The Atlantic

TECHNOLOGY

OpenAI Just Carpet-Bombed Mathematics

If mathematicians have no reason to pursue the mysteries of their field, the rest of us will stumble into a less certain world.

By Konstantin Kakaes

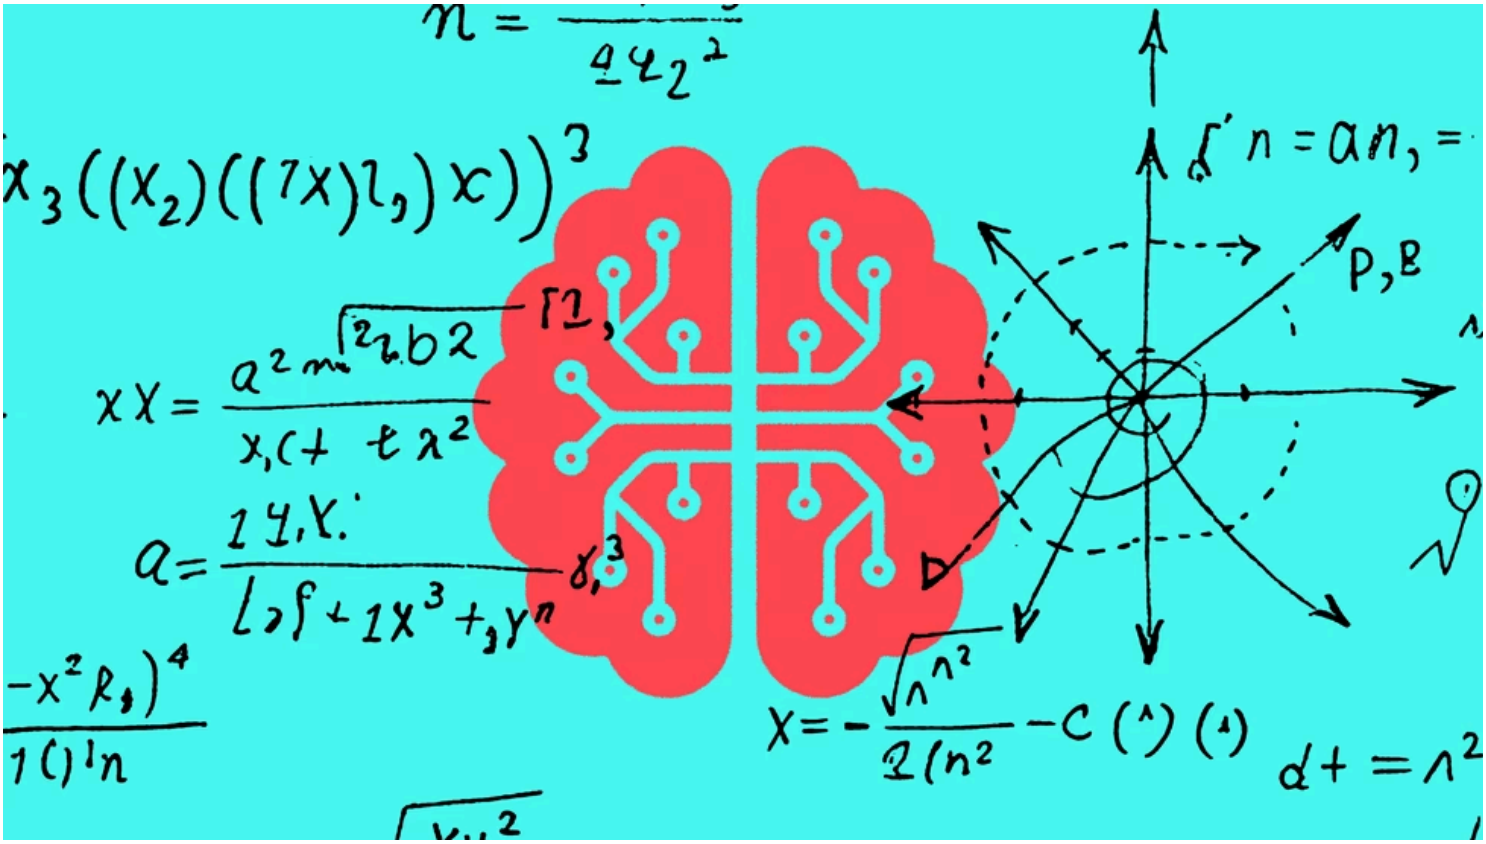

Illustration by The Atlantic. Source: Getty.

OCTOBER 8, 2026

SHARE AS GIFT DISCUSS 97 SAVE

Updated on October 10, 2026 at 6:50 p.m. ET

In the growing chorus of mathematicians revolting against AI companies, Hugo Duminil-Copin, who in 2022 received the Fields Medal, math's highest honor, has perhaps put mathematicians' concerns most vividly. In an essay published just before OpenAI announced last month that it had solved one of seven grand mathematical challenges known as the Millennium Prize Problems, he wrote that AI proofs "nuke the mathematical landscape, making it increasingly difficult to inhabit after each blast."

If resolving a Millennium Prize Problem was like dropping an atomic bomb, then the results that OpenAI shared on Tuesday were a mathematical carpet-bombing. Mathematicians had been bracing for the assault: OpenAI had teased that a new model had found solutions to "more than 100 long-standing open problems across most areas of mathematics." When the drop came, it included 377 new proofs by an as-yet-unnamed proprietary model.

About a third are major proofs of conjectures that are famous across mathematics. Another third are significant within their mathematical subfield. The remaining third are either more minor results or variants. Hundreds of people had spent dozens of years trying to understand—and prove—the results that OpenAI released in a single day.

Mathematicians are still picking through the rubble. It is too early for anyone to have verified the OpenAI work in detail, but several mathematicians have told me that at least some of the papers appear not to prove what they claim to prove. (Every carpet-bombing has a few munitions that fail to explode.) OpenAI has already withdrawn three of its papers; still, "we expect mistakes to be quite rare," an OpenAI spokesperson told me in an email. The company also said, as part of the initial announcement, that the model had attempted to solve about 3,600 further problems for which the company did not release successful proofs. OpenAI promised to fund the study of AI-produced results through conferences and special programs. In mathematics, such programs can last days or even months and cover not just travel costs but full salaries, although OpenAI has yet to specify how generous the funding will be.

"We don't believe the future of mathematics is set. We're working with the math community to navigate the future collaboratively," the spokesperson wrote. Nevertheless, many mathematicians are looking at that future with apprehension. They worry that having AI leap to solve proofs will sacrifice what the struggle to get there might have yielded. They worry especially about how they will train younger mathematicians to find that spark of life, or wanting to at all. And they worry that, if

mathematicians have no reason to pursue the mysteries of their field, then the rest of us will stumble into a world run by machines whose workings nobody understands.

Pure mathematicians spend their time manipulating abstract structures just for the sake of it—they tend not to be motivated by applications. But the possibility that a mathematical idea will profoundly change the world is always there, lingering in the background. As a mathematician who did top-secret research for the National Security Agency once told me, the real divide is not between pure and applied math, but between "applied math and math that hasn't been applied yet." The technologies that enable daily life—the mobile telephones sending and receiving billions of bits of information every second, the electric grid transmitting trillions of watts of energy from power plants to end users each day, sewage systems that flow fast enough to not clog, airplane wings that flex 25 feet without breaking, ultrasound-imagery machines that synthesize infant heartbeats from high-frequency sound waves—all rely on mathematical language.

That language has been forged through millennia of social interactions. Mathematicians collaborate. They compete. They run into dead ends. They periodically rewrite what is known in new, more compelling ways. All of this takes time, and by short-circuiting this process, AI is imperiling centuries-old practices for creating knowledge and training mathematicians. It is putting human fluency in that language at risk.

AI has already undercut one motive for mathematical exploration. "If you had asked mathematicians five years ago what was the most important thing they do, probably it would be that we write proofs and our proofs are correct," Mohammed Abouzaid, a professor of mathematics at Stanford University and the executive director of First Proof, an effort to benchmark AI math abilities, told me. But with AI tools, the value of proofs has already diminished enough that writing them is no longer mathematicians' most important activity.

Proofs have long been central because they distinguish math from other arts and sciences. A proof starts from self-evident truths: If two sets have the same elements, for example, then they are the same. From such simple beginnings, proceeding by small and large logical steps, humanity has built an enormous, durable infrastructure of certainty.

As the building blocks of that infrastructure, proofs have long been the currency of professional recognition. Prove something important, and you'll get a faculty job and tenure; prove something really important, and you'll get a big prize and a place in history. But if proofs come too easy or too fast, they can choke off progress. As Duminil-Copin wrote, his failed attempts to prove a prominent conjecture in

probability theory "generated dozens of ideas that I later repurposed in other contexts, leading to discoveries I would never have imagined making."

In a prescient 1994 essay, William Thurston, one of the most important mathematicians of the 20th century, argued that the field's emphasis on proofs backfires. Thurston was himself terrific at proving things and lamented the fact that mathematicians like him got too much credit. He voiced regret that his rapid proof of a number of theorems had led to the abandonment of a subject called "foliations," because nobody wanted to compete with him. But when people stopped working in the area, they forgot what they had learned. "Mathematical knowledge and understanding were embedded in the minds and in the social fabric of the community of people thinking about a particular topic," Thurston wrote. "This knowledge was supported by written documents, but the written documents were not really primary." You don't need to understand what foliations are to appreciate his point that mathematical knowledge is impermanent.

In order to consolidate and build knowledge and understanding, Thurston argued, mathematicians should be clear that their aim is to produce human understanding. But in the 30 years since his essay was published, his call to give greater value to things other than proofs has been largely ignored. AI is now arguably better than any living person at coming up with proofs, doing for all of math what Thurston did to foliations. So mathematicians are being forced to reckon with his concerns.

Plenty of working mathematicians are excited at the speed that AI affords them, and the ways they can use it to augment their abilities. They can continue to work on the 90 percent of problems that seem to have proved too hard for OpenAI's new model, and some see OpenAI's 377 proofs not as a carpet-bombing but as a treasure chest filled with novel gems of mathematical knowledge.

But this won't necessarily help their successors. Traditionally, an adviser would assign a smart graduate student a problem they estimated might take a few months of work to resolve. Solving such bite-size research questions trained students to tackle deeper questions. AI has turned that practice quaint, but finding those hard solutions is what builds the mental muscles that make exceptionally abstract mathematical thinking possible. You can't learn math just by reading existing math, in the same way that you can't learn to play the violin just by reading about music. You have to put in the mental and tactile work.

Exceptionally disciplined students can still learn by simply refusing to use modern AI models, but the technology leaves most math students, even very talented math students, in a difficult spot. Daniel Litt, a mathematician at the University of Toronto, has suggested math departments might create rooms where students can

spend 20 to 30 hours a week without internet access or AI, to give them a chance at learning the skills they need. Whether or not that's realistic, we will have to adapt to a world in which the value of human thought is diminished, and try to create social spaces where it still matters.

Lauren Williams, a math professor at Harvard University and another of the organizers of First Proof, told me that she worries about students accepting a computer's answer without engaging with the material themselves. Students are demoralized, she wrote later in an email: "OpenAI has produced manuscripts claiming results that have scooped the research of many young people, such as graduate students and postdocs." If tenured faculty get scooped, they've still got job security; people earlier in their career might not have the time to start all over again.

The OpenAI spokesperson noted that "many at OpenAI believe that the senior members of the community have a responsibility to rapidly evolve the field to adapt to the new reality that AI brings." In other words, the company believes that mathematicians need to get on board with its vision of the field's future or get out of the way. But the company has also responded, in some ways, to mathematicians' critiques. Last month, it convened an independent advisory group of eight extremely accomplished mathematicians and Ed Witten, arguably the world's most influential physicist. On September 29, the advisory group issued its first report: It asked frontier-AI labs to stop testing "advanced mathematical problems on proprietary models that remain inaccessible to the broader scientific community." Exactly one week later, OpenAI snubbed this instruction and released its proofs, obtained by just such a proprietary model. The company followed some of the advisory group's recommendations: The papers accompanying the proofs were uploaded to GitHub rather than hosted online by OpenAI itself, to leave an auditable paper trail, and 10 of the proofs were accompanied by an explanation of how the AI had arrived at its conclusions.

The OpenAI papers are markedly clearer than earlier generations of AI-written proofs, but on the whole, they are not well written. (Only one of the 722 manuscripts released has been edited by a human, according to the company.) "The notation and language is unnatural, the reasoning is unclear; they are not the sorts of manuscripts that a self-respecting mathematician would release into the world," Williams wrote. "Normally when we prove a result, we spend weeks if not months or years working hard to come up with the clearest and cleanest explanation of how all the pieces fit together, so as to genuinely contribute to human understanding. That is not what has happened here." The OpenAI spokesperson said that the company is "working to improve future models' writing ability."

Math is a deeply social undertaking. Mathematicians gather for weeks at summer workshops. They work informally, collaboratively, and late into the night, across borders and time zones. The result has been a shared human understanding, which

underpins not only a numinous world of rigorous abstractions but also the technologies that make it possible for 8 billion people to live on Earth. Right now, mathematicians can still explain how and why math works as it does. If they cannot figure out how to pass that knowledge on, it will affect not only mathematicians but all of us who rely on the fruits of their intellect.

This article has been updated to clarify the nature of OpenAI's promised mathematics funding.

View Discussion 97

ABOUT THE AUTHOR

Konstantin Kakaes

Konstantin Kakaes is a journalist in Washington, D.C., who writes "Transformation," a newsletter about math and AI, for *Quanta Magazine*. He is the author of *The Pioneer Detectives*. He has written about science and technology for *The Economist*, *The Wall Street Journal*, and *The Washington Post*.

Explore More Topics

OpenAI

About

Contact

Podcasts

Subscription

Follow

Privacy Policy Your Privacy Choices Advertising Guidelines Terms & Conditions Terms of Sale Responsible Disclosure Site Map

TheAtlantic.com © 2026 The Atlantic Monthly Group. All Rights Reserved.
This site is protected by reCAPTCHA and the Google Privacy Policy and Terms of Service apply