

Homework — Types of AI

Estimated time: ~30 min · No code required

Three short exercises. The goal is to move the four-lens framework from "something you read about" to "something you can apply in five seconds when you see a new AI tool." You'll categorise real AI headlines, compare two live tools side-by-side, and test a local model if you have the hardware.

How to do these

Keep the theory page open in another tab — the interactive grid is a useful reference. No installs required for Tasks 1 and 2. Task 3 is optional and only for people who want to try local AI. Take notes as you go; the self-check questions at the end of each task tell you what "done" looks like.

Task 1 — Classify five AI headlines

Below are five recent-style AI headlines. For each one, write down all four labels: *Capability / Output / Modality / Where*. There are no trick questions — if the headline is ambiguous, write down your best guess and note why it's uncertain.

1. **"Google's new model beats humans on a protein-folding benchmark."**
2. **"Runway releases Gen-3, a tool that generates 10-second video clips from a text prompt."**
3. **"Banks deploy AI to flag unusual card transactions in real time."**
4. **"Researchers release Llama 4 weights publicly; anyone can download and run it."**
5. **"ElevenLabs voice clone reads your book to you in your own voice."**

✓ **Self-check: for each headline you can fill in all four labels (Narrow/General/Super; Predictive/Generative; Text/Image/Audio/Video/Multi; Cloud/Local/Either) and give a one-sentence reason for each label. If a label is genuinely ambiguous, you can explain why. Compare with a friend or colleague who has also done this lesson — you'll likely agree on 90% of the answers.**

Task 2 — Side-by-side: cloud vs cloud, same prompt

This exercise makes the "same underlying model, different doorway" concept concrete by sending the *same prompt* to two different cloud text tools and comparing the results.

1. Open claude.ai in one tab. Sign up for a free account if you don't have one.
2. Open chatgpt.com in a second tab. Sign up for a free account if you don't have one.
3. Send **exactly this prompt** to both:

Explain what a spam filter does in two sentences, as if you're talking to someone who has never heard of AI.

4. Read both replies. In your notes, write down:
 - Which reply was shorter?
 - Which felt more natural or easier to understand?
 - Did either one use technical jargon the spec asked it to avoid?
 - Did they give meaningfully different explanations, or essentially the same one?
5. Now pick **one more prompt of your own** — something you genuinely want to know about AI — and repeat the comparison.

What you're looking for

The models are different products trained by different companies, so they'll have different tone and style. But they're both cloud-hosted, generative,

text-focused, narrow AI tools. The differences between them are about training, not about the fundamental category.

✓ **Self-check:** you've sent the same prompt to both tools and can describe at least two concrete differences in the responses (length, tone, vocabulary, structure). You also tried a second prompt of your own choosing. You can answer: "Are Claude and ChatGPT the same type of AI?" (Yes — both Narrow, Generative, Multi/Text, Cloud) and "What's different between them?" (company, training data, some capability trade-offs, pricing).

Task 3 — (Optional) Try a local model

This task is optional. It requires 8–16 GB of free RAM and about 4–8 GB of disk space for the model download. If your computer is modest (old laptop, less than 8 GB RAM), skip this and come back to it in Module 4 where we set up local tools more carefully.

1. Go to ollama.com and download Ollama for your operating system (Mac, Windows, or Linux). Install it. It's free and open source.
2. Open a terminal (Mac: `Cmd+Space` → type "Terminal"; Windows: press `Win+R`, type `cmd`, Enter).
3. Type this command and press Enter:

```
ollama run llama3.2
```

Ollama will download the Llama 3.2 model (about 2 GB) and then drop you into a chat prompt. This may take a few minutes the first time.

4. Type the same prompt you used in Task 2 and press Enter. Compare the response to what Claude and ChatGPT said.
5. In your notes, write down:
 - How long did the response take? (Cloud responses are usually faster on a typical laptop.)
 - Did the quality feel similar, better, or worse?

- Did anything leave your machine? (Answer: no. This is the privacy payoff of local models.)

6. When done, type `/bye` to exit.

If you're on a slow computer

Llama 3.2 on 8 GB RAM with no GPU can take 30–60 seconds to produce a response. That's normal. It's not broken — it's just your CPU doing what usually takes a specialised GPU. If it's painfully slow, quit and skip to Module 4 where we cover hardware requirements in more detail.

✓ **Self-check: you ran a real AI model on your own hardware and got a response. You can say with confidence: "All processing happened on my machine. My prompt was never sent anywhere." You noticed a speed difference compared to cloud. You now have a real, working example of the Local / Open category from the theory page.**

What "done" looks like

You've finished this homework when you have in your notes:

- Five classified headlines with all four labels filled in and a brief reason for each (Task 1).
- A written comparison of Claude vs ChatGPT on the same prompt, plus a second prompt of your choosing (Task 2).
- (Optional) Notes from running a local model, including your observation on speed and the fact that nothing left your machine (Task 3).

With that, you have hands-on experience of at least two real AI tools, a framework for categorising any new AI tool you encounter, and — if you did Task 3 — you've also run a

model entirely on your own hardware. That last part puts you ahead of most people who use AI every day.