Learning

Learning should continue all the way through your project: mathematicians learn mathematics as long as they do mathematics. Learning consists of mastering theoretical material from books and papers and doing exercises. While some sources and exercises may be provided by your mentor or advisor, you may ask AI for more. For instance, you may use the following prompts:

1) "*What are the best sources to learn about Y?*"

2) "*What is known about the following question? (*formulate the question precisely*). Are there any references?*"

3) "*I have been reading B (*upload the file*) and on page N I have trouble understanding the proof of Theorem Z. Can you explain this proof to me in more detail?*"

4) "*Can you give me the basics about T and provide some exercises?*"

5) "*What is the motivation for the notion A?*"

6) "*How is P related to Q?*"

7) "*Can you explain the main ideas of R?*"

On the other hand, using AI to **solve** the exercises given by your mentor or advisor is **a bad idea**. The exercises are designed for you to gain technical skills needed for your research. Think of an athlete who wants to run a marathon, but for training uses a motorcycle!

In any case, in most learning environments trying to pass AI output as your exercise solution or being unable to explain your solution because you rewrote an AI solution without sufficient understanding constitutes **an inappropriate use of AI** and will be viewed as a **violation of academic integrity**.

For beginning researchers, a good practice is trying a problem yourself before asking AI to solve it, or asking AI for hints rather than full solutions first. In the end, the goal is not to use you and AI as a machine to churn out as many theorems as possible, but rather to make progress on the project while also giving you training and experience in doing research. If you entrust everything to AI, how will you develop your own research skills?

After all, discovering new mathematics is the most enjoyable part of this whole endeavor. So why outsource this pleasure to a soulless chatbot? Imagine that you go to a fine restaurant and instead of savoring the delicious dishes, put them through a chewing machine and then swallow the resulting goo. How much fun would this be?

Writing and proofreading

At the writing stage, you can use AI for helping you with LaTeX and with diagrams or pictures, although it can be hard to explain what you want drawn. Sometimes it is better to sketch it on paper by hand, scan with your cell phone, upload to AI and ask to generate a PNG file or a LaTeX code for the diagram.

You may also use AI for proofreading. You can upload your text and say something like: "*This is my paper. Read it very carefully and point out misprints, mistakes, mismatches of notation, English issues, etc. Read sections X.Y to Z.T. Give me comments.*" It's advisable to ask AI to read a small portion at a time (say, 10 pages), then the quality may be better. I've been using this method for my own texts and those of my students, and I strongly recommend it (e.g. ChatGPT 5.5 Extended thinking does a great job!). Of course, it does not replace good old human proofreading, but it's nice to do an AI proofreading (followed by corrections) before the human one!

AI models are generally good at catching misprints, correcting English, finding mismatched and undefined notation, but less good at finding subtle mathematical errors, although on occasion they can do that as well. However, if you want AI to check your paper seriously, you should ask it to check individual proofs, and keep adding details and clarifications until it no longer finds issues. But make sure you prompt the AI to be critical and not to please you (and make its job easy) by saying the proof is correct without thorough checking!

At the same time, copy-pasting significant chunks of unpolished AI output into your paper is **a bad idea**.[3] You should really do the writing yourself!

One more reason for this is that AI can tacitly copy something from someone else's work without reference. So if you copy-paste this output into your text and it gets publicized in that form, this may constitute **plagiarism**. And even though you'd do this inadvertently, you'd still be responsible for it!

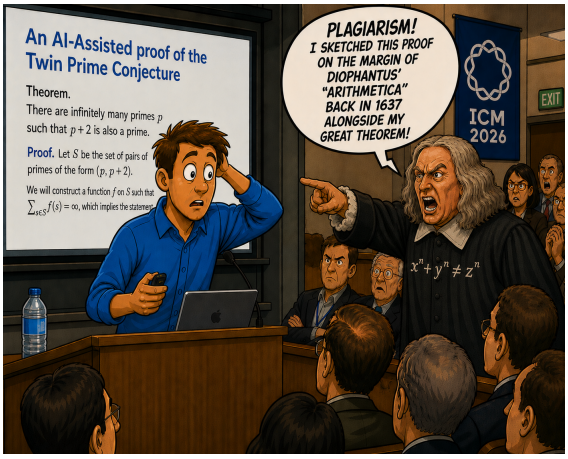

[3] Circulating unchecked AI output among a research group is acceptable, but only if its unverified status is clear to everyone involved.

For this reason, if you plan to use its output even in paraphrased form, it is a good practice to **always ask AI to list the sources it used** and then **independently verify them** (on arXiv.org, MathSciNet, etc.) and check if a reference to any of them is in order. Also it is a good idea to aggressively request the most related literature multiple times, and to **properly cite these original works in the reference section of your article**. For instance, it's easy to just cite Google in the reference section, but no one does this. It's far better for people to list the original resources that they are using, as this helps others understand the scope of the literature. You also might appreciate some credit for your ideas when they are being used!

Note that when asked to proofread, by default some LLMs will prefer to rewrite your text themselves, instead of giving you comments. This is for many reasons very undesirable, as explained above. So you should make it clear to the LLM that it should just give you a list of comments/corrections and not try to rewrite the text for you.