

Day 6 Student Exercises: Embeddings and the Frontier

Name: _____

Date: _____

Formative, not graded. Predict before you run a cell; discuss with your neighbour where asked.

Part A: Predict before you run (during the labs)

1. **Embeddings.** Of the six earnings snippets, which pair of companies will have the highest cosine similarity? Why?

2. **Cosine vs dot product.** Two embedding vectors point the same way but one is twice as long. Are the texts similar? Which measure tells the truth here?

3. **RAG retrieval.** For the query “revenue growth”, which snippets do you expect to be retrieved, and what will the final prompt look like (how many chunks, where does the question go)?

4. **Grounding.** Ask the model the question with no retrieved context. Predict: a correct answer, or a confident guess?

Part B: Discussion exercises

EX-D6-1. Attention by hand. Sentence: “The bank raised its rate after the central bank’s decision.”

- (a) For the word “rate”, which other words should attention weight most, and why?

- (b) What tells the model the first “bank” is a commercial bank and the second is a central bank?

EX-D6-2. Will an LLM alone get this right? Sort each into: LLM alone / needs RAG or a tool / plain script.

- (a) summarize this earnings-call transcript (text given): _____
- (b) Apple’s exact revenue in last week’s results: _____
- (c) compute the 30-day volatility of these returns: _____
- (d) draft a polite email declining a meeting: _____
- (e) is this filing’s tone bullish or bearish (text given): _____

EX-D6-3. Design a finance agent, then break it. For a “quarterly revenue growth” agent:

- (a) List its think-act-observe tool-loop steps:

- (b) Name one step where a silent error would corrupt the final number:

- (c) Name one guardrail you would add:

Exit ticket: name one thing you can do now that you could not on Monday. _____