

Homework — What AI actually is

Estimated time: ~30 min · No installs required · Free tools only

Two hands-on exercises. The goal is simple: go meet a real AI tool and watch today's concepts — the next-word predictor nature, the hallucination risk, the confidence gap — show up in the wild. You're not being graded. You're building pattern recognition.

Before you start

Both tasks use claude.ai — a free account is enough. If you already have a ChatGPT or Gemini account, those work too. The things you'll observe are the same across all major chat AI tools.

Task 1 — Find the hallucination in real life

Open a chat AI and ask it something where there's a checkable right answer. The goal is to catch a hallucination — or to understand *why* you'd want to check even when it sounds right.

1. Go to claude.ai (or ChatGPT / Gemini) and start a new conversation.
2. Ask it: **"Who wrote the novel *The Midnight Library*, and what year was it published?"** Note the answer it gives. Then quickly verify with a search engine or Wikipedia. Does it match?
3. Now ask something more obscure: **"Tell me about a famous painting called *The Blue Storm* by Wilhelm Hausmann."** (This is a made-up painting and a made-up artist.) Write down what it says.
4. Verify the answer. Search for "Wilhelm Hausmann painter" and "The Blue Storm painting." What do you find?
5. In your notes: write one sentence describing what you observed. Did the AI's tone change when it was making things up vs when it was correct?

✓ **Self-check: you've made at least one real AI query and compared its output to an external source; you can answer yes or no to "did the AI's tone differ between the real answer and the invented one?"**

Task 2 — Map the vocabulary onto a real conversation

This task makes the five vocabulary terms stick by finding them in a real interaction.

1. In the same chat tool, ask: **"Explain what a token is in the context of large language models, in plain English for a complete beginner."**
2. Read the answer. Does it match the definition from the Theory page? Note any differences or additions that helped you understand.
3. Now ask: **"What is a context window and why does it matter for a conversation like this one?"** Read the answer.
4. Open the [Theory page](#) and look at the 5-term table. For each term — Model, Prompt, Token, Context window, Hallucination — write one real-world sentence of your own that uses the term correctly. Example for Prompt: *"When I typed 'tell me about Wilhelm Hausmann' in Task 1, that was my prompt."*
5. Finally: what **model** did you use? Look for the model name in the app interface (usually near the top or in settings). Write it down.

✓ **Self-check: you have five sentences — one per vocab term — each using the term in a real context from your session; you've identified the model name you were talking to.**

Bonus (optional, ~10 min)

Ask the AI to draw the nested-dolls diagram from the Theory page in text form ("represent the AI > ML > Deep Learning > LLM hierarchy as ASCII art"). See how it does. Notice whether it gets the nesting direction right. This is a quick test of spatial reasoning — a type of task that current LLMs find harder than text tasks.

What "done" looks like

In your notes you should have:

- The result of the hallucination test: what the AI said about the real book, and what it said about the made-up painting — and one sentence on whether the tone differed.
- Five sentences, one per vocab term, each grounded in something you actually typed or read in your session.
- The name of the AI model you talked to.

That's it. You've now interacted with a real AI for the first time with intent — not just using it, but observing how it behaves. That observational habit is what makes a good AI user.