



# The Infinity Project

*How to use AI and mathematics to prove and improve science and security* | **Patrick Shafto**

# The Infinity Project

*How to use AI and mathematics to prove and improve science and security* | **Patrick Shafto**

---

This essay is part of [The Launch Sequence](#), a collection of concrete, ambitious ideas to accelerate AI for science and security.

## Summary

Mathematics is the foundation of immense progress, underlying revolutionary technologies from the computer to cryptography, AI, and the internet. Providing firmer mathematical foundations to science and cybersecurity could lead to transformative advances in those fields. The Infinity Project will broaden access to mathematics by grounding math in computer programs. Advances in programming languages have demonstrated the possibility of full formalization of traditional natural language mathematics in formally verified computer programs that automatically check correctness.

The Infinity Project is a \$112.5 million grant contest among mathematics institutes. The project capitalizes on advances in AI translation and computer programming to — within five years — formalize the entirety of mathematics, creating the infrastructure for anyone to automatically produce and verify a mathematical proof.

# Motivation

The Infinity Project will leverage AI to radically increase the application of mathematics to improve science. Mathematics has contributed unquantifiable value to society, despite being accessible to only a select few, due to the rigorous training required to produce correct proofs.<sup>1</sup> AI has the potential to make the power of mathematics accessible to all by grounding math in computer programs and automating proofs.

Mathematics is an underappreciated source of value for society. The computer, the foundation of a \$295 billion industry (hardware alone, in [2024](#)), was a [mathematical](#) object before it was a physical object.<sup>2</sup> Cryptography, the foundation of internet security (and therefore commerce), is premised on mathematical [problems](#). The multibillion-dollar AI industry is built upon [linear algebra](#). Google's \$265 billion per year [advertising empire](#) is run by the [mathematics](#) of game theory. Mathematics offers a natural language for establishing truth, but it's so difficult to produce valid proofs that most people can't use it effectively, limiting the field's impact.

There is no shortage of problems across science and cybersecurity that would benefit from advances in mathematics. For instance:

- Science is a chain of human-mediated processes: hypothesis generation, experimental design, data collection, and analysis. Sciences including [biology](#) and natural and [artificial intelligence](#) lack rigorous mathematical theory, with [demonstrated negative](#) consequences. The National Institutes of Health (NIH) alone is a \$48 billion annual investment in science that could be more efficient and effective with increased rigor and automation, which could be achieved by extending statistical inference from sets of parameters to entire mathematical theories using computer programs.
- Data and cybersecurity breaches are common (more than [3,200 data breaches](#) annually), expensive ([\\$4.9M per data breach](#) on average in 2024),

---

<sup>1</sup> Roughly 2,000 PhD Mathematicians [graduate](#) each year, compared with approximately 21,000 [medical](#) doctors.

<sup>2</sup> Turing, A. M. (1938). On computable numbers, with an application to the Entscheidungsproblem. A correction. *Proceedings of the London Mathematical Society*, 2(1), 544-546.

and a threat to industry and national security. A primary vector of breaches is hacking, which exploits security flaws in software, and sometimes hardware, systems. Formal verification of these systems has the potential to greatly reduce data and security breaches, but lacks a general, automated framework for proving correctness. Whereas formal verification is pursued on a case-by-case basis, mathematics is the natural way to prove general, automatable results. Mathematical advances could save billions of dollars each year and secure systems central to national security.

AI has the potential to radically broaden accessibility to (and applications of) proof-based mathematics by grounding natural language math in computer programs. Mathematicians have increasingly used formal verification, in computer programs such as [Lean](#), to [confirm](#) that small aspects of their proofs are correct, and to [distribute](#) large problems across people in a trustworthy manner. The computer programs verify that a proof is correct, but to do so, users must [translate](#) the entire proof from natural language math to computer code — a slow, tedious, and highly expertise-intensive task. While modern AI is notorious for [hallucinating facts](#), failing on [long tasks](#), and its limitations in [complex reasoning](#), it has excelled at [translation](#) and [computer programming](#). If we could leverage AI's successes at translation and computer programming for mathematics, we would alleviate the expertise bottlenecks that prevent mathematics from being broadly accessible, with transformative implications.

## Solution

The Infinity Project will address the main challenges that prevent AI from radically broadening access to mathematics. AI translation succeeded based on the enormous availability of data. Similar success in mathematics requires producing training data in the form of manual translations of math statements into computer programs. The second and third challenges are that math requires performance of long tasks with complex reasoning, which are limitations of current AI, due in large part to the lack of training data. Fortunately, mathematics is structured to decompose long, complex reasoning into smaller, simpler parts. Therefore, systematic training of AI systems on paired natural language math and computer

programs will address this limitation of current AI. Hallucinations will be naturally restricted by formal verification, which requires correctness.

Industry is unlikely to address this challenge for two reasons. First, math has no direct value. Math enables innovation, but is not itself innovation. For this reason, math has remained a public good for centuries. Second, the technical expertise required to lead this effort is of unprecedented scale. Success will require shared effort of mathematicians across mathematics, an organizational effort that requires the trust of the community.

The Infinity Project is structured as an award contest among mathematics institutes. Institutes will propose collaborative efforts together with academia, industry, and society that address the main challenges, and enable the creation of tools and infrastructure to realize the transformative implications of mathematics for society. Mathematics institutes, (such as the [Institute for Advanced Study](#), [SLMath](#), and the [Simons Institute for Theory of Computing](#), potentially alongside new institutes) will organize teams, lead proposals, and serve as leads. As the intellectual homes of mathematics in the US, these are trusted partners within the math community, capable organizers of the mathematical community, and home to the technical expertise required to lead.

Proposals will be evaluated based on their teams and plans for impact. Teams will be composed from academia, industry, and society. Academia will provide intellectual leadership, research, training, and personnel to address the challenges. Industry will provide complementary intellectual leadership, infrastructure, and valuable applications. Interests of society will be represented by scientists, governments, and/or community organizations. The purpose of these broad coalitions will be to foster interest in mathematics and AI, and to include transition partners who use the technology and contribute back to the projects as citizen scientists.

A board of advisors will be selected from the worlds of mathematics, philanthropy, AI, and industry. The advisors will be selected based on prior achievements in their fields, and their investment in the topics of mathematics and AI. The board will select the winning proposal(s) and provide oversight throughout the program's duration.

## Project activities

The technical work is fourfold: creating a high-quality library of formalized math, using AI models to further generate formal mathematics, training math-augmenting AI agents, and demonstrating societal implications. An \$112.5 million investment would enable the award recipient(s) to hire a core team of mathematicians and AI engineers dedicated to the project, marshal a wide network of collaborators to assist and support the project, and demonstrate societal implications. Specifically, the Infinity Project would:

- **Establish a library of formalized and natural language mathematics.**  
Natural math statements must be translated manually, by human experts, to equivalent statements in a computer program capable of formal verification, such as Lean. Manual translation is necessary to facilitate the training of AI systems in mathematics and to ensure connection to mathematical practice by humans. This can be achieved by experts in a decentralized fashion, with a central coordination committee, organized by the lead institute.
- **Further generate formal mathematics.** AI assistance will be important so models can accelerate the formalization of mathematics, and back-translate formalized mathematics into natural language mathematics. Correctness can be ensured via formal verification of computer programs. This will radically accelerate the formalization of mathematics and high quality data on which to train AI agents.
- **Develop math-augmenting AI agents.** Math-centric AI agents will be trained on the library of formalized and natural language math in order to automatically generate novel mathematics. The key challenge is for the mathematics to not merely be true, but to introduce useful concepts that help solve other math problems. The ability to introduce new, useful concepts is especially important for broadening the accessibility and usefulness of mathematics across domains.
- **Demonstrate implications for science, technology, and society.**  
Demonstrations will be developed to illustrate the transformative implications of mathematics and AI. Formalization of scientific pipelines such as lab work in biology and technological objects such as computer

programs will serve as demonstrations of the potential of broader mathematical foundations. Tutorial materials and outreach programs will be developed to foster education and widespread adoption across the economy and society.

Multiple complementary efforts exist to support the Infinity Project. The Lean Programming language is supported by the [Lean FRO](#), and focuses on the language itself, rather than mathematics. The [mathlib](#) community is a collection of volunteers who have begun the seeds of formalization of mathematics in Lean. Google DeepMind has been working on AI for math, though focused on [formal](#) rather than natural language math.<sup>3</sup> Amazon has significant investments in formal [verification](#), but not for math and only on a case-by-case basis. DARPA has been [promoting](#) formal verification as critical for defense and security applications. The DARPA [expMath](#) program focuses on innovations in AI related to abstraction in mathematics, but is not as ambitious as the breadth of the Infinity Project.

## Recommended actions

The US government should:

- **Convene a governmental advisory board.** Mathematicians and AI researchers should be brought together with members from philanthropy, industry, and broader society to provide strategic guidance to the project, and interface with the broader scientific, technical, educational, and governmental communities.
- **Establish dedicated funding for mathematics centers.** Allocate \$112.5 million over five years to transform mathematics with AI. This funding should combine elements dedicated to specific activities, as well as sums for sustainment beyond the five-year period.
- **Leverage shared interests.** Industry, philanthropy, academia, and existing governmental agencies have a shared interest in mathematics and AI. Technology companies large and small (Amazon, Google, OpenAI, [xAI](#)) have

---

<sup>3</sup> Deepmind has four teams: AlphaGeometry, AlphaEvolve, AlphaProof, and PDEs. Of these, AlphaGeometry and PDEs have narrow foci. AlphaEvolve is not general to mathematics. AlphaProof has been working in Lean but strictly with formalized statements, and up to this point on high-school level (International Math Olympiad) problems.

invested in this space, as have philanthropic organizations including Renaissance Philanthropy and the Simons Foundation. US academic institutes are home to the world’s leading mathematicians and scientists and are supported by a diverse array of government agencies. However, the scope of the challenges requires government-scale investment and leadership.

## Funding requirements

We estimate that the Project will require roughly \$112.5 million over five years:

| Category                        | Cost per unit             | Annual total | 5-year total |
|---------------------------------|---------------------------|--------------|--------------|
| <u>Personnel</u>                |                           |              |              |
| 6 Mathematicians                | \$250K                    | \$1.5M       | \$7.5M       |
| 5 AI Scientists / Engineers     | \$500K                    | \$2.5M       | \$12.5M      |
| 2 Information Scientists        | \$250K                    | \$0.5M       | \$2.5M       |
| 2 Programmers                   | \$250K                    | \$0.5M       | \$2.5M       |
| 4 Executives and Administrators | \$250K                    | \$1M         | \$5M         |
| Benefits and Taxes              | 50%                       | \$3M         | \$15M        |
| AI Services                     | \$4M                      | \$4M         | \$20M        |
| Grants to Community             | \$250K x 20               | \$5M         | \$25M        |
| General and Administrative      | 25% of Total Direct Costs | \$4.5M       | \$22.5M      |
| <b>Total</b>                    |                           | \$22.5M      | \$112.5M     |

We expect that a substantial portion of this cost could be offset by shared investments from industry, academia, and existing governmental agencies. Industry could supply both compute resources and technical expertise. Academia can supply expertise and can pay for advanced capabilities services.

Governmental agencies can redirect existing funds for science and technology to this project.

The Infinity Project can produce an extremely high return on investment. Large portions of science and technology are bottlenecked by imperfectly designed processes. Mathematics is a natural solution because it has established transformational value, and many of the challenges we face could be addressed by making mathematics broader and more accessible in two key areas: the sciences, and securing software in computers, cars, refrigerators, pharmaceutical plants, and the electric grid. Just as improved rigor enabled multibillion-dollar industries founded on physics and computers, the Infinity Project promises new multibillion-dollar industries built on novel, reliable, and scalable mathematical foundations.

## Further resources

- National Academies of Sciences, Engineering, and Medicine, [\*Artificial Intelligence to Assist Mathematical Reasoning: Proceedings of a Workshop\*](#), 2023.
- Lean-lang.org, [\*Lean\*](#), n.d.  
The Lean programming language and Theorem prover.
- DARPA, [\*Ushering in the New Era of Cyber Resilience\*](#), n.d.  
A formal methods colloquium.
- Kaiyu Yang et al., [\*Formal Mathematical Reasoning: A New Frontier in AI\*](#), arXiv, 2024.  
A position paper.

# Appendix

This appendix contains a detailed example of what formalization of science might mean, concretely and practically, in operational technology systems and in behavioral science research, together with the potential implications.

**Operational technology systems** play a key role in areas that rely heavily on instrumentation, measurement, and process, such as many sciences. An example from biology is the Western blot, a technique in molecular biology used to identify specific proteins from among a complex mixture. The key idea is to separate proteins by size, which requires mounting samples on a gel, dying the cells, and applying an electrical field. The results are used to confirm, for example, HIV tests, and play a role in the development of drugs for diseases like leukemia. Automation means that what was once a manual process by humans can be conducted by machines. Western blot, and similar machines, are critical to modern biological science and the clinical and pharmacological fields that commercialize these processes for human health.

Machines that automate the Western blot also represent attack surfaces for cyber threats. Modern machines contain computer chips and are connected to the internet. Therefore, they can be hacked with the goal of incapacitating the industry or harming significant portions of the population. Formal verification of machines would be an important step toward protecting them, and us, from cybersecurity threats. Moreover, formalizing such machines is the first step toward treating them abstractly as mathematical objects, and proving statements about more complex processes.

Currently, the workflow of an automated Western blot process links multiple machines, with humans intervening at various points. At each of these steps, failure is possible. It is also possible that failures can arise as a consequence not of any single step, but through interactions between different steps. If individual components were formalized, it would open the door to formalization of multiple steps of the process, through mathematical concepts such as composition. This is of special importance for large-scale operations, such as pharmaceutical production, where most or all of the sequence of operations is conducted by machines. The result of formalization will be assurances regarding safety,

abstraction of common processes and structures, and consequently much greater potential reliability of science and manufacturing in biology.

**Behavioral sciences** study why people do what they do, with the goal of improving education, health, and society. Researchers observe human behavior, identify factors that might influence it, design experiments to test their ideas, and evaluate real-world interventions. For example: What makes certain subjects difficult for students to learn? What helps people make healthier choices? What support systems actually work?

The behavioral sciences have faced a serious problem: [many studies don't replicate](#). When other researchers try to repeat an experiment, they often get different results. This "replication crisis" undermines the credibility of the entire field.

One major culprit (though by no means the *only* culprit) is something called "[p-hacking](#)," or [torturing the data long enough until it confesses](#): Researchers collect data and then run dozens of different statistical analyses, trying various combinations until they find something that looks significant. Maybe they exclude certain participants, or analyze the data in an unorthodox way, or look at unexpected relationships between variables. Eventually, by pure chance, they'll find something that appears meaningful, but it's actually just statistical noise. This practice is now recognized as scientific malpractice because it dramatically increases the likelihood of publishing "findings" that aren't real, and therefore don't replicate.

To combat this problem, behavioral scientists increasingly use [preregistration](#). Before running an experiment, researchers write down exactly what they plan to do: their hypotheses, how they'll collect data, and — crucially — exactly how they'll analyze it. This document gets time-stamped and stored publicly. The idea is simple: if you commit to your analysis plan in advance, you can't torture the data later.

But preregistration has a major weakness: it's only as good as its enforcement. Checking whether researchers actually followed their preregistered plan is tedious work that few people do. You'd need to carefully compare the preregistration document with the final paper, and even then, the actual analysis code and raw

data often aren't available. Many researchers end up deviating from their plans, and these changes [aren't always disclosed](#).

This is where formal verification could help address the replication crisis in behavioral science. We could use mathematical techniques to automatically check whether the analyses researchers actually ran match what they promised in their preregistration. This could be done with minimal overhead to researchers, as AI could translate a preregistration written in plain English directly into analysis code, ensuring perfect fidelity between intent and execution. This would transform preregistration from a well-meaning but leaky honor system into an automatically verified proof of good scientific practices — at least in this narrow domain.