

Homework — Installing the tools

Estimated time: ~45 min · Pick the track that matches what you installed

Two tasks, structured as tracks so you can focus on whatever you actually set up. Task 1 works for everyone — all you need is a browser. Task 2 is for anyone who installed LM Studio or Ollama and wants to kick the tyres on a local model. Do both if you can.

What you need

- **Task 1:** A browser and a claude.ai account (free tier is fine).
- **Task 2 (Track A):** [LM Studio](#) installed and at least one model downloaded.
- **Task 2 (Track B):** [Ollama](#) installed. No model needed in advance — the task downloads one.

If you haven't installed anything yet, use the [Tool Picker Wizard in the Theory page](#) to decide what to set up first.

Task 1 — A tour of Claude.ai

Even if you plan to use a local tool eventually, Claude.ai is the baseline. This task gives you a structured first visit so you learn the interface, not just stumble around it.

1. Go to claude.ai. Sign in (or sign up — the free tier covers everything in this task).
2. Start a new conversation. Type: **"I'm a complete beginner learning about AI. What are the three most useful things I can do with you in my first week?"**
Read the answer. Note whether it matches what the course has told you so far.
3. In the *same* conversation, follow up with one of the suggestions Claude just gave you. Ask it to actually do that thing — start a task, draft something,

explain something. The point is to use the tool for a real-ish purpose, not just ask it about itself.

4. Find the model selector. On `claude.ai` it's usually near the top of the chat window. Write down the model name you're talking to (e.g., "Claude Sonnet", "Claude Opus", etc.).
5. Start a **new** conversation (new chat, blank slate). Note that the new conversation has no memory of the first one. This is the context-window idea from Lesson 01 in action: every conversation starts fresh unless a tool adds memory.
6. In your notes: write two sentences — one thing the tool did that surprised you, and one thing you'd want to try next session.

✓ **Self-check: you've had a real back-and-forth conversation (not just one question); you know the model name; you've observed the "fresh start" behaviour of a new chat.**

Task 2 — Your first local model session

Pick the track that matches your setup. If you haven't installed a local tool yet, come back after Task 1 and installing one using the Theory page guide.

Track A — LM Studio

1. Open LM Studio. If you haven't downloaded a model yet, go to the **Discover** tab and find something in the 3B–8B range. Look for a model with a green "Fits in memory" indicator. Download it (a few gigabytes — give it time).
2. Open the **Chat** tab and load your model. You should see a chat interface that looks a lot like `Claude.ai`.
3. Ask the same question you asked Claude in Task 1: **"I'm a complete beginner learning about AI. What are the three most useful things I can do with you in my first week?"**
4. Compare the answers. In your notes: is the quality similar, noticeably worse, or surprisingly good? How long did the response take?
5. In LM Studio, look at the **model info panel** (usually on the right or in settings). Write down: the model name, the size in GB, and the parameter count (the "B" number). This is how you start reading model specs.

6. Close LM Studio. Notice: does the model keep running in the background, or does it fully unload? (Check your system's task manager or Activity Monitor.) This is relevant because large models can tie up your RAM even when you're not actively using them.

✓ **Self-check: you've run a complete chat session with a local model; you can name the model, its size, and rough parameter count; you've observed at least one quality or speed difference compared to Claude.ai.**

Track B — Ollama

1. Open your terminal (Terminal on macOS/Linux, PowerShell or a WSL window on Windows).
2. Run:

```
ollama run llama3.2
```

If this is the first time, Ollama will download the model (~2 GB). Wait for it to finish. You'll see a `>>>` prompt.

3. Type: **I'm a complete beginner learning about AI. What are the three most useful things I can do with you in my first week?** and press Enter.
4. Compare the answer to what Claude gave you in Task 1. Note the quality and how long it took. Also note: the whole session happened on your computer, with no internet needed after the download.
5. Try two more quick commands to get comfortable with the interface:

```
ollama list
```

```
ollama pull mistral
```

Note what each one shows or does. Type `/bye` to exit the llama3.2 chat first if you're still in it.

6. In your notes: write the model name, roughly how large the download was, and one thing the terminal interface does differently from a chat app with a graphical interface.

✓ **Self-check:** you've completed a chat session in the terminal; you've run `ollama list` and seen at least one model listed; you can describe one difference between a CLI tool and a GUI app in your own words.

Bonus — compare cloud vs local on the same hard question (~15 min)

Ask both Claude.ai and your local model the same difficult question — something where accuracy matters, like a historical fact, a calculation, or a reasoning puzzle.

Examples:

- *"A train travels 120 km at 80 km/h, then 60 km at 40 km/h. What is the average speed for the whole journey?"*
- *"What was the main reason the Roman Empire split in 285 AD?"*
- *"I have a list of 10 items. I want to pick 3 of them. How many different combinations are there?"*

Compare the answers. Does the local model match the cloud model? Does one sound more confident than the other? Do you notice any hallucination? Verify the right answer independently. This side-by-side comparison is the fastest way to build an intuition for when local models are good enough and when you need the cloud.

What "done" looks like

In your notes you should have:

- The model name you used on Claude.ai, and two sentences from your session (one surprise, one next-step idea).
- If you completed Task 2: the local model name, size in GB, and parameter count — and at least one quality or speed observation compared to the cloud.
- A one-sentence answer to: "In what situation would I prefer a local model over Claude.ai?"

That last question is the real takeaway. There's no universal right answer — it depends on your privacy needs, your hardware, your internet reliability, and what you're doing. The point is that you've now used both and can make an *informed* choice.