Using LLMs for Learning in Applied Math

Heather Wilber

Sept. 3, 2026

Overview

These notes discuss the use of LLMs in applied mathematics classes by undergraduate students. A recorded lecture covering these notes is available here.

In these notes, I highlight ways in which LLMs can be useful to the learning process, as well as ways that they can inadvertently cause harm. With this in mind, we consider the following spectrum:

Useful/Good	Potentially Useful	Dangerous
Enhances learning	Be careful of pitfalls	Derails learning

We will define two important terms to be used throughout this discussion:

> **(1) Text:** Any non-interactive resource that supplies a bank of useful knowledge. Classically, this is your textbook. It might also be videos, lecture notes, noninteractive lectures, or web resources. Texts help us organize, transmit, and prioritize content and knowledge.

> **(2) Teacher:** An expert that mediates learning through interactions with learners, and measures student ability through assessment. Interaction may be via conversations, feedback, and working together, but it also occurs through curriculum design choices, homework design, and assessment design. These choices are meant to prompt growth or measure performance with many learning goals in mind. Your professor is of course a teacher, and so are your TAs.

While education involves much more than simply these two, these tend to be the dominant modes through which undergraduate students experience STEM and mathematics classroom content. *One may notice that, oddly enough, LLMs seem to live somewhere between or among these two terms, whereas outside of LLMs, they are clearly distinct.*

A long history of academic work exists on best practices in education involving the two resources of texts and teachers. We are only at the beginning when it comes to understanding best practices in education involving AI-based resources. However, we can make some inferences based on non-AI-based educational theory.

Two guiding questions

To help think about which kinds of practices are helpful and which are harmful, let's keep in mind two central questions:

1. How does AI change text navigation?

- Text navigation classically involves interacting with resources that do not interactively respond back. Skills such as summarization, synthesis, sifting between important and less important information, and the recasting of passive information into "owned" information are all exercised and developed when working with texts.
- AI-based text navigation involves giving an AI text and using the AI to navigate that text. It allows text to be interactive. It can not only be queried (like all texts can!), but it can be responsive to the query! This is a powerful change in how we classically navigate texts. (This also blurs the distinction between text navigation and teacher interaction.) This can be helpful in some ways, since we can use AI to converse about text, clear up confusion, or connect disparate ideas across texts.
- AI-based text navigation can potentially eliminate important cognitive tasks from the work of reading/interacting with texts: it is extremely good at summary, synthesis, and breaking down challenging ideas. Education theory implies that offloading this aspect of textual navigation *as a learner* is dangerous! It impacts the development of critical thinking. It also fails to directly expose the student to important patterns and ways of encountering/structuring information in text that might be critical to developing expertise in a given field.
- It can foster a misleading and false sense of accomplishment. Because AI synthesizes/summarizes and overall does a very good job of explaining texts, students interacting with it can come away from sessions thinking they understand content well when in fact their understanding is far more surface-level than they realize. The same thing happens in lectures: the lecture makes great sense in the moment. It isn't until you start the homework that you expose holes and misunderstandings in your initial conceptualizations. Finding and patching these holes to unify understanding is needed for true learning, and most homework assignments are designed with this purpose in mind.
- AI can hallucinate. While this is less of a concern in well-trodden undergraduate mathematical subjects, it can happen and cause confusion. Pairing AI with good fact-checking and sanity-checking habits (e.g., working through simple examples) can alleviate this issue. When AI gets something wrong, point it out. It often can correct the issue and supply some nuance about what went wrong, which can be fruitful for learning.

2. How does AI function as a subsitute for teacher interaction?

- **Positives:** AI has the ability to interact with a learner in a teacher or tutor-like mode. As such, it provides extensive additional resources in settings where the availability of teachers/tutors is limited. It can mimic modes of inquiry well and guide students through points of challenge and confusion using well-established teaching moves, including Socratic questioning, providing additional context and examples, and interactive cooperative modeling for how to think through challenges. However, to get an AI to behave this way requires some care in setting it up (see AI setup video), as it is not the "natural default mode" of most LLMs.
- **Potential Pitfalls:**
  - AI used this way is a *simulacrum* of human interaction. Human interaction is extremely complicated and involves many systems and evolutionary traits/physiological aspects of communication that are subtle. We do not know how certain aspects of human–human

interaction contribute to learning. We also do not know how well such aspects of the interaction are captured by the AI simulacrum.

- AI is not an expert educator that develops curriculum with your long-term growth/future career trajectory in mind. Good teachers understand that they are facilitating multiple long- and short-term goals that learners may only partially be aware of. This is because learners do not always know what skills are needed for long-term success in a given field. Even if they do know certain skills are needed, it is not always obvious how to build and improve upon those skills over long time frames. Teachers are meant to think about this issue. Most curricula are designed to challenge students and foster uncomfortable, but achievable growth in a range of ways. It is unclear how constant access to AI changes that. In the real world, access to human tutors and teachers is limited. This forces student independence, and it makes tutoring interactions strongly focused on major roadblocks. With a 24-hour, always-available tutor, over-reliance becomes an issue to worry about. Moreover, students may not know how to prompt or guide an AI in a way that addresses less obvious aspects of growth that a course is meant to foster.
- AI can hallucinate. It is convincingly an expert, even when it is sometimes wrong. See bullet point in question 1 above.

Strategies for AI Use in Math Education

With these issues in mind, here is a small list of some ways that AI can be or is commonly used by students. Each strategy is ranked according to our color spectrum.

Good / Safe Strategies

These strategies apply AI to challenge you in productive ways.

- **Use AI to create practice quizzes and problem sets.** Give AI lecture notes or text and have it produce practice problems for you to complete. You can also ask it for feedback on this practice work.
- **Use AI to create exit tickets.** You can ask AI to write you an "exit ticket" for a topic covered by a set of lecture notes or subsection of a textbook. This is like a mini-quiz with quickly-answerable, key questions about the lectures that you should be able to answer if you are comprehending the material. It is MUCH more valuable than a summary, and can be combined with a prompt for you to write a short summary yourself (which is a great way to digest lecture material). These exit tickets can be stored like a note-set. They collectively serve as a quick way to keep yourself on track, as well as a useful study resource for test preparation.
- **Flipped class.** Let the AI impersonate a student. Teach the student material that you should know well (e.g., to prepare for a final, or to solidify your understanding of a concept you think you've got). This forces you to do the cognitive work of identifying key information, putting concepts into your own words, and designing materials to test understanding.
- **Concept checker.** Give AI lecture notes or a textbook excerpt. Then, tell AI you want to explain these ideas in your own words. Share your explanation and have it give critical feedback on how well your ideas work, where they might be misleading or wrong, and how you can improve your descriptions. After a few rounds back-and-forth honing the explanation, you can have the AI produce a summary note set of your explanations.

Medium / "Be Careful" Strategies

Without some care, these strategies can be harmful.

- **Guided inquiry.** This is talking through a text/idea with AI, and asking it to teach you about the concept. This is probably the most common way LLMs are used by motivated students that are trying to learn material. It is not the most effective one because it offloads a lot of cognitive work to the AI, reducing the effectiveness of the learning process. *It is absolutely not the same mental workout one gets from reading, interpreting, synthesizing, and struggling with a text independently.*
- **Getting unstuck.** Another common use case. It's better to interact with humans (e.g., on Piazza, in office hours, in small study groups) when you can.

  However, if you do decide to turn to an LLM, get the most out of the experience by treating it like a visit to office hours. *Remember, the problem is probably designed to get you somewhat stuck. It's part of the learning the teacher is intending for you!* Tell the AI you don't want it to just solve the problem, but instead to help you get unstuck so that you can continue making progress by yourself. Tell it what you've done and where you feel stuck. Ask for a hint or a nudge in the right direction.
- **Testing/debugging code.** Often code bugs are simple, silly things — a good use for AI! If you have tried debugging yourself and can't find the issue, give it to AI and see if it can identify the problem. Guide it to focus on debugging, not commenting on other aspects of your code. The caution here is that debugging is a skill you want to develop, so try debugging yourself before turning to AI.

Bad / Danger Zone

These are all examples of cognitive offloading that can be more detrimental to your learning than it might seem.

- **Summarizing lecture notes and text for you.**
- **Solving your problems before you solve them yourself.** Don't generate an answer key until you are done and have tried hard to check your work with reasoning/sanity checks — this is part of problem solving!
- **Writing your code (or central subroutines) in full for you.** Later, when you are a professional working in a job, this can be a good use of AI because you will understand how to very effectively drive the LLM, when to intervene, how to break down the pieces and test components, modify the output, etc. But as a learner, it's best to code yourself.
- **Optimizing your code when it runs but is slow or otherwise unsatisfactory.** Part of the challenging work and problem solving process is the design of code, which involves some tinkering and investigation that should not be bypassed.
- **Using AI as a homework buddy that is constantly open and available every time you begin working.**
- **Using AI as a replacement for office hours and peer study groups.** Social interaction around problem solving is different from working with an AI. Why and how socialization affects learning outcomes is fuzzy and complicated, but it is well-acknowledged as an important aspect

of learning. It also gives you important practice with technical communication, teamwork, and social skills, which will be needed in almost every STEM-based career.