

Ilia Binder (University of Toronto)

CMS President

In May 2026 an internal model at OpenAI produced a counterexample to a 1946 conjecture of Erdős on unit distances, a problem that had stood for eighty years. The result is real: a companion paper by external mathematicians checked it, and Gowers said he would recommend it for publication in the *Annals* without hesitation. Some of us were impressed and some alarmed; I was both. But it is worth being exact about what happened, because the details are the argument. The machine found the construction. Humans verified it, simplified it, and traced the ideas it rested on back to earlier work of Ellenberg, Venkatesh and others — work the model did not identify or credit; establishing that lineage remained a human scholarly task. And only months before, executives at the same company had claimed on social media that its system had solved ten of Erdős's problems, a claim withdrawn within days when it turned out to consist largely of rediscovering known results, as the mathematician who keeps that list pointed out. The triumph and the embarrassment teach the same lesson: the machine can now find real mathematics, and it cannot by itself tell you what is new, what is sound, or whose work it builds on.

A second example arrived in July, and it cuts a different way. Levent Alpöge announced a counterexample to the Jacobian conjecture, open since Keller posed it in 1939: an explicit polynomial map from $\mathbb{C}^3$ to itself whose Jacobian determinant is the constant -2 , and which nevertheless sends three distinct points to the same image. The conjecture is therefore false in dimension three and, by adjoining identity coordinates, in every dimension above it; the plane case remains open. The question came to him from Akhil Mathew, and the counterexample came out of a few hours of work with Anthropic's Claude Fable 5. Mathematicians checked the computation at once, and within days it had been formalized in a proof assistant and further consequences had been drawn from it.

Two features of this episode deserve our attention. The first is that anyone can check it: one computes a determinant and watches three points collide, and many mathematicians had done so within a day. A conjecture of eighty-seven years was settled in an afternoon, not because anyone trusted the machine, but because the object it produced was cheap to verify. That is a property worth wanting in a result, and worth asking for. The second is what checking does not give you. Written down, the map looks like a miracle. Terence Tao then reconstructed it: the map comes from multiplying polynomials together, and it fails to be injective because the same product can arise from different factorizations. Nothing miraculous remains — but the miracle disappeared only when a human explained it. The machine produced the example; a human supplied the reason. That division of labour may be the shape of much that is coming, and it is why the rest of this note dwells on understanding and not on certification alone.

The Leiden Declaration on Artificial Intelligence and Mathematics appeared on the second of June, endorsed within a day by the International Mathematical Union. I have signed it, and I ask the Canadian Mathematical Society to take it seriously.

Its central worry is not science fiction. Present-day systems produce arguments that look like proofs and are not: fluent, plausible, and wrong in ways that are hard to see. Our literature is a building in which every floor rests on the one below, and enough cheap, widely copied errors near the foundation make it unsafe to stand in. To this the Declaration adds two further concerns. Models reproduce human results without crediting the humans who found them. And proofs generated inside proprietary systems, which academics cannot open, would hand the direction of mathematics to companies rather than to mathematicians. The remedies it proposes are modest and correct: disclose the use of AI, keep publishing in peer-reviewed journals, and require that a result be checkable without proprietary access. A proof that only a paid model can verify is not a proof the community owns.

I want to be careful here, because it is easy to be misread. The Declaration is not a complaint against the machines, and neither is this note. Several of its authors build these systems for a living. Used well, an AI is a fast and tireless assistant, and I have no wish to send it back. The question of whether mathematicians will use these tools is already settled; we do. The questions that remain are practical: under what rules, and who writes them.

There is even a technology that meets the first worry head on. A proof written in a system like Lean and checked against its library, Mathlib, cannot hide a gap: the kernel accepts nothing on plausibility. To be precise, it guarantees the statement as encoded, not that we encoded the statement we meant — that check remains ours — but the merely plausible step, the kind that reads well and fails, is exactly what it refuses. It is not yet a general remedy. Most of mathematics is not formalized at all, and formalizing a serious theorem has until recently been slow work. That is changing fast, and not in a simple direction: the unit-distance disproof was formalized in Lean within a week of its announcement, at first leaning on two results not yet in the library, and within five weeks it had been formalized outright, with nothing assumed beyond the axioms — by another machine, which wrote more than a million lines of Lean in three weeks, against the two and a third million that nine years of human work put into Mathlib. Verification is becoming cheap. Whether it becomes ours or theirs is a separate question. A formal proof, moreover, can be correct and still tell you nothing about why it is true. What we need beside verification are tools that help a reader understand a proof and not merely certify it, and here the first steps exist. Patrick Massot's blueprints already turn a formalization into a human-readable document, with an automatically drawn map of how its parts depend on one another; they have served projects from Scholze's liquid tensor experiment to Tao's work on Freiman–Ruzsa. His current work aims at documents in which the reader chooses how much detail to see, from a one-line summary to the complete argument. That is a direction a society should encourage: not only an AI to be policed, but an interactive mathematics to be built.

Building that is only one of the society's tasks. The Declaration names funders, journals, institutions, societies, and for a national society these are not abstractions. The CMS publishes journals, runs meetings, awards prizes, sustains a competition pipeline, and speaks to government. It is exactly the place where a principle becomes a practice, or fails to. Signing a declaration costs nothing. Carrying it out is work, and the work is ours.

Consider teaching, which is most of what we do, and where the confusion is greatest.

Most of the students we teach will never become mathematicians. We teach calculus and linear algebra to engineers, to economists, to biologists. For years we half-justified these courses by telling students they would one day need to compute the integrals themselves. That was never quite true, and it is now plainly false: for the standard integrals we assign, freely available software is faster than our students will ever be, and usually more reliable. What is left is the reason that was always the real one. A person who cannot reason cannot check the

machine, and a person who cannot check the machine is at its mercy. The service role of mathematics has not weakened. It has lost its excuse and kept its purpose, which makes it more important than before, not less.

Training mathematicians is harder, and here I have less confidence and more worry. The temptation is to answer with a list of tools — learn this proof assistant, learn to write that prompt — and to mistake the list for a plan. Tools change every few months; the mathematician we are trying to form takes years. What must be protected is the very thing the tools most threaten: the long, unpleasant, private struggle with a problem that will not yield. Mathematical maturity is not handed over in lectures. It is built in exactly those hours, and a student who gives them to a machine has bought a grade and sold an education.

Problem solving deserves its own word here, because it is the training ground for all of this. Lectures transmit definitions and theorems; only problems teach students how to begin when no one has shown them what to do. Here my background speaks: I come from the Russian tradition of mathematical circles and olympiads, in which solving problems was not a preparation for mathematics but its first form, and I have seen what that tradition produces. The useful distinction is between an exercise and a problem. An exercise asks students to repeat what they have been shown, and machines now do exercises instantly. This changes their role, not their worth: done honestly, exercises still build fluency, as scales do for a musician, but they can no longer carry the weight of mathematical formation, and a submitted exercise no longer tells us whether the student could have done it unaided. A problem gives them nothing but the question, and working through it builds the judgment, the taste and the stamina that nothing else builds. If we mean what we say about protecting the struggle, then problem seminars and circles belong at the centre of our programs, not at their margins.

There is also the harder matter of keeping them at it. Students who watch a machine dispose in seconds of a problem that cost them a week may be forgiven for asking why they should bother. We owe them an honest answer, and we have one: the reward of mathematics was never that no one else could solve the problem; it is understanding, and understanding is not something a machine can have in your place. Our task is to let students feel that difference early, before discouragement hardens into departure — and not only undergraduates, since the doctoral student wondering what a thesis is for deserves the same answer. Motivation is not a soft concern; without it, nothing else we plan will matter.

And there is the calibration, which is genuinely new. Students who never touch these systems will be helpless before long; students who reach for them at the first sign of difficulty risk never building anything of their own. What we must teach is timing: struggle first and consult the machine after; use it to check and to explore, not to think in your place; notice the moment when it stops helping you think and starts thinking instead of you. They should also know that these systems are trained to be accommodating. When Tao published the record of his own session with a model while working through the Jacobian example, readers were struck by how steadily it assured him that he was right. A tool that approves of every step may harm a student's judgment more than one that occasionally errs: an error invites checking, while agreement passes unnoticed, because it is what we wished to hear. Overuse does not produce a weaker mathematician. It produces none. The Declaration says the goal of the subject is human understanding, and in the classroom that is not a slogan: it is the struggle we must refuse to automate away, even as we teach students to reach for the tool once they have done the work that makes it a tool rather than a substitute. These two demands pull against each other, and I see no clean way to satisfy both. The society will not resolve this tension. It can at least keep us honest about it.

Some of this the CMS can act on now. The most important step is also the simplest and the one most fully in our hands: require, in the Canadian Journal of Mathematics and the Society's other journals, a short statement of which AI tools an author used and where. Disclosure should be proportionate — substantive use in discovery, proof, computation or exposition is what matters, not spelling correction — and the policy should state what the Declaration takes for granted: a system cannot be an author, and responsibility for a paper rests with the humans who sign it. It costs us nothing and enacts the Declaration directly. Next, and more urgent than it sounds, we should give referees written guidance on both halves of one problem: not to upload manuscripts under review to external AI services, a confidentiality risk that journals are already confronting; and to ensure that the mathematical judgment expressed in a report is their own — help with language is one thing, delegated assessment is another. We should make it editorial policy that a published result be verifiable without proprietary access. And the Society should enter a conversation it has so far only watched: what consent and attribution ought to mean when our journals and archives become training data for proprietary systems.

In teaching our influence is only advisory, since curricula belong to departments, but nowhere are the stakes higher, and the remedy is not the same in every classroom. In proof-based courses we can move the weight back into the room: oral examination, proofs at the board, students defending their own solutions. In the large service courses this does not scale. I have taught calculus to eight hundred, and I will not be oral-examining eight hundred; there the honest answer is proctored, in-class assessment, or a frank admission that we do not yet have a clean one. A problem set done at home is now a test of the student's subscription, not of the student — it is no longer evidence of unaided work — and to pretend otherwise only teaches students that we are not paying attention. Competition mathematics, the olympiad pipeline the Society already supports, trains exactly the reasoning that has gained value, and should be funded as such.

Advocacy is the slowest of these, and here I want to be a Canadian about it. It is easy to write that the danger is dependence on three or four American companies, and there is something to that. But Canada has companies and institutions of its own: Cohere is Canadian, and Mila, the Vector Institute and CIFAR are a good part of why this country counts in the field at all. We are not bystanders to this technology, and we have some leverage over its direction; our sharpest problem is not a foreign menace but the quiet export of our best students to laboratories that pay in multiples of a professor's salary. The Society should press NSERC and the government on both fronts at once: to keep funding the training of human mathematicians, and to keep research infrastructure open and verifiable, so that the direction of the subject stays with the mathematical community rather than drifting to whoever owns the largest machines. One more point, small only in appearance: whatever policies the Society adopts must exist in both official languages, and we should notice that the models themselves are not equally capable in both. If these tools are becoming part of how mathematics is learned and done, a gap between their capabilities in French and in English is a gap in access for our francophone students and departments, and it belongs on the same list of inequities the Society already watches.

In the spirit of the Declaration I have just endorsed, I should make a disclosure: I used an AI system to criticize and edit drafts of this note; its argument and conclusions are my own.

I have no illusions about the shelf life of these pages. Given the speed of the last two years, the specifics may be stale two months after they are printed, and the whole note may read as naive in a few. But the particulars — the tools, the companies, the fixes — were always going to date. The principles will not: disclose what the machine did, keep proofs verifiable by anyone, protect the human understanding that is the point of the whole enterprise. Those I ask the Society to adopt now, together with a date for revisiting the rest. And when the revision comes and these pages look quaint, that too will be mine.

Copyright & Permissions

The Canadian Mathematical Society grants permission to individual readers of this publication to copy articles for their own personal use. Use for any other purpose is strictly prohibited. To obtain a license for anything other than copying articles for personal use, please contact the Canadian Mathematical Society to request permissions or licensing terms.

Canadian Mathematical Society — 616 Cooper St., Ottawa, ON K1R 5J2, Canada