Where is mathematics going?

Kevin Buzzard, Imperial College London

Simons presidential lecture, 24th September 2025

Summary of the talk

- Where mathematics is now;
- What's happening today;
- Where is mathematics going?

By “mathematics” here I mean mathematical research.

Preliminaries

I thank the Simons Foundation for giving me the opportunity to speak here.

Background: I'm a maths professor in London; in 2017 I became disillusioned with the way mathematics research was being done, and started looking to computer science for solutions.

What mathematics looks like

Over 2000 years ago, Euclid wrote down some axioms.

He then proceeded to make logical deductions.

Each deduction became a fact, usable in future deductions.

And today in mathematics departments around the world, we are still doing the same thing.

So mathematics is now incomprehensibly large.

What mathematics looks like

Euclid wrote down his ideas in a series of self-contained books.

A lot of early mathematics was self-contained in this way.

But by the 1800s things were beginning to “pile up”.

By 1900 mathematicians had proved the main theorems of global class field theory.

The proofs were long, technical, and relied on a lot of earlier work (for example, Galois theory).

Even today the proof takes an entire graduate level textbook to explain.

By 2004, mathematicians had proved the Classification of Finite Simple Groups.

We're still trying to write the $\approx 10,000$ page proof down coherently.

Where is mathematics now?

I claim that the mathematical ideas of humans have now become so complex that traditional methods of documenting it are now struggling to cope.

Every day I get an email from `arxiv.org` containing links to hundreds of pages of number theory articles which were uploaded in the previous 24 hours.

It is not uncommon to see 100- or even 200-page papers in my area, perhaps with 50-100 references.

Referee work is unpaid and it is “not the job of the referee to check all the details”.

The goal of our top undergraduate programs seems to be “get to the 1940s”.

Contrast this with computer science, which can get undergraduates to the forefront of modern understanding.

Does mathematics have problems?

This depends on who you ask.

Our papers are (rarely) incorrect, (sometimes) incomplete and (often?) are not being refereed carefully.

Key information, missing from the paper literature, is "known to the experts", who will die.

Historians of mathematics will tell you that this has been the state of things for at least a century.

Mathematicians are hence used to this state of affairs.

I ask: does that make it OK?

Can we do better?

Technology and mathematics

People have been doing mathematical research for thousands of years, mostly without technology.

The advent of "computer as calculator" in the 1960s had some effect.

But perhaps less than you think.

More recent innovations are "computer as generator of mathematics" and "computer as checker of mathematics".

More precisely: the large language model and the computer proof assistant.

How do these tools interact with research mathematics today?

The large language model

This is ChatGPT/Claude/Gemini/Deep Seek etc.

They are systems which can generate human-readable answers to questions, via "next token prediction".

General opinion seems to range from "these things will solve every problem known to humankind" to "these things will destroy humankind".

But how good are they at mathematics, today?

We are not seeing the hype yet

The hype machine (an essential component of the funding model) says that LLMs can solve International Mathematical Olympiad problems, they have become good guessers at “what is this number” problems, and hence within a few years they will be superhuman at mathematics.

The truth *today* is that LLMs have never told us anything profound which we didn't know already.

They can be used to brainstorm ideas, and they occasionally make helpful observations.

Hallucinations remain a major achilles heel.

Turns out that humans want LLMs to be sycophantic.

But one hallucinated false claim can destroy an entire mathematical argument.

The computer proof assistant

Mathematics can be thought of as a puzzle game played on an infinite board.

An *interactive theorem prover* (ITP) or *computer proof assistant* knows the rules of the game, and can check mathematical arguments, as long as they are written in its language.

Examples of ITPs: Lean, Rocq, Isabelle etc.

These tools have existed for decades and were mostly being used by computer scientists.

The 2020s have seen the adoption of these tools (and in particular the Lean theorem prover) by mathematicians.

Recent work of Fields Medallists has been checked in Lean.

Writing the code is currently a manual translation process.

LLMs and ITPs

Just like the LLM, the ITP has told us nothing profound which we didn't know already.

Neither of these technologies has had a “Deep Blue moment”.

Note that these systems have very different strengths and weaknesses.

LLMs can write text by themselves, but are not always accurate.

ITPs rely on external sources for input, but have superhuman accuracy.

The no-brainer idea: let's put them together.

Putting them together

A “raw” language model has no difficulty multiplying one-digit numbers together.

However, a “raw” language model cannot multiply ten-digit numbers together accurately.

Idea: you can “enhance” the model by teaching it to write python code to do such calculations.

The result: ChatGPT can now multiply ten-digit numbers together accurately.

Putting them together

A “raw” language model has no difficulty proving Pythagoras’ theorem.

However, a “raw” language model makes logical errors or leaves gaps (“vibe proving”) when attempting to prove a modern conjecture.

Idea: can we enhance the model by teaching it to write Lean code to check its proofs?

Can we eliminate hallucinations in LLM output, like we eliminated arithmetic errors?

Such an idea could extend well beyond mathematics.

Obstructions to this idea

This idea (LLM + ITP) is of course not new.

When speaking about it in the past, I have highlighted two major obstacles.

The first obstacle: there is not enough Lean code to train on.

My impression is that this is no longer an obstacle.

Example: several companies have now independently developed systems which can autonomously solve IMO problems in Lean.

Obstructions to this idea

But here's another obstruction to doing research mathematics in an LLM/ITP hybrid:

No modern ITP even understands the *statements* of what is going on in most modern mathematical research.

Hundreds, and probably thousands, of modern mathematical definitions are missing.

How about getting an LLM to just write these definitions by itself?

Problem: there are *no guardrails* when making definitions, so hallucinations are again possible.

Two projects I am involved in

I am involved in an EPSRC/UKRI-funded project whose goal is to manually formalise the papers of Wiles and Taylor–Wiles proving Fermat's Last Theorem (pic).

This will force us to teach Lean's mathematics library some of the definitions being used in modern number theory.

I am also involved in a Renaissance Philanthropy "AI for Math" project whose goal is to manually formalise the *statements* of all the recent theorems in the top maths journals.

This will force us to teach Lean's mathematics library many of the definitions being used in modern mathematics.

These are my attempts to make progress on what I think is an extremely important problem.

Where is mathematics going?

I suggested that Lean's mathematics library was missing perhaps several thousand relevant definitions.

This is not a large number.

One could imagine both humans and LLMs working on formalising these definitions.

Until recently there was a serious obstacle to this vision.

Thanks to the generosity of Alex Gerko we may well have a solution.

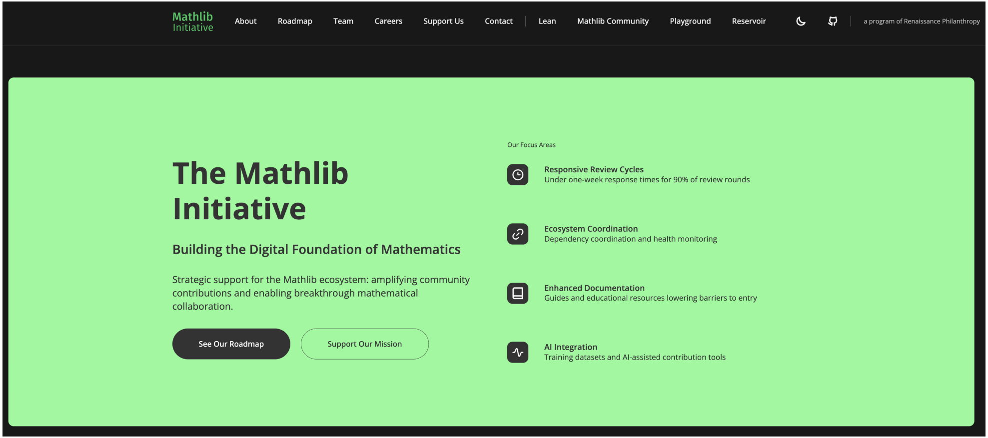

mathlib-initiative.org

Where is mathematics going?

Once we have many modern definitions, we can turn to autoformalizing the theorems in the literature.

Note: "get a system to prove the Riemann hypothesis" is currently science fiction.

But "get a system to translate the literature into Lean" is a more feasible proposal, which I believe will ultimately help to bring that science fiction closer to reality.

Example: the `math.inc` startup just announced that their system had successfully formalized a human proof of the prime number theorem.

The way it worked: humans broke the proof up into small pieces, which they typed up in a LaTeX "blueprint", and then a machine wrote the Lean code.

If the machine got stuck, then humans wrote more details.

Key question: will the systems get good enough to not get stuck on every line?

Where is mathematics going?

Mathematicians have a choice.

They can continue the way it has always been.

Or they can choose to upgrade their system.

Upgrading will take time and money.

But it has now got to the point where it seems viable to me.

So we need to decide: is “mostly right” good enough?

Or should we now choose “definitely right”?

Thanks for your attention.