

Your first 30 days with AI

Nobody trained you. This is what to do about it, on your own, in about twenty minutes a week.

One rule throughout: every step uses work you already have to do. None of this is practice. If you find yourself inventing exercises, you've drifted.

Week 1 — One task, low stakes

Goal: one real thing finished with AI, on work that was already on your list.

Pick a task that is repetitive, low-stakes, and *checkable* — meaning you'll know immediately if the output is wrong. Meeting notes, a first draft of a routine email, reformatting something.

Don't pick your hardest problem. The point this week is not to be impressed; it's to find out where the tool is and what it does badly.

- ☐ Read the *Can I put this into AI?* card before you paste anything
- ☐ Do one task end to end
- ☐ Write down one thing it did badly

That last box is the one that matters. Everyone remembers the impressive part; the useful knowledge is the failure.

Week 2 — Give it context

Goal: notice that the same request works better when you say more.

Take the task from week 1 and do it again, but supply what the model couldn't have known: who it's for, what happened before, what you're actually trying to achieve, an example of a good one.

Compare the two outputs side by side.

- ☐ Redo week 1's task with three sentences of context added
 - ☐ Note the difference — this is the whole skill, and it isn't prompt wording
 - ☐ Start a file called `prompts.md`. Paste in the version that worked.
-

Week 3 — Start checking properly

Goal: stop trusting fluent output.

Use the *three checks* card on something real before you send it. Find the load-bearing claims. Ask what the model couldn't have known. Recompute one number.

- ☐ Run all three checks on one piece of work
 - ☐ Do the "what couldn't it have known" check deliberately — it's the one nobody does
 - ☐ Notice whether you caught anything. Sometimes you won't. That's fine — the habit is what you're building.
-

Week 4 — Make one thing repeatable

Goal: stop starting from scratch.

Take the task you've now done three times and save the setup — the context, the format, the instructions — somewhere reusable. Your tool may have a feature for this; a text file works just as well.

- ☐ Save one reusable setup
 - ☐ Use it once without editing it
 - ☐ Decide what your *second* task is going to be
-

After 30 days

You'll have one task genuinely faster, one reusable setup, a checking habit, and a short list of things the tool is bad at. That's a real base, and it's more than most people around you will have.

Two things to keep doing:

Add to the library when you explain something twice. Not when a task is important — when it recurs.

Watch your own trust. The moment worth noticing isn't when you start using AI more. It's when you stop reading the output carefully. If you can't remember the last time you disagreed with something it produced, go back to the three checks.

What this deliberately doesn't include

No prompt-engineering course, no tool comparison, no list of 50 use cases.

Those feel productive and mostly aren't. Four tasks done on real work will teach you more than any amount of reading about it — which is, itself, the thing most AI training gets wrong.

Related reading: How to Learn AI at Work When Nobody's Training You and How to Pick Your First AI Task at Work (Without Getting Burned)
