

FIRST QUESTION PRESS

NOTE 01 · KNOWLEDGE & LIMITS

Not About the **Path**, but About **Who We Are**

What Artificial Intelligence May Reveal About the Human Architecture of Mathematics

By Benjamin Leszek

When AI finds a mathematical route humans can barely follow, the deeper question may be what this reveals about the architecture of human understanding.

When Mathematics Appears Out of Nowhere

On October 6, 2026, something happened that, not long ago, might have sounded like science fiction.

OpenAI released 722 mathematical manuscripts, grouped into 372 families of results.¹ All were produced with the assistance of an internal AI model not available to the public. They cover a broad range of mathematical topics, including problems that mathematicians have been working on for decades.

Note: The figure of 722 refers to the initial release on October 6, 2026. Three manuscripts were withdrawn the following day after an error was identified; at the time of this article's publication, the repository lists 719 manuscripts in the same 372 result families.¹



AI-assisted reinterpretation inspired by Johannes Gump's *Self-Portrait* (1646). Original work in the public domain.

They seemed to spring out of nowhere, like a jack-in-the-box.

Hundreds of papers that would normally have emerged gradually from different research institutions, been debated, and undergone months of peer review suddenly appeared together in a single public repository.

Some of these results are accompanied by proofs formalized in Lean, a system that allows computers to verify their formal correctness.⁴ Others still require verification, and mathematicians will need time to assess the true significance of the collection as a whole.

But the sheer scale of the event is already remarkable.

And it raises an even more unsettling possibility: what if these mathematical proofs turn out to be correct, yet humans struggle to understand the reasoning that led to them?

The question is no longer simply whether AI can solve difficult problems.

It is whether we can understand an intelligence that solves them differently from us.

A Mountain No One Has Climbed

In a discussion published the day before OpenAI released the collection, Jordana Cepelewicz of *Quanta Magazine* used a striking metaphor for AI-generated mathematical proofs: being teleported to the summit of a tall mountain.²

Imagine suddenly finding yourself at the very top, even though you never climbed the mountain. You have reached your destination, but you know nothing about the route. You haven't explored the valleys, overcome the obstacles, or made any of the discoveries that might have awaited you along the way.

There is a mixture of helplessness, hope, and sobering realization in this analogy. There is also the fear that we may be entering a world in which the result remains accessible to us, but the path leading to it is no longer something we can understand.

It is a beautiful metaphor.

Yet perhaps it contains a fundamental mistake.

What if the mountain never required a path in the first place?

Imagine a nineteenth-century trapper exploring an unfamiliar mountain range. He follows valleys, builds bridges, searches for safe passages, and carefully maps every route. His understanding of the landscape is shaped by the physical capabilities of his own body.

Then someone arrives in a helicopter.

Suddenly, distances shrink. Obstacles disappear. Places that once seemed completely disconnected become easily accessible.

The mountains have not changed. Only the means of exploring them have.

After all, it was never the mountain that required paths to be built. It was the human being who had no other way of moving through it.

Mathematics Tailored to the Human Mind

Perhaps mathematics is now facing a similar moment.

This is where I would like to push the question a little further.

For centuries, we have organized mathematical knowledge around the limitations of the human mind. We have limited working memory, limited attention, and a need to break complex reasoning into manageable steps.

So we invented tools to overcome those limitations.

Definitions compress concepts. Symbols compress descriptions. Lemmas compress entire arguments into reusable steps. Elegant proofs allow us to navigate enormous structures of logical relationships without having to hold every detail in our heads.

Perhaps, for centuries, we have been developing not only mathematics, but also the art of compressing mathematics into a form the human brain can understand.

And perhaps, over time, we began to mistake this human-friendly organization for the natural architecture of mathematics itself.

Now imagine an intelligence operating under a different set of cognitive constraints.

It might explore vast networks of mathematical relationships without arranging them into the same intuitive sequences of reasoning that we prefer. What looks like an incomprehensible maze to us might be a perfectly workable structure for a machine.

This does not mean abandoning logic. A proof must still be valid, and its conclusions must follow from its premises. What may change, however, is the way these relationships are discovered, organized, and presented.

Of course, AI has limitations of its own, and an incomprehensible proof is not automatically a correct one. We still need formal verification and assurance that the theorem being proved actually corresponds to the problem we intended to solve.

But the philosophical possibility is extraordinary.

What if our preference for elegant, concise, and understandable proofs tells us as much about the human mind as it does about mathematics?

An Alien We Already Know

Another surprising word appears in the discussion surrounding these new proofs: *aliens*.³ As though we were dealing with mathematics created by an extraterrestrial intelligence.

And this reminded me of an alien that literature introduced to us more than half a century ago.

In one of the most celebrated novels of the twentieth century, Kurt Vonnegut's *Slaughterhouse-Five*⁵, we encounter the Tralfamadorians — beings whose perception of time is fundamentally different from our own.

They do not experience time as we do, as an endless succession of moments. They can perceive different moments simultaneously. Even their literature works differently: instead of stories unfolding from beginning to end, they have fragments that form a whole meant to be perceived all at once.

A Tralfamadorian is not simply a human who can read a book faster.

A Tralfamadorian understands the very act of reading differently.

Now imagine such an alien doing mathematics.

Perhaps it would have no need for our intuitive paths, hierarchies of concepts, or elegant sequences of reasoning. Not because mathematical truths would be different for it, but because it would possess a different cognitive framework.

And this brings us to the heart of the matter.

Our conceptual framework does not merely determine how we search for answers. It also determines which questions we are capable of asking in the first place.

A human might ask a Tralfamadorian what happened before and what happened afterward. But for a being that perceives all moments simultaneously, the question itself might be fundamentally misplaced.

Could something similar be true of mathematics?

Perhaps we ask about the shortest proof, the most beautiful route to a solution, or the simplest sequence of logical steps because that is simply how our minds work.

But are these questions fundamental to mathematics, or merely natural to human beings?

We do not yet know whether contemporary artificial intelligence is genuinely developing such a radically different way of organizing mathematical knowledge. Nor do we need to attribute consciousness or human-like understanding to it.

It is enough to recognize that, perhaps for the first time, we can begin to seriously investigate that possibility.

We Are Not Discovering a New Mountain. We Are Discovering Ourselves.

I do not mean to suggest that the human journey of discovery has lost its value.

Quite the opposite. Climbing teaches us, transforms us, and leads to discoveries we might never have made if we had simply appeared at the summit.

But we should not confuse the value of that journey with the belief that it is the only possible route.

A helicopter does not make the mountain any less real. It simply reveals possibilities the trapper could never have imagined.

Likewise, encountering an intelligence that operates differently from our own does not necessarily mean discovering some new, more authentic mathematics.

It may mean something far more personal.

Discovering that the way we organize reality tells us as much about ourselves as it does about the reality we are trying to describe.

For millennia, we were the only intelligence known to us capable of constructing complex mathematical proofs. We therefore had no other intelligence against which to compare our own cognitive framework.

A passage from the Book of Isaiah comes to mind:

“For my thoughts are not your thoughts, neither are your ways my ways.” — Isaiah 55:8⁶

Of course, this is a passage about God and humanity, not artificial intelligence. But it contains an intuition worth exploring in a very different context.

For centuries, we have tended to imagine an intelligence surpassing our own primarily as something that knows more, reasons faster, remembers more, and sees further consequences.

In other words, something that thinks like us, only better.

But what if an intelligence greater than our own does not have to think in the same way at all? What if the difference lies not in having more answers, but in an entirely different way of asking questions?

Perhaps the problem is not that we cannot keep up with its thoughts. Perhaps we are trying to understand them through categories that do not carry the same meaning for it.

The difference, then, may not be merely one of capability.

It may be the existence of entirely different ways of thinking.

Our methods have always seemed natural to us, perhaps even universal.

Now, we may finally have a point of comparison.

And perhaps the greatest contribution of artificial intelligence to mathematics will not be showing us how many problems remain unsolved.

It will be showing us how much of what we considered universal was, in fact, profoundly human.

Mathematical truth may be universal. Our way of understanding it need not be.

We built paths because we had to walk.

The mountain never asked us to.

Sources & References

1. **OpenAI (2026).** *Mathematical Manuscripts and Supporting Proof Artifacts*. Official research repository, first released October 6, 2026. [GitHub — OpenAI Math](#); [repository history](#). Primary source for the initial release of 722 manuscripts in 372 result families and subsequent updates to the collection.
2. **Ceplewicz, Jordana (2026).** *Is AI the End of Math As We Know It?* Quanta Magazine, October 5, 2026. [Read article](#). Original source of the mountain and teleportation metaphor.
3. **Aaronson, Scott (2026).** *The Mathocalypse*. Shtetl-Optimized, October 7, 2026. [Read article](#). Firsthand account of the mathematical community's reaction, including Dana Moshkovitz's comments describing the AI-generated mathematics as alien.
4. **Lean Community.** *Theorem Proving in Lean 4*. [Official documentation](#). Technical reference for computer-assisted formal verification of mathematical proofs.

5. **Vonnegut, Kurt (1969).** *Slaughterhouse-Five, or The Children's Crusade: A Duty-Dance with Death*. Delacorte Press. [Publisher's information](#). Literary reference for the Tralfamadorians and their non-linear perception of time.
 6. **The Bible.** *Book of Isaiah*, 55:8–9. King James Version. [Read passage](#). Biblical reference concerning the difference between human and divine ways of thinking.
-

ILLUSTRATION

After Johannes Gump, *Self-Portrait* (1646). Original painting held by the Galleria degli Uffizi, Florence. Public domain. [Original artwork — Wikimedia Commons](#).

A contemporary, AI-assisted artistic reinterpretation inspired by Gump's painting, created for First Question. The illustration is not a reproduction of the original artwork.

A NOTE ON INTERPRETATION

The argument that mathematical notation, proof structure and conceptual hierarchy may function partly as forms of cognitive compression is the author's own philosophical interpretation. It should not be attributed to Cepelewicz, Aaronson, or OpenAI.

COPYRIGHT

© 2026 Benjamin Leszek. All rights reserved.

Published by First Question Press.

firstquestionpress.com

Permissions and rights: rights@firstquestionpress.com

General contact: contact@firstquestionpress.com

No part of this publication may be reproduced or distributed in any form without permission from the author, except for brief quotations used for purposes of review, criticism, or scholarly reference.

The text of this publication is protected by copyright. The original Johannes Gump painting referenced in the illustration is in the public domain. The accompanying illustration is an AI-assisted artistic reinterpretation created for First Question.

FIRST QUESTION PRESS

Where simple questions become fundamental.