Cultivating Curiosity for Neuroscience using Generative AI

In this assignment, you will use a large language model (LLM) of your choice (such as ChatGPT, Claude, or Gemini) to explore a neuroscience question you're curious to learn more about. Rather than simply asking the AI for an answer and moving on, the aim is to engage in a real back-and-forth conversation: asking follow-up questions, pushing back when something seems unclear or too simple, understanding what is known and not yet known about a topic, and seeing how deeply you can understand a topic through dialogue. You'll submit the full transcript of your conversation, a brief note checking one of the LLM's claims against a course reading or outside source, and a short reflection (~350–500 words) on what you learned and how the conversation shaped your thinking.

The purpose of this assignment is twofold. First, it's an opportunity to go deeper on a topic in neuroscience that truly interests you. The brain is complex, and curiosity is one of the greatest gifts we have for learning about it. Second, and perhaps more importantly, it's a chance to develop critical habits around AI use that will serve you well beyond this course. LLMs are powerful learning tools and we can expand our thinking a great deal by using them, but they are not textbooks — they can oversimplify, "hallucinate" by confidently asserting things that are wrong, and flatten scientific complexity into tidy summaries. Learning to use them well means learning to question them, probe their reasoning, and verify their claims. That skill is, in many ways, a skill of scientific thinking.

This assignment is not about generating polished AI outputs or arriving at the "correct" answer. A messy conversation full of dead ends, confused follow-ups, and moments where you pushed back on something that didn't quite make sense is far more valuable than a smooth, linear exchange that stays at the surface level. What we're looking for is the quality of your curiosity and the honesty of your reflection, not the quality of what the LLM produced.

A few practical notes before you begin: LLMs occasionally decline to answer certain questions or give overly cautious responses. This is especially common with topics like drug use, mental health, or self-harm, even when the question is framed scientifically. If this happens, try rephrasing your question with more explicit scientific framing (e.g., "from a neuroscience research perspective...") or simply choose a different angle on your topic. Additionally, for security reasons, treat LLMs the way you would any public platform: don't enter sensitive personal information you wouldn't want stored or shared. Please do not share any personal medical information with an LLM and avoid asking it to interpret symptoms or provide any kind of diagnosis.

Instructions

1. Choose a question: Pick a specific neuroscience question you actually want to understand better. It should be open-ended and complex enough to sustain a real conversation, not answerable in a few sentences or bullet points. Some examples of questions at the appropriate level of depth are listed below. You are, of course, not limited to these questions and may choose a topic of your own interest.

   Example Questions:
   Why do we dream?
   How does stress affect memory?
   Why do habits form and how can we break them?
   What causes hallucinations to occur?
   What is intelligence?
   How do psychedelic drugs affect the brain?
   Why does music evoke emotion?
   How do infants learn language?
   How does meditation change the brain?
   Why does chronic pain persist?

2. Have a conversation with an LLM: You may use any LLM you would like. Your conversation should demonstrate multiple back-and-forth exchanges, at least one critical engagement with what the LLM is saying (pushback, skepticism, or request for clarification), and some evidence that your understanding has evolved. There is no minimum number of exchanges you must have, but the length of your submission should show that you made a sincere exploratory effort to understand a new topic.

*** At the outset, you should know that the quality of your questions (called 'prompts') will determine the quality of the output you get from the LLM. Good prompts are specific, ask for evidence of mechanisms and real-world examples, invite the AI to flag areas where it may be uncertain, request competing theories (if relevant), ask how scientists study this topic experimentally, ask what remains unknown, and identify common misconceptions about the topic. ***

3. Verify one claim: Identify one scientific claim from the conversation and fact-check it against a verifiable source (textbook, reputable website, scientific study, etc.). Note whether the LLM was accurate, incomplete, or misleading.

4. Write a reflection: Write a 350-500 word reflection on your exchange. You should not use an LLM to help you write the reflection. You may wish to explore some of these questions, but you do not need to explore all of them:

What did you learn that surprised you? How did your understanding evolve during the conversation? What kinds of prompts produced the most useful answers? What kinds of prompts produced surface level answers? Did the LLM ever confuse you, oversimplify, or contradict itself? What was its tone when it gave erroneous information? How did it respond if you pointed out that it was incorrect? How trustworthy did the AI seem? What do you think are the strengths and weaknesses of using LLMs to learn about new topics? Did this interaction change how you think about using AI? Did this interaction change how you think about learning itself?

5. Submit three things: the full transcript of your conversation (either copy-paste, or you can use the export function on many LLMs), your verification of an AI claim, and your written reflection as a single document.

Rubric (5 points total)

Comprehensive LLM transcript showing curiosity and critical engagement: 2 points

Verification of AI claim from an independent source: 1 point

Depth and rigor of written reflection: 2 points