption in your area".
- Judge the network at three months, and expect to extend it.

WHERE WE COME IN

We run champions programmes with pre-built material, so champions spend their time helping people, not building content.

Champion time: the plan and the reality



From our programmes. The day job wins, so design the role for the time people actually have.

15 Somewhere to go that stays current, and a way to hear what people need

People need one place to go when they want to learn more or share what they've found. And you need to know what they're asking.

WHAT WE'VE SEEN

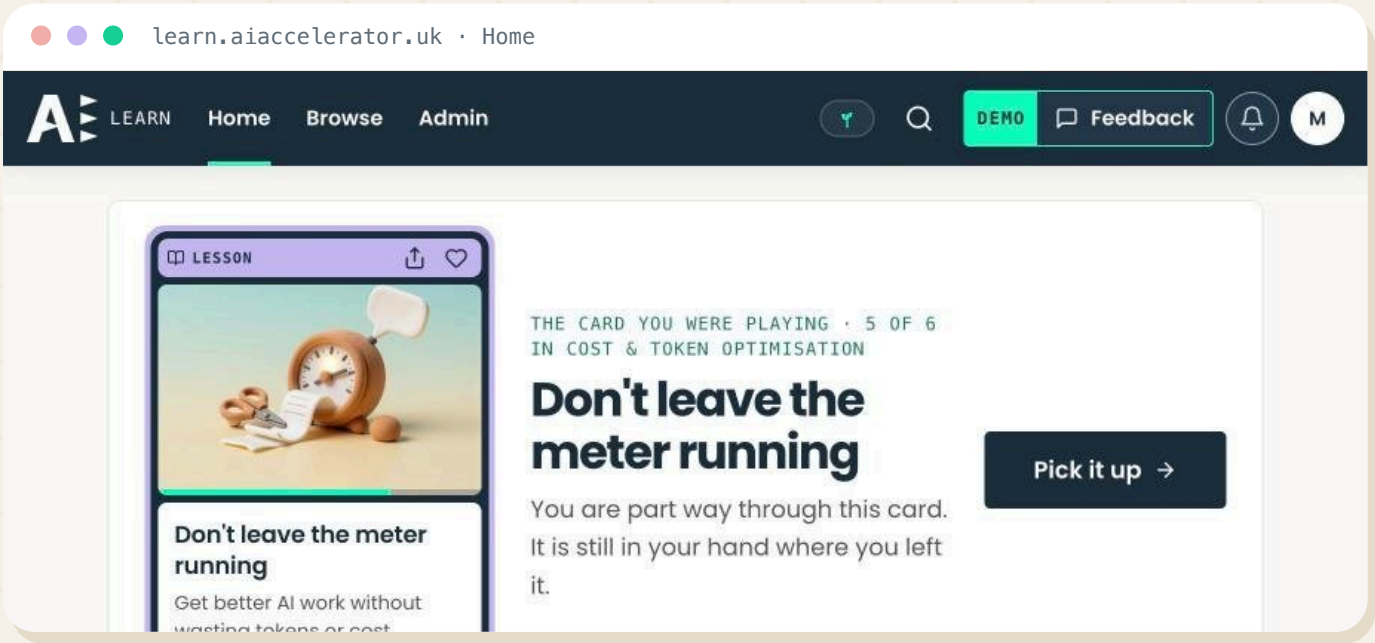
Internal AI hubs mostly end up quiet: people visit at launch and don't come back, because nothing changes and nobody posts. What works is a source that keeps updating, pitched at the right level for each person, plus a steady read of what people are asking. You'll learn more from the questions in a busy webinar than from any survey you send.

WHAT TO DO

- [x] Give people one place for learning and sharing, and make sure it changes every week.
- [x] Point each person to content that fits their role and level. A beginner and an experienced user shouldn't get the same path.
- [x] Treat questions as data. Read webinar Q&A, session chat and support requests every week, and let them set the next topics.

WHERE WE COME IN

Our platform recommends content for each person's role and level, and our weekly content is shaped by what people across all our clients are asking.



Learn: each person's hand of cards, what's trending, and what the organisation has dealt them

16 Measure adoption with five metrics, and set a target for each

Most programmes report one number: how many people logged in, and call it active users. It's the easiest number to get and it tells you almost nothing. Someone who opened the tool once and someone whose week runs on it count the same. Measure four signs that the habit is forming, and one that it's paying off.

WHAT WE'VE SEEN

On one client's weekly dashboard, two teams had the same number of "active users". They looked nothing alike once we asked how often, how deeply, how widely and whether anyone was building. verify

WHAT TO DO

Count everyone the programme is aimed at and take a baseline before you start. Then set a target for each measure against that baseline. The targets below are starting points, not benchmarks.

Report time-back stories alongside the numbers, because sponsors remember stories. If your tool only reports weekly, or not by team (lesson 12), use a one-question Friday pulse until you can get better data.

Consistency is it a habit?	The share of people using AI on 3 or more working days a week.	**Starting target:** half by month six, with the "no habit" group shrinking every month.
Depth are they doing real work with it?	How much people do with AI when they use it, in bands set from your own baseline.	**Starting target:** the typical user moves up one band by month three.
Breadth are they using more than chat?	The share of regular users who use more than one kind of feature: files, projects, skills, agents.	**Starting target:** one in two regular users by month six.
Build is anyone making things others use?	Shared skills and assistants that other people use every week. Count the users, not the builds (lesson 8).	**Starting target:** every team has at least one shared skill that most of the team uses weekly.
Value is it paying off?	One business measure for each priority workflow, such as cycle time, cost per case or turnaround, baselined before you start and agreed with finance.	**Target:** set it with your CFO before the pilot, along with the point at which you'd stop.

17 Plan with what you know, and review it every quarter

As Ethan Mollick puts it, "assume this is the worst AI you will ever use." verify: Mollick Every plan you make is built on tools that will be overtaken, so the plan needs a rhythm for changing. Commit to the best plan you can make today, and look for what would prove it wrong.

WHAT WE'VE SEEN

We scoped our own intranet, with interactive pages and dashboards for the whole business, as a months-long project. It took a week. verify The plan wasn't wrong when we made it. What we could do had changed by the time we started.

WHAT TO DO

- [x] Plan with the best information you have. Don't wait for the market to settle, because it won't.
- [x] Review the plan every quarter. Ask what has changed in the tools, in what your people are asking for, and in what the measures show (lesson 16). Then update it.
- [x] When something leaps ahead between reviews, decide in writing: adopt it now, after testing it in your own environment (lesson 19), or hold until the next review, with the reason.
- [x] Build on what lasts (how AI works, habits, judgement), and treat tools and features as replaceable parts (lessons 4 and 20).
- [x] Give every quarter a finish line. The plan has no final version, but each quarter should say what done looks like in 90 days, and what would make you stop.

WHERE WE COME IN

We run the quarterly review with programme leads: what changed in the market, what your people are asking, what done looks like for the next 90 days, and what to change.

PART 3

Don't do this

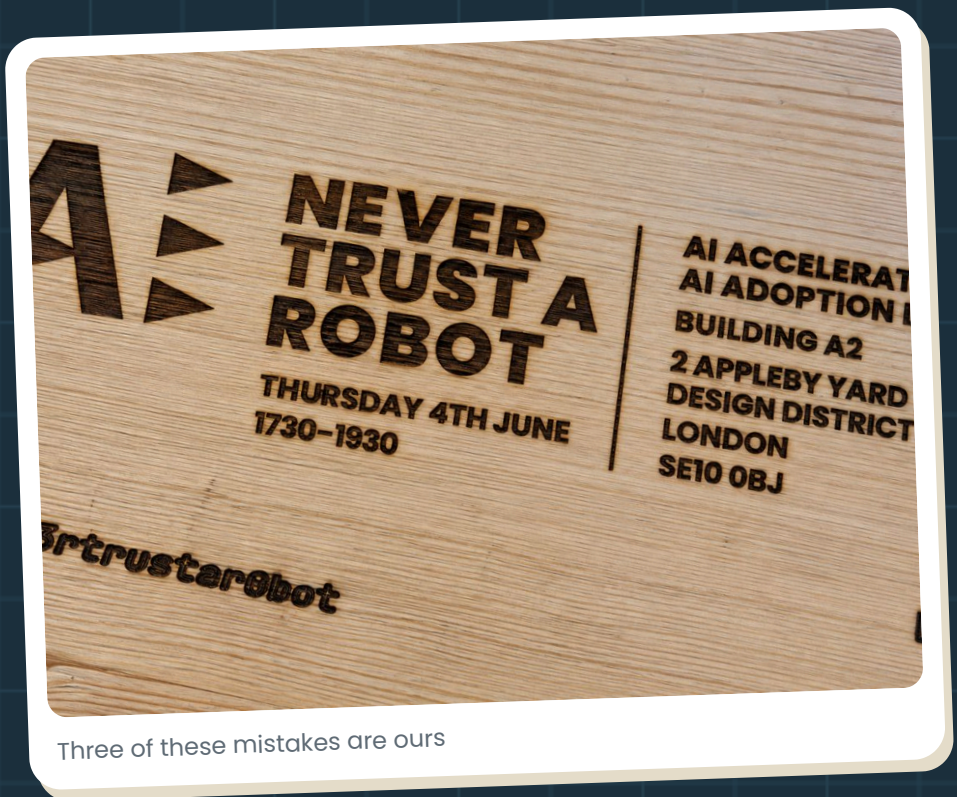
18 Don't build your own AI learning platform

19 Don't teach or demo anything you haven't tested where it will run

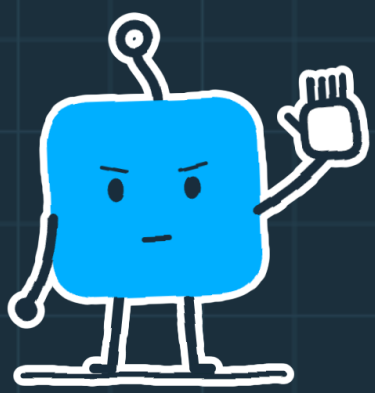
20 Don't build methodologies for a moving target

21 Don't deliver a story that isn't yours

22 Don't put everyone in the same room at the same level



Three of these mistakes are ours



18 Don't build your own AI learning platform

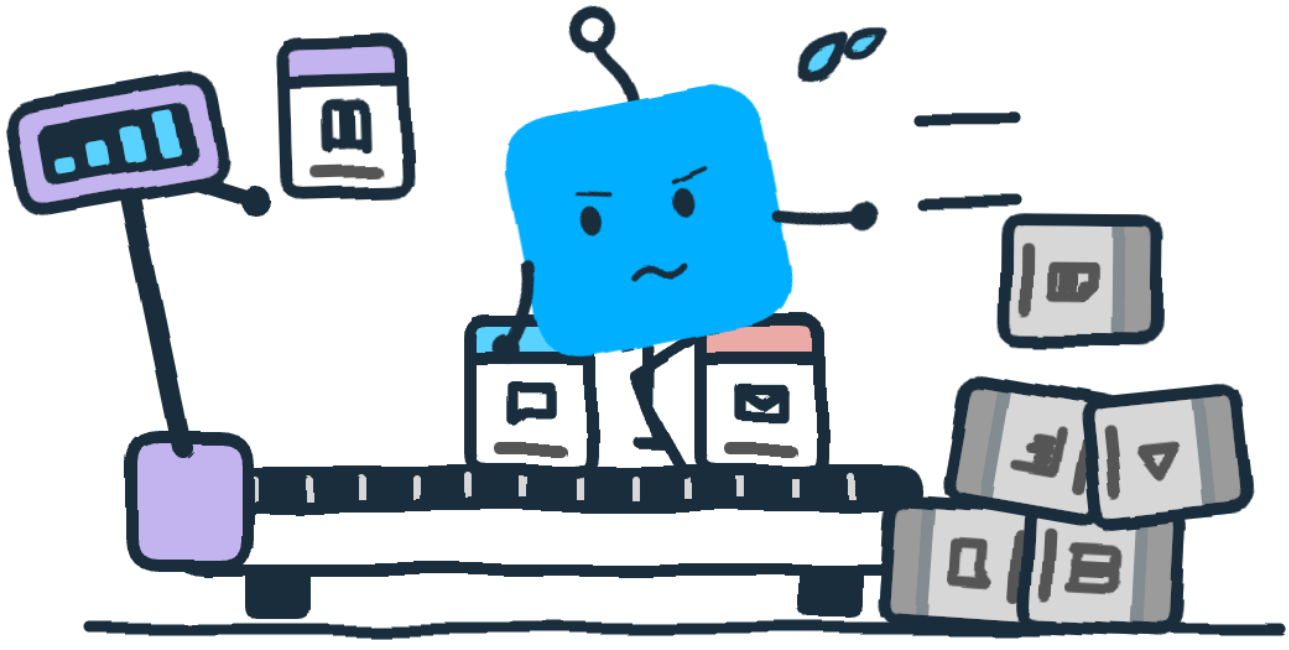
It seems sensible: put the AI training on your own intranet, in your own branding, built for your people. We've helped clients do exactly that.

WHAT WE'VE SEEN

There was too little content, and it was the wrong content: generic material, not what people were asking. It was also out of date almost as soon as it went live. A new model arrived, a feature moved, a vendor changed what was included. The client was then on a treadmill of making content, catching up, and making more, and most people stopped visiting. Nobody in-house is full-time on this, and nobody is set up to keep pace.

INSTEAD

Use a source that someone else keeps current, and spend your own effort on the parts only you can do: your policies, your use cases, your people.



The content treadmill: every week there is more to make, and last month's goes stale.

19 Don't teach or demo anything you haven't tested where it will run

OUR MISTAKE

What a vendor says a product does, and what your people can actually do with your licences, settings and security policies, are often different things. We learned the demo half of this the hard way.

WHAT WE'VE SEEN

Features shown in vendor demos that weren't switched on for the people being trained, or weren't in the licence they had. Agentic features still "under evaluation" by IT months after they were announced. And our own mistake: we built an AI tool for a client to use live in a workshop, couldn't test it inside their systems beforehand, and on the day it didn't perform. We had a fallback exercise and the session held, but it was a lesson.

INSTEAD

- [x] Test every feature you plan to teach in your own environment, with an ordinary user's account, before you build training around it.
- [x] Test every demo where it will run, with the access your audience has. If you can't, have a fallback ready and say so.

20 Don't build methodologies for a moving target

OUR MISTAKE

We've been here. Early on we built frameworks, scaffolds and methodologies for everything: how to run discovery, how to structure a programme, how to sequence the learning. Some of them were good. Most were out of date before we'd finished them, because the technology underneath kept moving.

WHAT WE LEARNED

A methodology built for this year's AI is a constraint next year. There's an element of working it out as you go that every organisation has to accept. Keep principles fixed and methods light, and expect to change them.

21 Don't deliver a story that isn't yours

OUR MISTAKE

WHAT HAPPENED

We've delivered sessions built on someone else's material, where the slides were handed over but the story behind them wasn't, and there was no full dry run. It didn't land.

INSTEAD

Whoever presents needs the author's narrative, not just the deck, and a full rehearsal with the people who know the material. If a session is going to run several times, record and transcribe the first one and fold what was said back into the notes.

22 Don't put everyone in the same room at the same level

Your organisation contains people who have never opened an AI tool and people who've been using it daily for years. Put them in the same session and you'll bore one group and lose the other.

WHAT WE'VE SEEN

Advanced users rated foundation sessions lowest, and some said so. Beginners in advanced sessions went quiet. One session format didn't fit a leadership team, a frontline team and a technical team.

INSTEAD

Find out where people are (lesson 12), then offer different routes: a foundation for everyone who needs it, and a faster track for those who don't. Let people check in with peers at their own level.

Two routes, one programme

Foundation — For everyone who needs the basics first →

Faster track — For people already using AI every day →

Find out where people are first (lesson 12), then let them check in with peers at their own level.

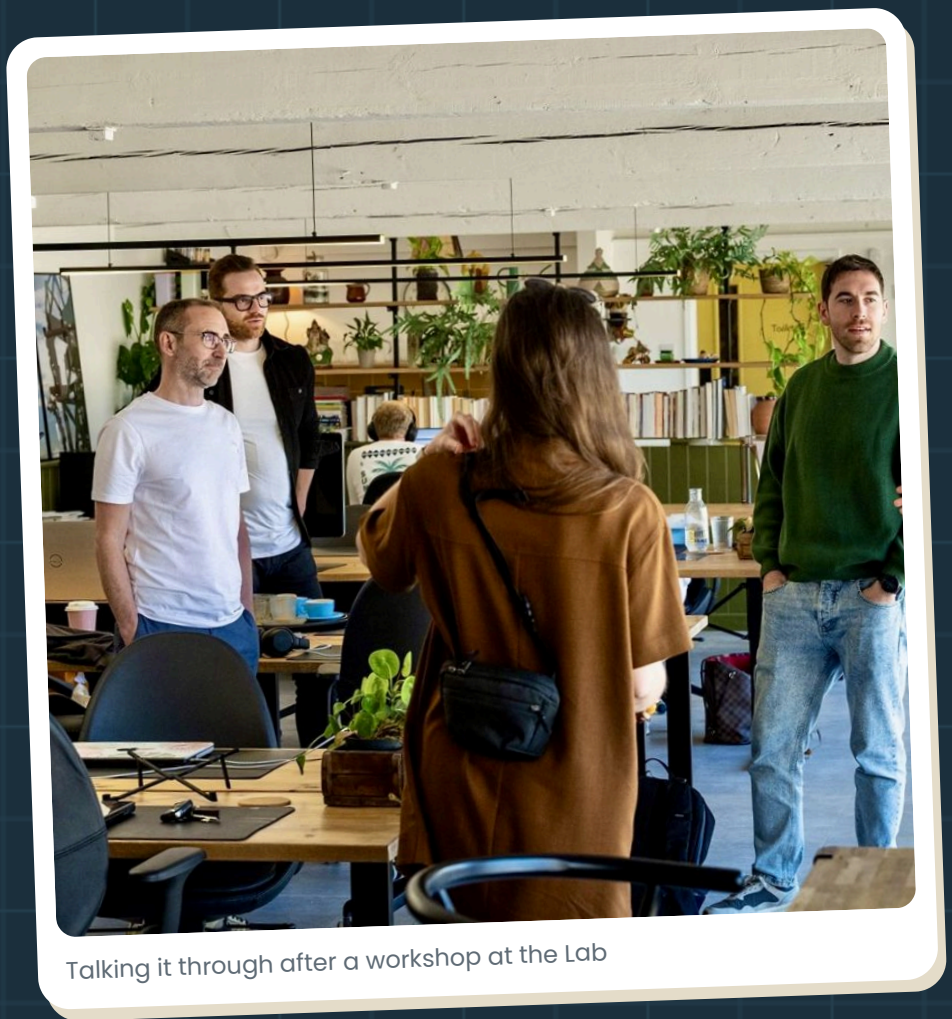
PART 4

Don't worry about this

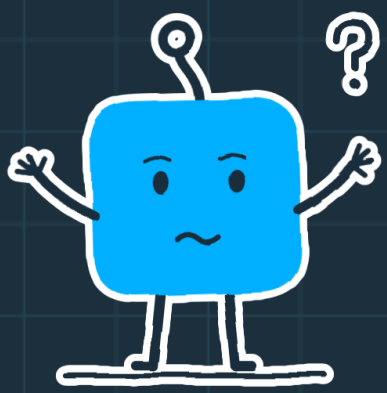
23 The perfect tool, or the perfect prompt

24 A centre of excellence before the crowd moves

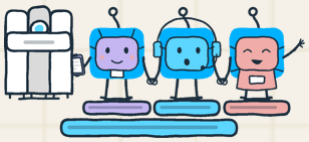
25 The late adopters



Talking it through after a workshop at the Lab



23 The perfect tool, or the perfect prompt



People lose months debating Copilot against ChatGPT against Claude against Gemini, and you'll be offered prompt formulas and courses on writing the perfect prompt. Both are methods built on how AI used to work (lesson 17). The tools leapfrog each other, and each model release needs less coaxing.

WHAT WE'VE SEEN

What appears in one tool turns up in the others soon after: skills are now a core part of what Microsoft offers. And given the same context, a plainly worded request now does about as well as an engineered formula verify. What makes the difference is the context: the document, the audience, what good looks like. What people need to learn has moved on, from writing prompts to handing work to AI and checking what comes back.

INSTEAD

- Use the tool you've got, and revisit the choice every quarter as part of the plan review (lesson 17).
- Teach people to give AI context: what it's for, who it's for, the material it needs, what good looks like.
- Judge a template by its blanks. If they ask for your context, it's worth keeping. If they ask for magic words, it's out of date.
- Turn the context you keep giving into a skill, so nobody has to type it again (lesson 8).
- Ask any provider what happens to their prompt method at the next model release.

24 A centre of excellence before the crowd moves

A lab or centre of excellence sounds like the serious option. Without people across the organisation already using AI, it has nobody to serve and nothing to build on.

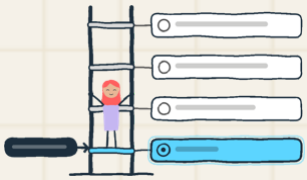
WHAT WE'VE SEEN

One client's lab never got going: it needed systems, access and governance that weren't in place, and the effort went elsewhere.

INSTEAD

Get the crowd moving first. The people who should be in your centre of excellence will reveal themselves.

25 The late adopters



Early sessions fill easily. Then demand drops, and it can look like the programme is failing.

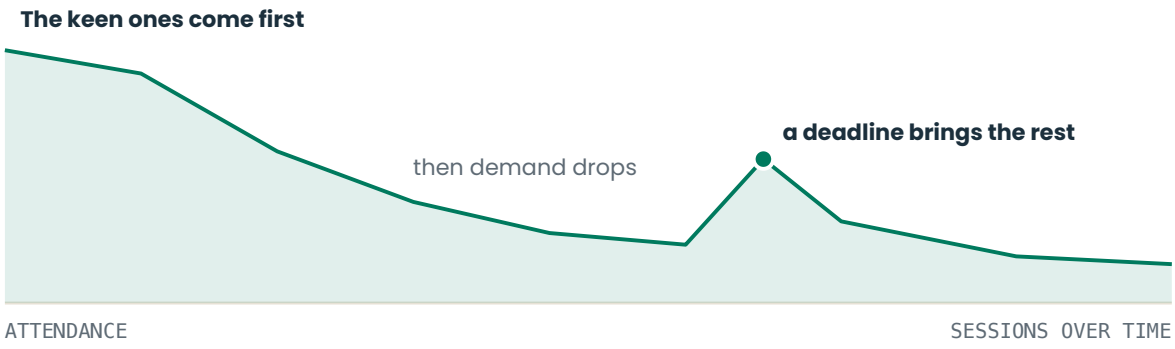
WHAT WE'VE SEEN

Attendance falls as you move from the people who were keen anyway to everyone else. It's normal. At one client a company-wide deadline for the foundation session brought the rest through.

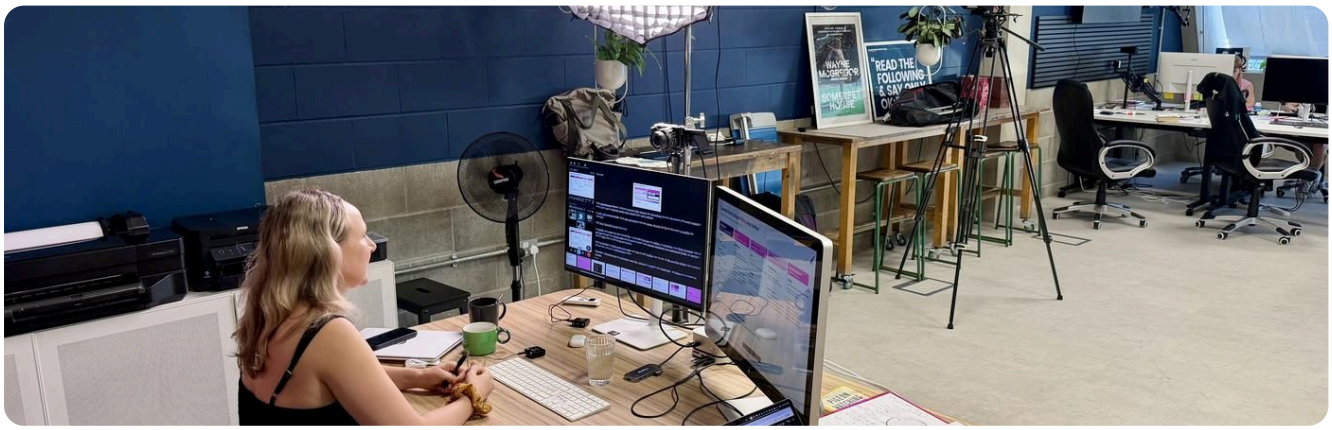
INSTEAD

Expect it, set a date, and keep the weekly slot going for the people who come later.

Why demand drops, and why that's normal



Illustrative shape, from what we've seen across programmes.



The weekly slot keeps running for the people who come later

Run adoption the new way



Most adoption programmes are still run the way training was run ten years ago: a launch email, a slide deck, a survey as the only feedback, an intranet page. Yet the tools you're introducing can transcribe, build, analyse and respond. Use them on the programme itself.

THE OLD WAY	THE NEW WAY
~~A survey as the only signal~~	Telemetry from the conversation: what people ask, what confuses them, what they come back to
~~A slide deck and a PDF handout~~	Interactive tools people can try: explainers they click through, exercises that respond, cards they work through at their own pace
~~An intranet page nobody visits~~	Artefacts built for your people: a page, a tool or a small site made for one team, in an afternoon
~~A training calendar planned a quarter ahead~~	A programme that steers itself: what people ask this week sets what you teach next week
~~Write-ups a fortnight after the event~~	Transcripts turned into an FAQ, a summary and next steps the same day

WHAT WE'VE SEEN

We run our own programmes this way. We transcribe our sessions and webinars, and the questions people ask shape the next session and our FAQ library. Our explainers are interactive, and this guide was built with the same tools we teach. When people see the programme itself being run with AI, they stop asking whether it's useful and start asking how.

WHAT TO DO

- [x] Record and transcribe sessions, with consent, and turn each transcript into an FAQ and next steps the same day.
- [x] Read the questions as telemetry: what's confusing, what's working, what people want next.
- [x] Build small interactive things instead of handouts: an explainer, a checker, a page for one team.
- [x] Show your people how you made them. It's the best demo you have.
- [x] Mind the edges: consent to record, where transcripts are kept and who can see them (lesson 11).

Where we come in. We build these for clients, and we teach their teams to build their own.

Where this is heading

So far, adoption has mostly meant training people to use AI tools. That's changing.

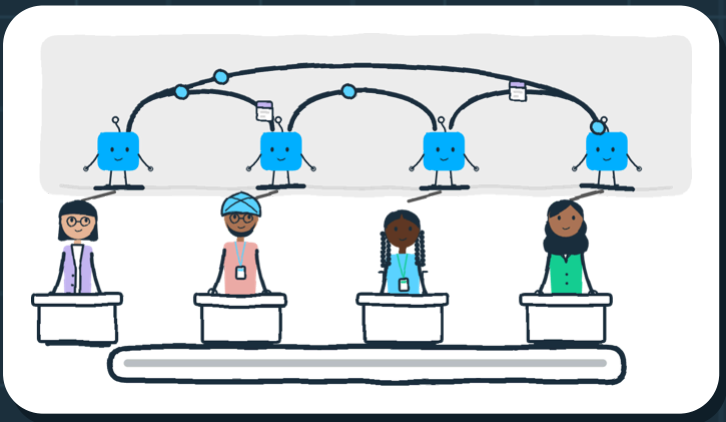
WHAT IT
knows
Your context, your documents, its memory of past work

WHAT IT
can do
The tools and skills it can use

WHAT IT
may do
Where it works, and what it's allowed to do

The organisations moving fastest are starting to build their own harnesses: the set-up around the AI that decides what it knows (your context, your documents, your memory of past work), what it can do (the tools and skills it can use) and what it may do (where it works and what it's allowed to do). Being able to do something isn't the same as being allowed to. Knowing which is which, for your organisation, is becoming a leadership question as much as a technical one.

The next step is AI that works across a whole organisation, not just for individuals: shared skills, shared context, people and AI working together on the same work. That multiplayer way of working is a big challenge, around data, permissions, accountability and how teams are organised. It's also the biggest opportunity in this whole area, and the organisations that have done the groundwork in this guide will be the ones ready for it.



You don't need to be there yet. You do need someone watching it for you.

The five common approaches, and where they fall short

Most organisations go about adoption in one of five ways. Each has something going for it. Each has a gap.

	GOOD FOR	WHERE IT FALLS SHORT
1 **Generic trainers and course libraries**	Good for basics at scale.	But they teach tools, not how to apply AI to your work, and the content dates quickly. There's rarely a route for each person that makes sense for their job, and rarely anyone who can answer the questions that matter.
2 **Large consultancies**	Strong on strategy and governance documents.	But the model is selling ongoing transformation and change programmes, and building AI capability in your people is better done by people who use these tools every day, or by your own people with the right support.
3 **Vendor adoption kits**	Free, well produced, and tied to one product.	They rarely go deep enough, they're written to sell the next licence tier, and they aren't written by people running a business on AI.
4 **Doing it in-house with L&D alone**	You know your organisation better than anyone.	But no internal team can keep up with the pace of change alongside everything else on its plate, and keeping content current becomes a full-time job nobody was hired for.
5 **Solo AI experts**	Often brilliant individual users with great tips.	But an organisation isn't one person. Working with AI across teams raises questions about data protection, GDPR, meeting transcripts, information leaking between systems and who's accountable for what, and an individual's experience rarely covers them.

How we work, and why

We're an AI capability partner. Our work comes down to three things: *Skills for your people. Skills for your AI. Judgement for your leaders.* Everything above shaped how we do it.

Practitioners, not trainers

Everyone who teaches for us uses these tools every day to run our own business and build for our clients. That's why we can leave the Q&A open, and why our answers come from experience rather than a script.

Always current

AI moves fast. We make sure you do too. Our platform and the content engine behind it are updated every week, and what we learn in one organisation's sessions improves what everyone gets. That's the answer to the treadmill in lesson 18: we run it once, for all of our clients, so you don't have to run your own.

Studio-produced, at scale

We run webinars from a professional studio, with a producer and a practitioner, from regular sessions for a hundred or more people to broadcasts for thousands verify. The questions people ask feed straight back into what we teach next.

Advice at leadership level

Our senior team works with boards and executive teams on direction: what's worth building, what to wait for, and what the next model release will make redundant. We bring experience from inside large organisations and from running an AI-native business ourselves.

We run our own business on AI

We're an AI-native business, and we use what we teach every day. Our intranet, with its interactive pages and dashboards for the whole company, was built with AI in a week (lesson 17). Our knowledge, our processes and our client work run on shared skills. When we tell you something works, it's because we're using it. When we tell you to wait, it's usually because we tried it.

Built for organisations, not individuals

Data protection, GDPR, information security, multi-team working: all of it is part of how we teach, because that's what it takes for AI to work in a real organisation.

NEXT **Start wherever makes sense**

Where to start

You don't have to commit to a programme to find out whether we're useful. Start wherever makes sense.

STEP 1

Try a webinar

A studio-produced session for your people, on the topic you choose, with live Q&A. **From £1,500 + VAT per session**. *Start here if you want to test the water or reach a lot of people quickly.*

STEP 2

Book a half-day session

A leadership strategy session, or hands-on training for a team. **£3,500 + VAT, plus expenses**. *Start here if your leadership team needs to get it first, or one team is ready to move.*

STEP 3

Give everyone the platform

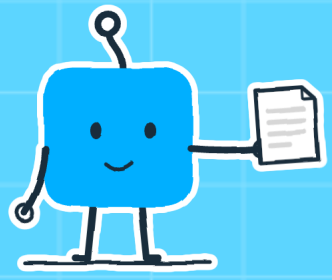
Short learning cards, content picked for each person's role and level, recorded sessions and a weekly flow of what's changed. **From £2,000 a month + VAT** for the whole organisation, priced by size, never per seat. *Start here if you want everyone learning from what's working elsewhere, without building it yourself.*

STEP 4

Design and run a full enablement programme

Leadership, foundations, champions, functional workshops, governance and measurement, designed with you and delivered with you, built so the capability stays with your people. **Priced on conversation**. *Start here if you've got the licences, the mandate and the appetite to do it properly.*

hello@aiaccelerator.uk



AI Accelerator Ltd. Client organisations are not named in this guide. The lessons come from our enablement programmes and our own business, 2025 to 2026.