Lost in Space-Time

Proving it Proofs, the central tenet of mathematics, occasionally have errors in them. Might computers be able to prevent this from happening? asks **Emily Riehl**

Emily Riehl is a professor at Johns Hopkins University, where she works on category theory, a subject about the relationships between mathematical objects. Lost in Space-Time is a monthly email newsletter from the frontiers of the cosmos

Emily's week

What I'm reading
Advanced Marathoning *by Pete Pfitzinger and Scott Douglas while preparing for the Baltimore Marathon.*

What I'm watching
Nothing, as I'm travelling to various summer mathematics conferences, but before I left home I stayed up late to finish season 2 of Yellowjackets.

What I'm working on
Formalising the results from an old joint paper of mine on infinite-dimensional category theory in a new experimental computer proof assistant called rzk.

ONE miserable morning in 2017, in the third year of my job as a mathematics professor, I woke up to a worrying email. A colleague had questioned the proof of a key theorem in a paper I had co-authored. The proof, he noted, appeared to rest on a tacit assumption that was not warranted.

Much to my alarm, I realised immediately that he was correct. After an anxious week working to get to the bottom of my mistake, it turned out I was very lucky. The theorem was true; it just needed a new proof, which my co-authors and I supplied in a follow-up paper. But if the theorem had been false, the whole edifice of consequences "proven" using it would have come crashing down.

The essence of mathematics is the concept of proof: a combination of assumed axioms and logical inferences that demonstrate the truth of a mathematical statement. Other mathematicians can then attempt to follow the argument for themselves to identify any holes or convince themselves that the statement is indeed true. Patched up in this way, theorems originally proven by the ancient Greeks about the infinitude of primes or the geometry of planar triangles remain true today – and anyone can see the arguments for why.

Proofs have meant mathematics has largely avoided the replication crises in other sciences, where the results of landmark experiments haven't held up when repeated. But as my experience shows, mistakes still occur. Ideally, a false claim would be caught by the peer review process. In practice, however, peer review in maths is less than perfect – not just because experts can make mistakes too, but because they often do not check every step in a proof.

MARSHALL CLARKE

This is not laziness: theorems at the frontiers of mathematics can be dauntingly technical, so much so that it can take years to confirm the validity of a proof. Vladimir Voevodsky, who received a Fields medal, the discipline's highest honour, noted that "a technical argument by a trusted author, which is hard to check and looks similar to arguments known to be correct, is hardly ever checked in detail". After several experiences in which mistakes in his proofs took over a decade to be resolved, Voevodsky's crisis of confidence led him to abandon his "curiosity-driven research" to develop a computer program that could verify the correctness of his work.

> **"Theorems at the frontiers of maths can be so technical it can take years to confirm the validity of a proof"**

This is known as a proof assistant, though it might be better called a "proof checker". It can verify that a string of text proves the stated theorem. The proof assistant knows the methods of logical reasoning and is equipped with a library of proofs of standard results. It will accept a proof only after satisfying each step in the reasoning process, with no shortcuts of the sort that human experts often use.

Computer proof assistants can be used to verify proofs that are so long human referees are unable to check every step. They can also be used to verify results in subfields so technical that only specialists understand the meaning of the central concepts.

Fields medallist Peter Scholze, for example, spent a year working out the proof of a theorem he wasn't quite sure he believed and doubted anyone else would have the stamina to check. To be sure his reasoning was correct, he posed a formalisation challenge in December 2020. The mathematics involved was so cutting edge it took 60,000 lines of code to formalise the last five lines of the proof – but nevertheless this project was completed this past July by a team led by Johan Commelin.

Could computers just write the proofs themselves, without involving any humans? At present, large language models like ChatGPT can fluently generate mathematical prose. However, the logic of these "proofs" tends to be nonsense. Researchers at Google and elsewhere are looking to pair large language models with automatically generated formalised proofs to guarantee the correctness of the mathematical arguments, though initial efforts are hampered by sparse training sets. But while machine capabilities are relatively limited today, auto-formalised maths is surely on its way.

In thinking about how the human mathematics community might wish to collaborate with computers in the future, we should return to the question of what a proof is for. It has never been solely about separating true statements from false ones, but about understanding why the mathematical world is the way it is. While computers will undoubtedly help humans check their work and learn to think more clearly – it is a much more exacting task to explain mathematics to a computer than it is to explain it to a kindergartner – understanding what to make of it all will always remain a fundamentally human endeavour. ■