

The Future of Mathematics

Jeremy Avigad

Carnegie Mellon University

October 9, 2026

“Mathematics underwent, in the nineteenth century, a transformation so profound that it is not too much to call it a second birth of the subject—its first birth having occurred among the ancient Greeks...”

Howard Stein, “Logos, Logic, and Logistiké: Some Philosophical Remarks on Nineteenth-Century Transformation of Mathematics”

“The report of my death was an exaggeration.”

Mark Twain

“May you live in interesting times.”

(traditional)

Overview

Advances in AI for mathematics over the last several months have been stunning.

I will discuss:

- AI for Mathematics
- Recent Events
- Mathematicians' Responses
- The Future of Mathematics
- The Role of Research in AI


AI for Mathematics

AI for Mathematics

In 2023, a workshop by the National Academies for Science, Engineering, and Medicine explored the promise and challenges of AI for mathematics.

The NSF Institute for Computer-Aided Reasoning was launched in September 2025 to address the challenges.

 **Institute for Computer-Aided Reasoning in Mathematics**




AI for Mathematics

We interpret the phrase broadly to include:

- interactive theorem proving and formalization
- automated reasoning and symbolic AI
- machine learning and neural AI

All three come together in neurosymbolic theorem proving.



AI for Mathematics

Formalization can be used to:


- verify correctness
- build reference libraries
- search for results
- support collaboration
- teach mathematics
- explore and discover new mathematics
- support uses of automated reasoning and AI

```
theorem riesz_lemma {F : Subspace  $\mathbb{K}$  E} (hFc : IsClosed F) (hr : r < 1) :  $\exists x_0 : E, x_0 \notin F \wedge \forall y \in F, r * \|x_0 - y\| \leq \|x_0\|$   
classical  
obtain ⟨x, hx⟩ :  $\exists x : E, x \notin F := hF$   
let d := Metric.infDist x F  
have hFn : (F : Set E).Nonempty := ⟨_, F.zero_mnorm⟩  
have hdp :  $0 < d :=$   
  lt_of_le_of_ne Metric.infDist_nonneg fun heq : d = 0 => hx ((hFc.mem_iff_infDist_zero hFn).2 heq.symm)  
let r' := max r 2-1  
have hr' :  $r' < 1 :=$  by  
  simp only [r', max_lt_iff, hr, true_and]  
  norm_num  
have hlt :  $0 < r' :=$  lt_of_lt_of_le (by simp) (hr')  
have hdl :  $d < d / r' :=$  (lt_div_iff hlt).mpr  
obtain ⟨y₀, hy₀F, hxy₀⟩ :  $\exists y \in F, \text{dist } x y < d :=$  hFn  
have x_ne_y₀ :  $x - y₀ \notin F :=$  by  
  by_contra h  
  have :  $x - y₀ + y₀ \in F :=$  F.add_mem h hy₀F  
  simp only [neg_add_cancel_right, sub_eq_add_neg]  
  exact hx this  
refine ⟨x - y₀, x_ne_y₀, fun y hy => le_of_lt ?_⟩  
have hy₀y :  $y₀ + y \in F :=$  F.add_mem hy₀F hy
```


AI for Mathematics

Symbolic methods include:

- SAT solvers
- First-order theorem provers
- Constraint solvers
- SMT solvers
- Computer-algebra systems



Formal Verification of the Empty Hexagon Number





AI for Mathematics

Machine learning includes:

- Data mining and pattern detection
- Reinforcement learning
- Evolutionary search
- PINNs
- Use of coding agents
- Formal and informal proving agents

$X(z)$: Unlabelled Bruhat interval




AI for Mathematics

Neurosymbolic theorem proving brings together LLMs, formalization, and symbolic automation.

One can ask for ordinary (informal) proofs, formal proofs, or both.

Google DeepMind's AlphaProof and AlphaGeometry performed at the level of a silver medalist on problems from the 2024 International Mathematics Olympiad.



Ancient History

As recently as January, I was saying the following in talks:

- AI has not had a substantial impact on mathematics.
- Nonetheless, there have been some striking results.
- These show that the technology is promising.
- AI will have an impact on mathematics in the long term.
- Mathematicians should learn about the technologies and play a role in their development.


Moral: Be careful what you ask for.

Recent Events

Recent Events

On February 5, 2026, eleven mathematicians posed benchmark research problems as part of the *First Proof* project.


By the time the solutions were posted on February 13, Google DeepMind and OpenAI had claimed surprisingly good performance using internal models.



The screenshot shows the 'FIRST PROOF' website. At the top left is the logo, which consists of a grid of circles with some filled and some empty, next to the text 'FIRST PROOF'. Below the logo is a navigation bar with links: 'Home', 'Third Batch', 'Second Batch', and 'Past' (which is underlined and has a dropdown arrow). The main content area has a heading 'FIRST BATCH' with a yellow dot icon. Below this is a paragraph: 'In baking, the *first proof*, or bulk fermentation process, is a crucial step in which one lets the entire batch of dough ferment as one mass, before dividing and shaping it into loaves.' This is followed by another paragraph: 'This project represents our preliminary efforts to develop an objective and realistic methodology for assessing the capabilities of AI systems to autonomously solve research-level math questions. After letting these ideas ferment in the community, we hoped to produce a more structured benchmark.' The final paragraph states: 'We presented a diverse set of 10 research-level math questions, drawn from algebraic combinatorics, spectral graph theory, algebraic topology, stochastic analysis, symplectic geometry, representation theory, lattices in Lie groups, tensor analysis, and numerical linear algebra. Each question arose naturally in the research process of the authors and has been answered with a proof of roughly five pages or less.'

Recent Events

On February 23, an AI startup, Math Inc., announced that its proprietary proving agent had completed the formalization of a proof of the optimality of the E_8 sphere packing, building on a blueprint, library, and partial formalization by Hariharan et al.



The screenshot shows the Math, Inc. website. The top section has a green header with the title "Formalising Sphere Packing in Lean" and the authors "by Chris Birkbeck, Sidharth Hariharan, Seewoo Lee, Gareth Ma, Bhavik Mehta, Maryna Viazovska". Below the header are buttons for "Blueprint (web)", "Blueprint", "Library", and "Partial". The main content area is white and features the title "Formalising Sphere P" in green. Below the title is a paragraph: "In 2022, Maryna Viazovska was awarded the Fields Medal for her work on sphere packing in dimension 8. This project forms part of a larger effort following her [original paper](#) and [followup](#) collaboration currently being led by Chris Birkbeck, Sidharth Hariharan, Seewoo Lee, Gareth Ma, Bhavik Mehta and Seewoo Lee. If you'd like to contribute, please contact us." Below the paragraph is a list of links: "Zulip chat for coordination", "Blueprint", "Blueprint as pdf", "Dependency graph", "Doc pages for this repository", "Project dashboard", and "Viazovska's original paper". The right side of the screenshot shows a dark blue sidebar with the Math, Inc. logo and a "Back Home" link. The main content area on the right is titled "Completing the formal proof of higher-dimensional sphere packing" and contains a paragraph: "Using Gauss, we have helped formally verify the sphere packing problem in dimensions 8 and 24 — certifying that the E_8 lattice and the Leech lattice achieve the densest possible arrangements of non-overlapping spheres in their respective dimensions." Below this paragraph is another paragraph: "These results, originally proved by Maryna Viazovska and collaborators, earned Viazovska the Fields Medal at the 2022 International Congress of Mathematicians. This is the only formalization of a Fields Medal-winning result from this century ([GitHub](#))." The bottom section is titled "The Problem" and contains a paragraph: "How densely can identical spheres be packed in n -dimensional space? In one dimension, the problem is trivial, and in two dimensions, elementary proofs have existed for decades. In three dimensions, a solution was conjectured by Kepler in 1611, but it was not proved until Thomas Hales settled it in 1998, in a heavily computer-assisted proof that subsequently took over a decade to formally verify."

Formalising Sphere Packing in Lean
by Chris Birkbeck, Sidharth Hariharan, Seewoo Lee, Gareth Ma, Bhavik Mehta, Maryna Viazovska

Blueprint (web) Blueprint Library Partial

Formalising Sphere P

In 2022, Maryna Viazovska was awarded the Fields Medal for her work on sphere packing in dimension 8. This project forms part of a larger effort following her [original paper](#) and [followup](#) collaboration currently being led by Chris Birkbeck, Sidharth Hariharan, Seewoo Lee, Gareth Ma, Bhavik Mehta and Seewoo Lee. If you'd like to contribute, please contact us.

- [Zulip chat](#) for coordination
- [Blueprint](#)
- [Blueprint as pdf](#)
- [Dependency graph](#)
- [Doc pages for this repository](#)
- [Project dashboard](#)
- [Viazovska's original paper](#)

Math, Inc.

[← Back Home](#)

Completing the formal proof of higher-dimensional sphere packing

Using Gauss, we have helped formally verify the sphere packing problem in dimensions 8 and 24 — certifying that the E_8 lattice and the Leech lattice achieve the densest possible arrangements of non-overlapping spheres in their respective dimensions.

These results, originally proved by Maryna Viazovska and collaborators, earned Viazovska the Fields Medal at the 2022 International Congress of Mathematicians. This is the only formalization of a Fields Medal-winning result from this century ([GitHub](#)).


The Problem

How densely can identical spheres be packed in n -dimensional space? In one dimension, the problem is trivial, and in two dimensions, elementary proofs have existed for decades. In three dimensions, a solution was conjectured by Kepler in 1611, but it was not proved until Thomas Hales settled it in 1998, in a heavily computer-assisted proof that subsequently took over a decade to formally verify.

Recent Events

On May 20, OpenAI announced that an internal model had disproved the Unit Distance Conjecture, settling a longstanding problem in combinatorics.

The proof was subsequently formalized in Lean, again using AI.



Recent Events

On June 2, a group of authors released the Leiden Declaration, building on discussions from a meeting on the Mechanization and Mathematical Research at the Lorentz Center in September 2025.

Leiden Declaration on Artificial Intelligence and Mathematics

This declaration calls for action to address the challenges posed by the use of artificial intelligence within mathematics research. It is the result of a community initiative and is [endorsed](#) by the International Mathematical Union (IMU).

[Sign this Declaration](#)[Read Declaration](#)[Endorsements](#)[View signatories](#)[News](#)

Share the declaration



DECLARATION TEXT

Leiden Declaration on Artificial Intelligence and Mathematics

2 June 2026 · [DOI: 10.5281/zenodo.20302944](https://doi.org/10.5281/zenodo.20302944)

Preamble

Technological developments have repeatedly transformed the practice of mathematics. Recent artificial intelligence technologies, including symbolic and neural methods for the generation and formalization of mathematics, may already have initiated a significant chapter in this long history. Among researchers, artificial intelligence has produced a wide range of reactions: enthusiasm for its potential to yield new discoveries; intimidation by the pace of developments; indifference to these rapid changes; and concern for the implications,

Recent Events

More AI-generated results appeared, including

- resolutions of Gromov's conjecture on sofic groups,
- the Jacobian conjecture in three or more dimensions,
- and the Hopf conjecture on complex structures on S^6 ,

more often than not, accompanied by Lean formalizations.

August 1, 2026 Publication

Ten advances in mathematics and theoretical computer science

[Read the paper ↗](#) [Read the reasoning walkthroughs ↗](#)

$\delta \leq \delta_{opt}$

R^d

PROBLEM HIGH-DIMENSIONAL SPHERE PACKING NUMBER 01/10 FIELD DISCRETE GEOMETRY

Recent Events

On July 19, Anthropic's Fable disproved the Jacobean Conjecture for three dimensions and higher, during the World Cup final.

 **Post**

 **levent** 
@__alpoge__

Follow  

hello there the jacobian conjecture is false thanx to my close friend akhil for asking about it and my other close friend fable for working during the world cup final

$$((1+xy)^3 z + y^2 (1+xy) (4+3xy), y + 3 x (1+xy)^2 z + 3 x y^2 (4+3xy), 2 x - 3 x^2 y - x^3 z): \mathbb{C}^3 \rightarrow \mathbb{C}^3$$
, has jacobian determinant -2 , and sends $(0, 0, -1/4)$, $(1, -3/2, 13/2)$, and $(-1, 3/2, 13/2)$ to $(-1/4, 0, 0)$

7:19 PM · Jul 19, 2026 · **40M** Views

 1.7K  8.5K  43K  11K 

Relevant  View quotes >

Recent Events

On July 23, Jacob Tsimerman was one of four mathematicians to be awarded a Fields Medal at this year's International Congress of Mathematicians.

The same day, the New York Times announced that he was leaving mathematics to work on AI safety.

Jacob Tsimerman Lives in Number Theory

Dr. Tsimerman won one of four Fields Medals, an award for top mathematicians under 40, for his work on the André-Oort conjecture, but is now changing his focus to artificial intelligence.

Listen · 4:47 min

Share full article

7



Jacob Tsimerman, a professor of mathematics at the University of Toronto, is one of this year's recipients of the Fields Medals, which recognizes top mathematicians under the age of 40. Simons Foundation

Recent Events

The next day, Terence Tao gave a plenary lecture, “Mathematics in the Age of AI.”

There was a panel on the topic one day later, moderated by Emily Reihl.

Mathematics in the age of AI

Public lecture, International Congress of Mathematicians 2026

Terence Tao


University of California, Los Angeles

July 24, 2026

Recent Events

On September 4, Anthropic announced that Claude had formalized a proof of Fermat's Last Theorem over a period of 11 days.

Kevin Buzzard's project to formalize it (taking some background theorems as axioms) is still ongoing.



Recent Events

On September 8, OpenAI constructed a solution to the Navier-Stokes equations that exhibits finite-time blowup.

This (apparently) solves a Clay Millennium Prize problem.

September 8, 2026 Research Publication

On the Navier–Stokes Millennium Prize Problem

[Read the paper](#)[Link to Lean formalized proof >](#)[Listen to article](#)

9:01

[Share](#)

We're sharing a solution to the Navier–Stokes existence and smoothness problem, one of the Millennium Prize Problems. This proof, produced by an internal OpenAI system, shows that the dynamics of the Navier-Stokes equations for fluid motion can develop a singularity in finite time. We're sharing both a writeup of the proof and a formalization in Lean.

Recent Events

Tristan Buckmaster, an expert in PDEs, had been making steady progress using the same strategy.

He worked with Levent Alpöge (Anthropic) and various AI models.

The incident raised several concerns.

Buckmaster was careful to credit Diego Córdoba and Luis Martínez-Zoroa for the strategy.

An N.Y.U. Mathematician Clashed With OpenAI Over a \$1 Million Proof

Tristan Buckmaster was on the path toward an important proof when one of the A.I. giants used its staggering resources to get there first.

Listen · 11:58 min

Share full article



594




Recent Events

On October 6, Open AI released a repository with 722 manuscripts for 372 problems, 235 of which have a Lean formalization (possibly partial).

Claimed results include

- the quasi-Riemann hypothesis,
- Hilbert's tenth problem over $\mathbb{Q}$, and
- Rokhlin's problem on multiple mixing.



Recent Events

The mathematics community is still digesting the results.

Kevin Buzzard recently analyzed reactions to AI in terms of the five stages of grief.

Still, there are also glimmers of optimism and enthusiasm.

Nobody seems to doubt that this is a watershed moment for mathematics.

Xena

```
-- part (a) -- delete one of these and prove the other.

theorem M1F_Sheet1_Q1a_is_T :  $\forall x : \mathbb{R}, x^2 - 3x + 2 = 0 \rightarrow x = 1$  := sorry
theorem M1F_Sheet1_Q1a_is_F :  $\neg (\forall x : \mathbb{R}, x^2 - 3x + 2 = 0 \rightarrow x = 1)$  := sorry

-- part (b) etc etc.

theorem M1F_Sheet1_Q1b_is_T :  $\forall x : \mathbb{R}, x = 1 \rightarrow x^2 - 3x + 2 = 0$  := sorry
theorem M1F_Sheet1_Q1b_is_F :  $\neg (\forall x : \mathbb{R}, x = 1 \rightarrow x^2 - 3x + 2 = 0)$  := sorry
```

[Home](#) [About Xena](#) [Student projects](#) [Installing Lean and mathlib](#) [Bluesky](#) [Useful links.](#)

← FLT: Anthropic has beaten me to it

To grieve, or not to grieve?

Posted on [October 1, 2026](#) by [xenaproject](#)


Recent events in the field of AI for mathematics have shown us beyond all reasonable doubt that the field is currently undergoing a rapid transformation, unlike anything that we have ever seen before. Language models are solving hard problems which humans could not, and mathematicians are reacting very differently to this news. I personally am extremely excited about the future of our field. However it is becoming clear to me that my views are not shared by everyone in the community: indeed, many of my colleagues seem to be upset. So let me start by outlining my understanding of their positions.

Mathematicians' Responses

Mathematicians' Responses

Proofs and Prompts was launched in August as “communal blog for mathematicians’ thoughts on how AI is changing what it means to do mathematics, from PhD students and professors, skeptics and enthusiasts alike.”


There have been more than 85 contributions to date.



Mathematicians' Responses

A few days before the OpenAI announcement, Terence Tao provided context on Mathstodon.

He also noted that the lessons learned along the way are more important than the solution itself, and that AI-generated solutions disincentivizes exploration.



Mathematicians' Responses

On September 11, 25 Fields Medalists issued a statement warning of a “severe misalignment of AI in mathematics.”

More than 8,000 people have endorsed it to date.

Math and AI.

DeclarationEndorsers

A Severe Misalignment of AI in Mathematics

Published 11 September 2026 · DOI [10.5281/zenodo.22737750](https://doi.org/10.5281/zenodo.22737750) ▶ English

Over the last few months, the mathematical capabilities of LLMs have improved dramatically, to the point that they can solve major outstanding problems in many fields of mathematics. However, the push by AI companies to solve mathematical problems as a benchmark is detrimental to the science of mathematics, and to the mathematical community. The goals of the AI companies and the goals of the mathematical community are severely misaligned. We see these as part of broader alignment issues impacting other scientific and creative professions, as well as the whole of society.

Research mathematics deals with understanding basic structures of shapes, numbers, and natural phenomena. Over the course of generations, it has built a large corpus of sophisticated ideas, methods, abstractions, and other tools to comprehend the mathematical landscape. In turn, modern technologies and sciences are based on mathematical tools.

Famous problems have often served as landmarks and lighthouses against which one can measure an improved understanding of this landscape. Solving one of these problems has been a certain sign of new insights and interesting methods, which would then be studied

ADD YOUR NAME

If you support this declaration, you are invited to endorse it.

Endorse the Declaration ↗

Your endorsement appears right after you sign in with ORCID, or after you confirm your academic email address.

Mathematicians' Responses

Tao has shared several guest posts on his blog.

We are not trying to meet some abstract production quota of definitions, theorems and proofs. The measure of our success is whether what we do enables *people* to understand and think more clearly and effectively about mathematics.

William Thurston, "On proof and progress in mathematics", 1994

What's new

Updates on my research and expository papers, discussion of open problems, and other maths-related topics. By Terence Tao

 [Subscribe to feed](#)

[Home](#) [About](#) [Career advice](#) [On writing](#) [Books](#) [Mastodon+](#) [Applets](#)

RECENT COMMENTS



thoughtfully829ed2a5...
on [AIM: an invitation to explore...](#)



cosmo5d7acd87d on
[247A, Notes 2: The Hardy...](#)



sagaistrue on [Joint Statement about Mat...](#)



Anonymous on [AIM: an invitation to explore...](#)



mjc249 on [AIM: an invitation to explore...](#)



mjc249 on [AIM: an invitation to explore...](#)

How AI does, and does not, change the way I do math

29 September, 2026 in [guest blog](#), [math.GM](#) | Tags: [AI](#), [Rachel Webb](#) | by [Terence Tao](#)

[This is a guest post by [Rachel Webb](#). This blog post was initially written in a different file format and converted using AI. — T.]


The way I do math experienced a great upheaval once before, halfway through my PhD. I was in a seminar with a group of graduate students and faculty developing new discipline-specific writing courses. The facilitator asked the question,

What does it mean to write well in your discipline?

Mathematicians' Responses

On September 17-18, Harvard's CMSA hosted a Summit on PhD Math Education in the Age of AI.

A week later, the group of 24 mathematicians produced an 8-page report with recommendations.



The screenshot shows the Harvard University Center of Mathematical Sciences and Applications (CMSA) website. The header includes the CMSA logo, a "SUPPORT CMSA" button, and navigation icons. The main content area features the title "Education in the Age of AI" and "Summit on PhD Math Education in the Age of AI". Below the title, there is a quote: "As we write (September 2026) the mathematics profession is at the beginning of a profound period of change, due to the rise of artificial intelligence. The effects on the mathematics community along many axes are overwhelming. In particular, our students face great". The page also includes a section titled "Summit on PhD Math Education in the Age of AI" and a paragraph describing the summit: "On September 17–18, 2026 a group of 24 mathematicians met at Harvard to grapple with the changing landscape for mathematics PhD programs in the age of AI. We produced recommendations given in the report below. We hope they serve as a starting point for discussions in each department about how to evolve their program. We welcome [feedback via email](#)." The page concludes with a section titled "Organizers" and a list of names: Dan Freed, Joe Harris, Mark Kisin, and Melanie Matchett Wood.

HARVARD UNIVERSITY
CENTER OF MATHEMATICAL
SCIENCES AND APPLICATIONS

SUPPORT CMSA

Education in the Age of AI

September 17–18, 2026

Summit on PhD Math Education in the Age of AI

Homepage / Summit on PhD Math Education in the Age of AI

As we write (September 2026) the mathematics profession is at the beginning of a profound period of change, due to the rise of artificial intelligence. The effects on the mathematics community along many axes are overwhelming. In particular, our students face great

Summit on PhD Math Education in the Age of AI

On September 17–18, 2026 a group of 24 mathematicians met at Harvard to grapple with the changing landscape for mathematics PhD programs in the age of AI. We produced recommendations given in the report below. We hope they serve as a starting point for discussions in each department about how to evolve their program. We welcome [feedback via email](#).

Organizers

- Dan Freed
- Joe Harris
- Mark Kisin
- Melanie Matchett Wood

Mathematicians' Responses

Mathematics Departments,
Journal Editorial Boards, and
Professional Societies have held
emergency meetings, instituted
policies, and issued statements.



[Home](#) / [News](#) / [ICIAM Statement on Mathematics and Artificial Intelligence](#)

ICIAM Statement on Mathematics and Artificial Intelligence

Discovery, Understanding and Trust

Executive Statement

International Council for Industrial and Applied Mathematics (ICIAM), September 2026

A moment of extraordinary opportunity, and responsibility


Artificial intelligence is rapidly changing mathematical research, scientific computing and engineering. ICIAM welcomes these developments and believes that applied and industrial mathematics must have a strong voice in shaping them. AI can search enormous spaces, generate proofs and counterexamples, accelerate simulation and support scientific discovery. But mathematical progress is more than the production of answers, and scientific usefulness requires more than impressive benchmark performance. Understanding, validation, reliability, attribution and human judgement remain essential. In industry, AI is increasingly being combined with mathematical models and simulation to enable faster design, real-time prediction, digital twins and decision support. For ICIAM, the central opportunity is therefore a partnership in which AI extends human capabilities while mathematics provides structure, insight and trust.

1. AI can accelerate discovery, but discovery is more than producing answers

Mathematicians' Responses

An increasing percentage of arXiv submissions in mathematics report substantial use of AI.

arXiv imposed a two-paper per month rate limit.



arXiv submission rates have steadily increased since our founding in 1991. In September, arXiv received 40,363 submissions in a single month – a new record.

The Future of Mathematics

The Future

Immediate disruption:

- What should mathematics look like going forward?
- What does a working mathematician do on a day-to-day basis (especially students and early-career researchers)?
- How do we assess work (especially for students and early-career researchers)?
- How should we train our students?
- What should we tell our students?
- How much should we spend on AI subscriptions?
- How do we deal with the flood of journal submissions?

The Future

I won't address these questions here.

I am confident that mathematics will continue to thrive and will emerge even stronger than before.

I will try to explain why.

The Future of Mathematics

5 October, 2026 in [guest blog](#), [math.HO](#), [opinion](#) | by [Terence Tao](#)

[This is a guest post by [Jeremy Avigad](#). This blog post was initially written in a different file format and converted using AI. — T.]

"Mathematics underwent, in the nineteenth century, a transformation so profound that it is not too much to call it a second birth of the subject—its first birth having occurred among the ancient Greeks.."

Howard Stein, in "Logos, Logic, and Logistiké: Some Philosophical Remarks on Nineteenth-Century Transformation of Mathematics"

"The report of my death was an exaggeration."

Mark Twain

"May you live in interesting times."

(traditional)

I recently attended a meeting of innovative science and technology startups supported by Convergent Research, the organization that oversees the Lean FRO, a nonprofit that develops the Lean theorem prover. The meeting was designed to stimulate discussion, and when I introduced myself as a mathematician, many participants were eager to talk about the impact of recent events in AI on mathematics and reactions in the mathematics community. They were surprised to hear that I find the tone of the community responses on blogs like this one and [Proofs and Prompts](#) generally positive and encouraging, even though we all recognize that fundamental aspects of our day-to-day professional lives are bound to change. These discussions have helped me shape some of the thoughts I would like to share here.

The Future

A narrow view of what mathematicians do:

- we try to solve problems;
- when the hard problems are too hard to solve,
- we make up easier approximations,
- solve them,
- and then vary the parameters.

This is the practice that has been disrupted by the events of the last few months.

The Future

History offers a broader view: mathematics is a culture of rigorous reasoning and communication, providing us with language and abstractions that let us think and communicate more reliably and efficiently.

It is the exercise of capacities that are so fundamental to our survival as a species that they are wired into our DNA.

Mathematical thought isn't going away any time soon.

The Future

The challenge: solving the kinds of problems we have been solving is decoupled from the goal of enhancing our mathematical understanding, given AI.

The question, therefore, isn't whether we still need mathematics, but rather how to pursue mathematical understanding in the age of AI.

I will provide three general answers.

Solve Harder Problems

There are two salient features of today's foundation models:

- they have seen the entire mathematical literature,
- they are tireless; a swarm of agents can make its way to a solution by trying countless variations.

The AI-generated solutions to open problems all have a similar character: they are problems that AI could solve by cobbling together available techniques.

(I am grateful to Matthew Ballard for this characterization, and the subsequent analysis.)

Solve Harder Problems

Can every interesting mathematical question be answered that way?

Probably not, and even questions that can may have more interesting and satisfying solutions that invoke novel ideas and insights.

Reinforcement-learning training regimes have systems chain conventional moves and learn, from a history of failures and successes, which ones are most promising in a given state.

Solve Harder Problems

The fact that the value of an action is graded solely in terms of the success of a final trajectory breeds superhuman cleverness but may miss creativity and higher-level strategizing.

In any case, there are still hard questions to be solved and ambitious research programs to pursue, and the possibility of making progress on them with the help of AI is exciting.

Think Bigger Thoughts

In 1853, the young Bernhard Riemann submitted three potential topics to his advisor, Carl Friedrich Gauss, to choose from for his Habilitationsvortrag, a lecture he was required to give to secure a teaching position at the University of Göttingen.

Gauss reportedly chose the topic for which Riemann was least prepared, to see what he would make of it.

Think Bigger Thoughts

The resulting lecture, “On the Hypotheses Which Lie at the Foundations of Geometry,” was published posthumously in 1868, and it revolutionized the field.

The lecture distinguished a space’s metric properties from its topological properties and introduced the general notion of a manifold, though the latter did not receive a fully rigorous treatment until the twentieth century.

Think Bigger Thoughts

The focus on intrinsic properties of a space—in Riemannian geometry, those determined by the metric and independent of embedding in a larger space—was key to Einstein's theory of general relativity decades later.

Riemannian geometry has had applications that Riemann himself could never have imagined, from robotics and medical imaging to statistical analysis.

From the point of reinforcement learning:

- The benefits of Riemann's decisions are diffuse and hard to isolate.
- The time horizon is huge.

Think Bigger Thoughts

Will AI eventually be able to make advances like these?

That's not even the right question to ask.

We should care about AI only insofar as the results are interesting and important to us, and, at the end of the day, it's up to us to decide what that means.

The values we assign to mathematical developments are embedded in our history and culture.

Think Bigger Thoughts

Imagine planning a trip to go backpacking in the Alaskan wilderness for exercise and recreation.

You might be happy to let AI help book your flight, but not to let AI take the hike for you and send you pictures.

In mathematics, what is at stake isn't our recreation but agency over our reasoning and deliberation.

Think Bigger Thoughts

Whether or not AI can think, it can't think for us.

We have our own lives to live; mathematics is our story to tell, and it's up to us to decide how to tell it.

With AI, there's even more to explore, and no shortage of avenues for discovery.

Try New Things

The technologies themselves raise new questions, and that we have a lot to learn about how to use them effectively.

This requires mathematicians to give up fixed notions of what it means to “do mathematics.”

The notion has been constantly changing.

The Role of Research in AI

Keep Mathematicians in Mind

Two goals of experimentation:

- Basic science: develop methods and understand mechanisms
- Deliverables: develop tools for users.

Both are important!

Develop More Efficient, Scalable Approaches

Some reasoning is best done at the level of natural language.

Other things are better done symbolically: calculating, checking correctness.

Better neurosymbolic combinations will get us further, with less cost and less environmental damage.

Focus on User Interaction

You have coding (largely) solved; now help programmers come up with the right specifications.

You have proving (largely) solved; now help mathematicians reason, explore, and discover.

Alignment involves going from

- performing tasks to
- helping people think about what they want and how to get it.

Collaborate with Mathematicians

Warning:

- We are spacy.
- We are opinionated and, at times, arrogant.
- We do mathematics because we want to.
- There are lots of things we don't care about.
- We don't care about citations.
- We don't care what people think.
- We don't believe in deadlines.

Try it. You might like it.