

1. Think about your Exercise 1 conversation where you asked the LLM a question, got its reply, then asked “are you sure?”. **Write the document that was fed to the model to generate the first word of its reply to the follow-up question.** (You don’t need to mark the boundaries between tokens.)
2. Now imagine a conversation where you provide the model an *image* (say, of a whiteboard) and ask it “what’s on the board” and to make it look like a chalkboard instead. **Sketch what the document might look like** for the input and reply.
3. Two discussion questions:
 - a. Someone once argued (on a video call I was on) that ChatGPT is reliable on questions about Scripture interpretation because it has been trained on the Bible and so many people’s writings about it. Is this a good argument? Why or why not?
 - b. Suppose a teacher suspects that an essay was written by ChatGPT. So they copy-paste the essay into ChatGPT and ask “did you write this?”. Good idea or not? Discuss.

Before you leave, pick a couple of these questions to react to:

1. What was the most important concept from today for you?
2. What was the muddiest concept today?
3. How does what we did today connect with what you've learned before?
4. What would you like to review or clarify next time we meet?
5. What are you curious, hopeful, or excited about?