

STATEMENT ON AI USAGE FOR PHD STUDENTS

As you are probably aware, AI agents and LLMs are becoming both a valuable and potentially problematic tool in mathematics research. On one hand, they can be useful in many ways, e.g., for surfacing arguments and results from areas of mathematics we may not be familiar with, sometimes leading to solving problems in unexpected ways. On the other, not everything the current models produce is mathematically sound, and even when the content is correct, obtaining accurate references can be unreliable, which may lead to forms of unconscious plagiarism with respect to the currently accepted notion of authorship in the community. There are also broader ethical concerns of social, environmental, political, economic, and public health nature around AI usage and the AI industry, which I suggest considering carefully if you have not.

Beyond these general concerns, I believe AI poses particular risks at the PhD stage, when you are still developing the mathematical maturity needed to evaluate its outputs. That maturity is built precisely by wrestling with problems on your own, enduring confusion, pursuing dead ends, and finding your way out. There is no shortcut to this process that does not cost you something essential.

For these reasons, PhD students working under my supervision, or graduate students doing reading courses with me, will be required to disclose whether they are using AI in their studies and research, and if so, how. Please do so at our meetings or in any write up you send me.

I strongly recommend limiting the use of these tools at this stage in your career and instead focus on developing your mathematical skills through independent thinking and careful, slow, reading of the literature, even if it may take longer. In fact, I suggest thinking about a question or problem on your own before looking into the literature. The same principle applies to attending conferences, using online resources, or asking a colleague for a solution to an exercise or problem.

My general advice is to not use a result from any source, AI or otherwise, without understanding it fully. Do not settle for partial understanding; this is, in a sense, the essence of what research in this field demands. Current AI systems produce plausible-sounding mathematics very quickly, which makes it easy to absorb errors without noticing or settle for partial understanding in order to keep generating more text, some of which will be correct, but many of which will not. As you know from our discussions and my style, I will always ask probing questions to ensure that you have a thorough understanding of every claim and argument you invoke, down to first principles.

I hold myself to the standards set out in the Leiden Declaration on Artificial Intelligence and Mathematics (leidendeclaration.ai), which I recommend you read.

Manuel Rivera, June 2026