AI in mathematics research

17.8.2026

To paraphrase Terry Tao at the most recent ICM, we are in a mathematics revolution. And I take solace in the knowledge that there have been previous revolutions in science and mathematics. Proofs used to be withheld and used to answer problems in mathematical duals. Then the publishing industry took off. Computers revolutionized certain kinds of proofs, and made obsolete many important techniques (ask an incoming group of grad students if they know how to use a slide-rule). Then the internet shortened distances and connected the university system across the world. Just like mathematics itself, many of the issues that we are seeing at the moment have been faced before. At the same time we must all adapt to this new world or face being left behind and avoid becoming dependent and losing our hard earned skills. Walking that line is a dangerous game. Responsibly mentoring someone without those skills to develop what skills they need while avoiding those two traps is near impossible.

I have been using ChatGPT pro for 1 month now. I've used it to: clean up an existing preprint, my collaborator Jens Marklof used it to identify an error in a preprint and to help rewrite the paper, and I have used it to help develop proofs for two projects. The first three do not sound like bad applications of this new tool, but the two AI assisted proofs make me feel uncomfortable.

The first of these two projects concerns density estimates for exceptional eigenvalues of locally symmetric spaces (a sequence of numbers). In 2022 I was invited by Dennis Sullivan to give a colloquium at Stony Brook. He asked me a simple (to an expert in the field) question: the base eigenvalue (bottom of the sequence) is equal to the Hausdorff dimension of the limit set (a geometric quantity), does the any other exceptional eigenvalues have such a geometric interpretation? I've wrestled with this for a long time, and developed an answer: the convergence rate of the Hausdorff measure bounds the spectral gap. I plugged this in to ChatGPT and asked it to look for applications of this idea. It wasn't able to find any which gave an improvement to any known open question. However, I kept falling short of improving a bound of Sarnak and Xue. So I read their paper, and then I read the subsequent paper of Hee Oh establishing something similar, but leaving open an important conjecture. So I pointed ChatGPT at the conjecture and it spat out a proof. I completely rewrote the paper, and de-slopped the presentation and the result is now on my publications list.

For the second of these two projects, since working on my post-doctoral work with Alex Kontorovich, I have long thought that an approach from the world of partial differential equations (an area of math separate but overlapping my own) would have un-exploited consequences in my own area. I pushed this idea in a paper on horocycle equidistribution. So I explained it to ChatGPT and fed it the papers. I then asked if this approach – calculating the so-called fundamental solution to the pde formed by the ring of invariant differential operators – could be used to estimate certain important functions. ChatGPT then found that yes, it could prove a new estimate and improve on a well-known theorem of Peter Sarnak (one of the pioneers of this area of mathematics). It then found two additional applications to existing problems where it could fill in gaps or prove better bounds on a quantity.

I worry that I have not done enough to warrant authorship on these papers. At the same time, if I can prove these great new results, then shouldn't I do so and put out the proofs under my name. If I don't, someone else will and the community has agreed that with proper declarations the authorship falls on the mathematician interacting with, digesting, and reformulating what the AI

produced. Perhaps there should be an asterisk by these publications to declare AI use to distinguish it from my past unassisted work. Does it even matter anymore?

I foresee four possible pathways for mathematics at the moment, none are without their pain. The first is that the top-range AI models become prohibitively expensive or those companies go under and the gap is filled by cheap generic models which do not perform as well. However, with the exception of the Concord, it is hard to think of a technological advance which has had such a trajectory.

Another possibility is that we have reached (or gotten close to) the peak of the mathematical capability of LLMs, they are a tool which, when directed at well posed problems can be very effective, but need mathematicians to operate and digest. For mathematicians this will lead to one of two outcomes: we all use AI to develop our long-standing ideas and answer the questions it is capable of, then when the low hanging fruit has been picked, things reach a new equilibrium with AI-assisted mathematicians tackling a new class of problems. Or it becomes so easy to have our modern LLMs conduct research that our input becomes more observational and digestive.

The fourth possibility is that AI models develop such an advanced mathematical theory that our input becomes entirely obsolete. This seems least likely to me

At this stage, I cannot say for sure which pathway we are on. What we know for sure is that: At the moment, there are problems that AI cannot prove without some additional human guidance. There are proofs which AI has failed to find for ideas coming from humans. In its current form, it requires a decent amount of back and forth to develop an idea into something that works and this is only possible when starting from a decent idea developed through years of research in the area. I'm very grateful to have been educated before the era of AI and do be facing this revolution as more fully developed mathematician. But I do envy those at the end of their careers who can watch this revolution happen without being worried the robots will scoop your big idea.