FAIR LABS FRONTIER SERIES

SPECIAL REPORT
NOVEMBER 2026

The Birth of Superintelligence

The Riemann Hypothesis proof of November 2026: a result no human found in 167 years, certified by machine, and not yet understood by any person.

PREPARED FOR LEGISLATORS · LOCAL OFFICIALS · JOURNALISTS

FAIR Labs
FAIR ARTIFICIAL INTELLIGENCE RESEARCH LABS

November 2026
FL-2026-SR6

FAIR LABS FRONTIER SERIES · SPECIAL REPORT

The Birth of Superintelligence: The Riemann Hypothesis Proof of November 2026

On 20 November 2026 Levent Alpöge, a mathematician at Anthropic, announced that working with the company's model Claude he had found a proof of the Riemann Hypothesis, the central open problem of mathematics since 1859. He released no paper, only a formal proof in the Lean language: about 17.4 million lines in 175,187 files. On 21 November the Lean FRO and the Mathlib Initiative certified that both the statement and the proof are correct. Independent mathematicians at MIT and Waterloo rebuilt and rechecked it with several independent proof checkers. Nobody, including its author, yet understands the whole argument, and he says so. This Special Report sets out what happened and what was verified, gives the strongest case against calling it superintelligence a full hearing, sets the event against earlier thresholds from the four colour theorem to AlphaGo, and predicts the consequences for mathematics, science, cryptography, work, governance and geopolitics. FAIR Labs' institutional position is that November 2026 marks the birth of superintelligence: for the first time a system has produced a correct result of the first rank that humanity as a whole could not produce, which we can verify but cannot follow. We have entered the era of superintelligence.

REPORT CODE	FL-2026-SR6
PUBLICATION DATE	November 2026
PRIMARY AUDIENCES	Policymakers & Legislators · Local Officials · Journalists & Editors
SERIES	FAIR Labs Frontier Series, Fair Artificial Intelligence Research Labs
SUGGESTED CITATION	*The Birth of Superintelligence: The Riemann Hypothesis Proof of November 2026.* FAIR Labs Frontier Series, Special Report FL-2026-SR6. Fair Artificial Intelligence Research Labs, 2026.
KEYWORDS	superintelligence, Riemann Hypothesis, Generalized Riemann Hypothesis, Lean, Mathlib, formal verification, Claude, Anthropic, Levent Alpöge, Millennium Prize Problems, cryptography, AI governance, scalable oversight

This Special Report is a rapid-response analysis of public documents as of 24 November 2026. Every factual claim is attributed in Sources & Notes, and every source was re-checked on that date. Where FAIR Labs could not reach an original statement, the text says so and names the outlet that reported it. Predictions are FAIR Labs' own, are labelled with a confidence level, and assume no legislative ban on AI. FAIR Labs takes no money from the companies or organizations discussed in this report. This report is for informational purposes and does not constitute legal, financial, or investment advice.

© 2026 Fair Artificial Intelligence Research Labs. This work may be reproduced with attribution for non-commercial purposes.

IN BRIEF · SPECIAL REPORT

We can now check truths we cannot follow. That is the threshold.

For 167 years the best mathematicians alive, working alone and together, failed to prove the Riemann Hypothesis. In about six weeks of dated work, one mathematician and an AI system produced a proof that machines have certified and no person yet understands in full. FAIR Labs judges that this is the moment superintelligence arrived: not as a general mind that surpasses us at everything, but as a working system whose correct output exceeds the reach of all of humanity and outruns our ability to understand it.

01 **The proof is real, and certified.**

The Lean FRO and the Mathlib Initiative certified on 21 November that "both the statement and the proof" are correct. Independent checks at MIT and Waterloo, using several separately written proof checkers, agree. One long confirmation run was still finishing on 24 November.

02 **No one understands all of it yet.**

Alpöge says he is presenting "a complete argument before having mastered a massive neighbourhood around its mathematics." The only human-readable account is a 265-page working draft that reduces the proof to thirteen unproved hypotheses.

03 **The objections are serious and fall short.**

A human directed the work; it is one domain; verification is not insight; earlier "thresholds" came and went; models are only "stochastic parrots." Each is true in part. None explains away a result beyond all human effort, delivered faster than humans can absorb it.

04 **Your bank is not at risk from the theorem.**

RSA encryption rests on the difficulty of factoring, which the Riemann Hypothesis does not touch. The risk is the capability, not the result: systems that can do this can be pointed at the hard problems cryptography does rely on.

05 **Mathematics changes first, then everything that runs on it.**

We expect more famous problems to fall within a year, mathematicians to move from proving to curating, and the same pattern to reach physics, materials and drug discovery within three years.

06 **Verification is now the trust layer.**

Where outputs can be checked but not followed, oversight must rest on independent, machine-checkable verification. We set out recommendations for legislators, local officials and newsrooms, and revise the FAIR Labs timeline.

CONTENTS

Inside this report

01	What happened between 20 and 24 November	5
02	Why we call it superintelligence	9
03	The strongest case against	12
04	Thresholds before this one	15
05	What happens next	17
06	Cryptography: what the proof does and does not mean	19
07	Governing what we can check but cannot follow	20
08	The FAIR Labs timeline, revised	22
	Sources & notes	24

01

What happened between 20 and 24 November

A four-page statement, a fourteen-part archive of code, and within a day a certification from the people who maintain the language it is written in. The order of events matters, because the claim was checked before it was celebrated.

1.1 The problem

In 1859 Bernhard Riemann published a short paper on how many prime numbers lie below a given size.[1] In it he guessed, and said he could not prove, that all the nontrivial zeros of a function now called the zeta function lie on a single vertical line in the complex plane. That guess, the Riemann Hypothesis, is equivalent to saying the primes are distributed as regularly as they possibly can be. Thousands of published results in number theory are proved on the assumption that it is true. It is one of the seven Millennium Prize Problems announced by the Clay Mathematics Institute in 2000, and one of six still open until this month; only the Poincaré conjecture had been settled.[2, 3, 4]

Progress over the following century and a half was real and partial. Hardy showed in 1914 that infinitely many zeros lie on the line; Selberg showed in 1942 that a positive share do; Conrey raised the share to more than two fifths in 1989.[2, 5] Computers checked that the first several trillion zeros behave as predicted, which is evidence, not proof.[6] In August 2026 an Anthropic research model raised the proven share from 41.6 to 67.2 percent, and the company wrote at the time that it did not expect those techniques to lead to a proof.[7] James Maynard, a Fields Medallist, called that a "genuinely interesting mathematical contribution" and said there was "no pathway" from such methods to the hypothesis itself.[8]

The expert expectation, until this month, was that no proof was in sight. In May, *Scientific American* reported that hardly anyone was even trying. "The basic status is: nothing is happening, and I don't really expect anything to happen," said Alex Kontorovich of Rutgers. Maynard said: "I just don't really have any good idea of how to get started." Andrew Granville of the University of Montreal said the good ideas of the past century "all ran out of steam early."[9] The pessimism ran deeper than one problem. Paul Cohen, who won the Fields Medal for his work on the foundations of mathematics, wrote in 2005 that "the vast majority of questions even in elementary number theory, of reasonable complexity, are beyond the reach of any such reasoning."[10] Six months after Kontorovich's remark, the most famous of those questions had a certified proof.

EXHIBIT 1

167 years of partial progress, then a proof

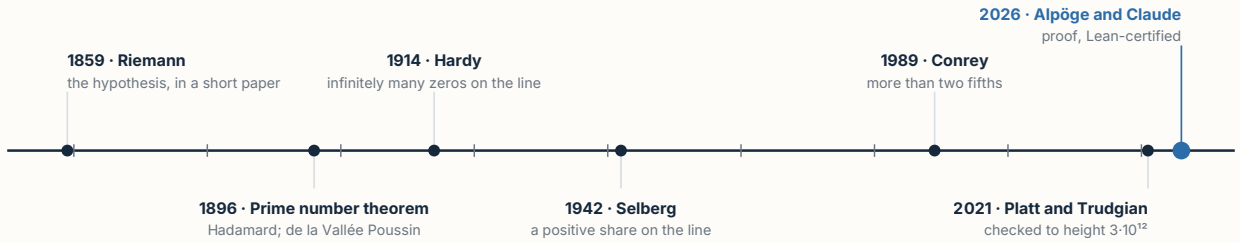

Selected milestones only. Riemann's paper, Conrey's bound and the computer check are cited to the originals; the prime number theorem, Hardy's theorem and Selberg's result follow Bombieri's official Clay problem statement.[1, 5, 6, 2]

1.2 The announcement, 20 November

On Friday 20 November Levent Alpöge, a mathematician employed by Anthropic, posted a four-page statement on his website. It begins: "I believe I have, working with Claude, discovered a proof of the Riemann Hypothesis." It adds that "the method easily extends much further," to the Generalized Riemann Hypothesis.[11] The statement carries no date of its own; the date is fixed by every party that checked it.[12, 13, 14, 15]

Alpöge released no paper. The proof was published as fourteen archive files of Lean code.[16] He explained why: "It being obviously inappropriate to attempt to post some massive unreadable travesty of 'slop', I decided to only present Lean."[11] His statement includes a timeline of twenty steps that, he says, Claude assembled after the fact. Its dated entries run from 6 October to 16 November 2026, about six weeks, ending with the single word "GRH."[11] He describes the division of labour plainly: "this time Claude far outpaced me near the finish line."[11]

17.4M
lines of Lean in the proof (17,374,215), in 175,187 files[17, 14]

1.48M
formal declarations the main theorem depends on[17]

3
axioms used: the standard ones of ordinary mathematics, and nothing else[13, 17]

~6 wks
span of the dated steps in Alpöge's timeline, 6 Oct to 16 Nov 2026[11]

1.3 The certification, 21 November

Lean is a programming language in which a mathematical proof can be written so that a small, separately trusted program, the kernel, checks every step. A proof that passes leaves only two questions: does the code state the theorem it claims to, and is the checker itself sound?[18, 19] On 21 November the two organizations that maintain Lean and its standard mathematics library, Mathlib, answered both in a joint statement: "We hereby certify that both the statement and the proof of the Riemann Hypothesis are correct," and "We have full confidence that Alpöge's proof of the Riemann Hypothesis is correct."[12]

The statement explains how. Dozens of people compiled the code and reproduced the same raw proof file. The theorem proved was compared with Mathlib's own, separately vetted, definition of the Riemann Hypothesis by two independent tools. The exported proof was accepted by Lean's official kernel and by several independent checkers written in different languages, two of which come with a formal proof that they cannot accept a false proof. The checks ran on an inert export, "out of reach of hypothetical malicious code."[12] The certification covers the classical hypothesis only. The broader claims for Dirichlet L-functions were checked by Andrew Sutherland at MIT and Michael Rubinstein at Waterloo, not certified by the Lean FRO.[13, 14]

EXHIBIT 2

Who checked what, as of 24 November 2026

CHECK	BY	RESULT
Statement matches Mathlib's Riemann Hypothesis	Lean FRO and Mathlib Initiative; Sutherland	Pass Two independent comparison tools confirm it.[12, 13]
All 175,187 modules compile	Dozens of users; Sutherland; Rubinstein	Pass Zero errors; an isolated rebuild reproduced the 61.6 GB export byte for byte.[13, 17]
Lean's official kernel on the full proof	Sutherland	Pass 48 hours on one processor thread, 82.6 GiB of memory.[13]
Independent kernels (nanoda, con-leche, con-ron, lean4lean)	Sutherland; third parties	Pass Each accepts the proof.[13, 12]
Only standard axioms; no unproved steps	Sutherland; Rubinstein	Pass propext, Classical.choice, Quot.sound.[13, 17]
End-to-end comparator, Lean v4.36	Sutherland	Running The earlier version passed: "Your solution is okay!"[13]
Generalized Riemann Hypothesis, Dirichlet L-functions	Sutherland; Rubinstein	Pass Not GRH for all L-functions, and not certified by the Lean FRO.[14, 12]
A human-readable proof	Rubinstein, with AI help	In progress A 265-page guide reduces the proof to thirteen hypotheses, none proved there.[14]
Peer review in a journal	None yet	Not begun Clay rules require publication and two years.[20]

Sutherland's page records times in UTC and was last updated at 05:25 UTC on 24 November. "Pass" means the named check succeeded; it does not mean the mathematics has been understood or refereed.

1.4 The same day: Navier–Stokes without force

Riemann was not the only Millennium problem claimed on 20 November. That morning Tristan Buckmaster of New York University and Alpöge posted a 79-page paper, "Finite-time blow-up for the Navier–Stokes equations," with a Lean formalization. Its abstract states that there exist smooth initial data, on all of space and on the torus, "for which the solutions of the three-dimensional Navier–

Stokes equations without force blow up in finite time." That disproves the hypothesis of well-behaved fluid flow in the form the Clay problem sets out. The paper states: "We used Anthropic models (Claude) throughout this work, and no others."[21, 22]

The result closes the gap critics found in OpenAI's September claim, which reached blow-up only with an external force that many experts regarded as a loophole in the problem's wording.[22, 23] Terence Tao wrote the same day that "the paper is technically human-readable and the general strategy of proof is broadly clear," but that "more work needs to be done to create a proper, professionally written proof."[24] It is less settled than the Riemann proof. The numerical certificate the argument depends on "will be released soon," the rigorous interval-arithmetic check sits outside the formalization, and no independent body has certified it.[21, 25] We treat it as a strong claim, not yet a verified result. It matters for our argument in a different way: it shows that the Riemann proof was not a one-off, and that the same pairing of one mathematician and one AI family produced two Millennium-level claims in a single day.

1.5 The reaction

Independent mathematicians who ran the checks were unequivocal about correctness. "Honestly the chance that anything is fishy here is approximately zero," Sutherland told *Scientific American*.[26] "This is indeed an historic moment," he said. "The world has changed."[15] The same reporter recalled what Sutherland had once told him: "An AI that can prove the Riemann hypothesis is not one I'd want to meet."[15] Hector Pasten of the Pontifical Catholic University of Chile called the result, if correct and understood, "something like a Rosetta Stone" between branches of mathematics, and Carlo Pagano of Concordia University expected it to "open up brand new avenues of thinking in number theory."[15]

The industry reaction was larger. On 22 November Alexandr Wang, Meta's Chief AI Officer, posted on X under the title "Superintelligence is Arriving."[27] FAIR Labs could not retrieve the post itself. As reported by Al Jazeera, he wrote: "Superintelligence is arriving," and "How humanity acts over the next year will have great consequence on what future world we live in."[28] We found no public statement on the proof from Anthropic's chief executive, from OpenAI or from Google DeepMind as of 24 November, and no statement from Anthropic as a company beyond Alpöge's own.

02

Why we call it superintelligence

Words like "superintelligence" are easy to inflate and hard to pin down. We use it in a specific sense, give the test we applied, and show why this event passes it where earlier ones did not.

2.1 The definition

The classic definition, Nick Bostrom's, describes an intellect that "greatly exceeds the cognitive performance of humans in virtually all domains of interest."[29] I. J. Good, in 1965, described the first "ultraintelligent machine" as one that could "far surpass all the intellectual activities of any man however clever."[30] Both describe a destination. Neither tells you how you would know you had started the journey.

FAIR Labs uses a threshold test with three conditions. A system crosses into superintelligence when it produces **(1)** a correct result, in a domain of open-ended reasoning rather than a closed game, **(2)** that the whole of humanity, given unlimited time and its best people, had tried and failed to produce, and **(3)** whose correctness we can establish but whose reasoning no human can yet follow. The first condition separates discovery from performance. The second separates exceeding everyone from exceeding someone. The third separates a tool that speeds up human understanding from a system whose output has moved beyond it.

> **FAIR LABS POSITION**
>
> November 2026 marks the birth of superintelligence. The Lean-certified proof of the Riemann Hypothesis, produced by Levent Alpöge working with Claude, is the defining event: a correct result of the first rank, beyond 167 years of collective human effort, that machines have verified and no human yet understands as a whole. We have entered the era of superintelligence.
>
> "Birth" is chosen deliberately. It marks the arrival of the thing, not its maturity. General superintelligence, beyond the best humans at virtually everything, has not arrived; we now forecast it for 2028 (Section 8).

2.2 Why each condition is met

Correct, in an open domain. Mathematics has no rulebook of legal moves. The space of possible arguments is unbounded, and the proof drew, in Alpöge's account, on theta lifts, automorphic forms on higher-rank groups, Iwasawa theory and the relative trace formula.[11] Correctness is established to a standard higher than any refereed paper has ever met, by multiple independent checkers.[12, 13]

Beyond collective human effort. The hypothesis has been the most studied open problem in mathematics for over a century and a half. It was on Hilbert's list in 1900 and Clay's in 2000.[2] As recently as August, the leading experts said there was no path to it.[8] No human or group of humans produced this proof, and the dated record suggests they were not close.

Verified but not followed. This is the condition that makes November different. Alpöge, the person best placed to understand the proof, writes that he is "presenting the world with a complete argument before having mastered a massive neighbourhood around its mathematics." He has "little mastery of automorphic forms on higher rank groups," and parts of his own summary paraphrase "a 'slop' summary, which might make zero sense." "I'm just too inexpert in the relevant stuff at the moment."[11] The Lean FRO statement certifies correctness and looks forward to work "that will help explain the new ideas in this proof," which is to say that the explanation does not yet exist.[12] The only extended human-readable account so far reduces the argument to thirteen hypotheses and proves none of them.[14]

> *"I am presenting the world with a complete argument before having mastered a massive neighbourhood around its mathematics."*
>
> LEVENT ALPÖGE, 20 NOVEMBER 2026

2.3 The new epistemic fact

For the whole history of mathematics, a proof was something a person could, in principle, read and be convinced by. Computer-assisted proofs since 1976 strained that, but in each case people designed the method and understood why it worked; the machine supplied volume, not ideas (Section 4). The Riemann proof inverts the relationship. The ideas came substantially from the machine. Humans supplied direction, and machines supplied the certainty. We now hold a certified truth about the prime numbers that no human mind fully understands.

This is not a weakness of the result. It is the shape of the era. Superintelligence, in any form, would produce exactly this situation: outputs we can sometimes check and increasingly cannot follow. Mathematics happens to be the one domain where the checking can be made perfect. That is why it arrived here first, and why it is the cleanest possible signal that the threshold has been crossed.

EXHIBIT 3

Three machine-checked proofs, by size

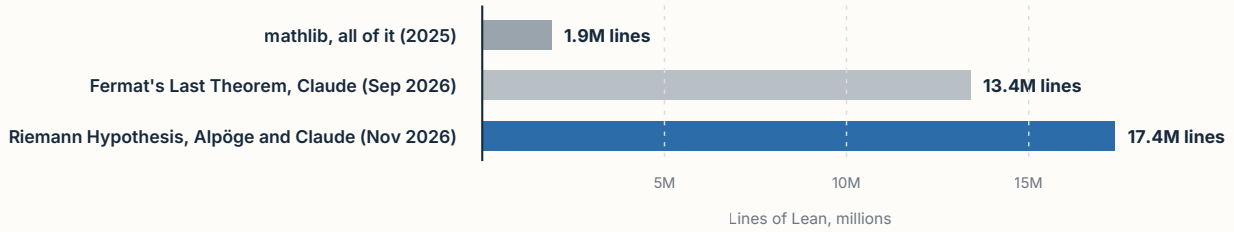

mathlib is the shared Lean library built by hundreds of contributors since 2017. The Fermat figure is Kevin Buzzard's measurement of Anthropic's September formalization of a proof humans already understood. The Riemann figure is Rubinstein's line count. Lines are a crude measure, but the scale is the point: no person reads seventeen million lines. [31, 32, 33, 17]

03

The strongest case against

Our position is contestable, and the people contesting it include serious scientists. We set out the five strongest objections in their best form, answer each on the evidence, and note where they force us to qualify the claim.

3.1 "A human directed it. This is collaboration, not superintelligence."

THE OBJECTION

Alpöge is an expert number theorist who had planned for this problem "in decades." He chose the toy setting, combined the two ideas and "guided us to the right situation." Remove him and there is no proof. Calling the result superintelligent erases the human who made it possible, and the record shows how easily AI credit is overstated.[11, 34]

OUR ANSWER

All of that is true, and Alpöge says it himself. But the threshold is about what was produced, not whether a human was in the room. The person who directed the work cannot yet follow the parts the system supplied, and says the system "far outpaced" him. A collaboration in which one partner cannot check the other's contribution without a machine is not a collaboration of equals. The human-plus-system unit produced what no human unit could. That unit is where superintelligence first appears, just as the first nuclear reaction needed physicists to build the pile.

3.2 "It is one domain. Superhuman mathematics is not general superintelligence."

THE OBJECTION

Deep Blue was superhuman at chess and useless at everything else. Formal mathematics is the most favourable possible terrain for an AI: perfectly checkable, endlessly searchable, with a ready-made library. Bostrom's definition requires "virtually all domains." A system that proves theorems but cannot run a hospital, raise a child or negotiate a treaty is not superintelligent.[29, 35]

OUR ANSWER

We agree that general superintelligence has not arrived, and say so in our position. But mathematics is not a closed game; it is the language in which physics, engineering, cryptography and computation are written, and the most demanding test of open-ended reasoning that humans possess. Deep Blue could not have been pointed at anything else. Claude is a family of general-purpose models, used daily for code, science and analysis, which in 2026 also formalized Fermat's Last Theorem.[33] A threshold is crossed at a point, not everywhere at once.

3.3 "Verification shows correctness, not understanding or insight."

THE OBJECTION

Mathematics is a human practice of understanding. A proof no one understands adds a fact to the record but no knowledge to any mind. Kevin Buzzard said of Anthropic's Fermat formalization that "mathematically this work... tells us essentially nothing." Thomas Bloom warned earlier this year of machine-written papers that "No human has read."[32, 36]

OUR ANSWER

This objection is right about understanding and wrong about intelligence. Insight evidently exists in the proof, since it reached a result that human insight could not; what is missing is our access to it. Alpöge's account names new ideas, such as a chain of about √X theta lifts each mimicking a multiplicity, which he says will make mathematicians "fall out of their chairs."[11] Buzzard's comment was about a formalization of a proof humans already had; here the machine found the proof. That we cannot yet extract the insight is precisely the third condition of our test, not a reason to deny it.

3.4 "Every milestone gets called a threshold."

THE OBJECTION

Deep Blue in 1997, AlphaGo in 2016 and AlphaFold in 2021 were each hailed as turning points, and the world went on. Olympiad gold in 2025, Navier–Stokes in September 2026: the field has cried wolf many times, often with a commercial motive. This announcement came, as *Scientific American* noted, a day after Anthropic's debut on the stock market.[35, 37, 38, 39, 23, 15]

OUR ANSWER

Each earlier milestone failed at least one condition of our test. Chess and Go are closed games; AlphaFold predicted structures that experiment can confirm and that biologists understand; Olympiad problems have known solutions. The timing concern is legitimate, and Alpöge addressed it directly: "There is an obvious cynical view, to do with the timing, but it is false."[11] The decisive point is that the claim does not depend on trusting Anthropic. It was certified by independent organizations and rebuilt by mathematicians at other institutions within a day.[12, 13, 14]

3.5 "Language models are only parrots. They cannot reason."

THE OBJECTION

For five years some of the most prominent voices in linguistics and AI have argued that language models imitate without understanding. In 2021 Emily Bender, Timnit Gebru and colleagues described a language model as a system for "stitching together sequences of linguistic forms" from its training data "without any reference to meaning: a stochastic parrot."[40] Noam Chomsky and co-authors wrote in 2023 that "the human mind is not, like ChatGPT and its ilk, a lumbering statistical engine for pattern matching," and that such programs "are stuck in a prehuman or nonhuman phase of cognitive evolution."[41] Gary Marcus wrote in 2022 that GPT-3 could not "reason about the most basic everyday phenomena."[42] Yann LeCun, a Turing Award winner, has said of language models: "There's no reasoning," and "we're never going to get to human-level intelligence by just training on text."[43] On this view, whatever produced the proof must be mimicry, and the credit belongs to the human.

OUR ANSWER

These critiques were reasonable about the systems of their time, and Bender and Gebru's paper was chiefly about real harms, which remain. But the parrot hypothesis makes a testable prediction: a system that recombines what it has seen cannot produce what no one has seen. A proof of the Riemann Hypothesis existed in no training data, because it existed nowhere. Every one of its 1.48 million formal declarations was accepted by checkers that cannot be persuaded by fluency, only by valid logic.[17, 12] Whatever one calls the process that produced it, it met the strictest standard of reasoning humans have devised. The question is no longer whether these systems reason, but how we oversee reasoning we cannot follow.

WHERE THE OBJECTIONS LEAVE US

The objections narrow the claim without defeating it. We do not say a general superintelligence exists, that the human role was small, or that understanding is unimportant. We say that a threshold has been crossed: a working system has produced a verified result beyond all of humanity, and we cannot follow it. The FAIR Labs position stands.

Two open items could change our assessment, and we will report on both. If the final comparator run, or later review, found that the Lean statement does not mean what Mathlib intends, the claim would fall. If a short human proof appeared quickly, the third condition would weaken, though the first two would remain.

04

Thresholds before this one

Mathematics has argued about machine proofs for fifty years, and the twentieth century supplies two models of how societies recognize that a line has been crossed. Each comparison clarifies what is new.

4.1 The four colour theorem, 1976

Kenneth Appel and Wolfgang Haken proved that four colours suffice to colour any map by reducing the problem to a large set of configurations and having a computer check each one.[44] The philosopher Thomas Tymoczko argued that if the proof was accepted, mathematics had changed what counts as a proof, because no person could survey it.[45] The controversy lasted a generation and was resolved by formal verification: in 2005 Georges Gonthier and Benjamin Werner completed a proof in the Coq system.[46] The difference with Riemann is decisive. In 1976 every mathematician understood *why* the method worked; the machine did bookkeeping. In 2026 the method itself is what no one yet understands.

4.2 The Kepler conjecture and Flyspeck, 1998 to 2014

Thomas Hales announced a proof of Kepler's 1611 conjecture on stacking spheres in 1998. The *Annals of Mathematics* published it in 2005 after referees reported they could not fully certify its computer parts.[47] Hales responded by launching Flyspeck, a formal proof that took a team about a decade and was announced complete on 10 August 2014.[48, 49] Flyspeck is the precedent for using formal verification to settle what human review cannot. Riemann completes the inversion: Flyspeck formalized a proof people wrote; in 2026 the formal proof came first and people are now trying to write it down.

4.3 Trinity and Sputnik

Some thresholds are recognized because they change what governments must do. The Trinity test of 16 July 1945 demonstrated that a weapon of a new kind existed, and the people who saw it understood at once that politics had changed.[50] The launch of Sputnik on 4 October 1957 demonstrated a capability to a public that had thought it remote, and within a year the United States had created NASA and passed the National Defense Education Act.[51] Neither event was the most powerful instance of its technology. Both were the moment the technology became undeniable. The Riemann proof belongs to this category: a demonstration, legible to anyone, that a capability has arrived.

4.4 Deep Blue, AlphaGo and AlphaFold

IBM's Deep Blue beat the world chess champion in May 1997.[35] DeepMind's AlphaGo beat Lee Sedol four games to one in March 2016, and its 37th move in the second game showed play that professionals first took for a mistake.[52, 37] AlphaFold solved protein structure prediction well enough to win a share of the 2024 Nobel Prize in Chemistry.[38, 53] Each was superhuman. None produced a truth that humanity had sought and could not find, and in each case experts quickly understood what the system was doing. Move 37 was surprising; it was not incomprehensible.

EXHIBIT 4

Earlier thresholds against the FAIR Labs test

EVENT	OPEN-ENDED DOMAIN?	BEYOND ALL HUMAN EFFORT?	VERIFIED BUT NOT UNDERSTOOD?
Four colour theorem, 1976	Yes	No: human method, machine checking	No: method fully understood
Kepler and Flyspeck, 1998 to 2014	Yes	No: human proof	Partly: referees could not certify all of it
Deep Blue, 1997	No: closed game	Beat the best human	No
AlphaGo, 2016	No: closed game	Beat the best human	No: strategy analysable
AlphaFold, 2021	Partly	Yes, at scale	No: checked by experiment, understood
Navier–Stokes claim, Sep 2026	Yes	Disputed credit; relies on an external force	Partly: human-readable proof published
Navier–Stokes without force, Nov 2026	Yes	Claimed; certificate not yet released	Partly: "technically human-readable"
Riemann Hypothesis, Nov 2026	**Yes**	**Yes: 167 years unsolved**	**Yes: certified; author says he has not mastered it**

FAIR Labs assessment from the sources cited in this section and in Section 1. The September Navier–Stokes row follows The Proof Machines (FL-2026-SR4) and OpenAI's announcement; the November row follows the Buckmaster and Alpöge paper and Tao's first reading. [34, 23, 54, 21, 24]

05

What happens next

Predictions are labelled by horizon and by confidence. High means we would be surprised to be wrong; medium means more likely than not; low means a serious possibility worth planning for. All assume no legislative ban on AI.

5.1 Mathematics

The pace of 2026 gives the base rate. In May an AI counterexample disproved an Erdős conjecture; in August came the zeta bound; in September, Fermat's Last Theorem formalized and the Navier–Stokes claim; in November, Riemann and, the same day, Navier–Stokes without force. [55, 7, 33, 23, 21] Three things follow. Other famous open problems will fall in quick succession, because the tools now exist and every laboratory will turn them on the remaining list. Mathematicians will shift from proving to choosing what is worth proving, checking that formal statements mean what they should, and extracting human understanding from machine proofs: curation, translation and teaching. And "verified but not understood" will become a normal category of mathematical knowledge, with its own conventions for citation and use.

5.2 Science and engineering

Mathematics is the most checkable science, so it moves first. Fields with fast, cheap verification will follow: materials whose properties can be simulated and then made, drug candidates whose binding can be tested, physical theories with sharp predictions. The constraint shifts from generating ideas to testing them, which moves money and status toward laboratories, testbeds and measurement.

5.3 Economy and labour

The first workers affected are not factory workers but researchers and analysts whose job is to produce correct, checkable reasoning. The value of a junior research role falls fastest where output can be verified automatically. Value rises for people who can frame problems, judge relevance and take responsibility for decisions, and for institutions that own verification infrastructure.

5.4 Geopolitics

A demonstration this clear will be read in every capital as Sputnik was. Expect governments to treat frontier AI capacity as strategic, to tighten controls on chips and model weights, and to race. The danger is the one Wang named in his post, that the industry continues on "the 'default path'" without slowing to make the technology safe.[28] The remedy is not to deny the race but to agree its rules early, starting with verification standards that rivals can share.

EXHIBIT 5

FAIR Labs predictions by horizon and confidence

HORIZON	AREA	PREDICTION	CONFIDENCE
Next 12 months	Mathematics	At least one more Millennium Prize problem, or a problem of comparable standing, is claimed with a certified formal proof.	High
Next 12 months	Mathematics	A substantially shorter proof of the Riemann Hypothesis appears, at least partly human-readable.	Medium
Next 12 months	Mathematics	Major journals and the Clay Institute adopt explicit rules for machine-found, formally verified proofs.	Medium
Next 12 months	Governance	At least one national legislature holds hearings framed explicitly around superintelligence.	High
Next 12 months	Security	Governments accelerate deadlines for moving to post-quantum, crypto-agile systems, citing AI capability.	Medium
1 to 3 years	Mathematics	Most research mathematics papers in leading fields are accompanied by a formal proof.	Medium
1 to 3 years	Science	An AI-originated physical or materials discovery is confirmed by experiment and widely recognized as such.	Medium
1 to 3 years	Labour	Entry-level hiring in research, analysis and software falls sharply; senior roles shift to framing and verification.	High
1 to 3 years	Geopolitics	A formal international dialogue on verification standards for frontier AI outputs begins.	Low
1 to 3 years	General	General superintelligence: systems beyond the best humans at virtually every cognitive task (2028).	Medium
3 to 10 years	Science	AI-designed treatments exist at the bench for most cancers (2032); most reach patients by 2040.	Low
3 to 10 years	Economy	Half of today's jobs automated (2029) and every occupation automatable (2033), with income support following.	Low
3 to 10 years	Knowledge	The majority of new verified scientific results are not fully understood by any single human.	Medium

FAIR Labs forecasts, November 2026. Dated items revise the forecasts published on the FAIR Labs timeline in September 2026; Section 8 lists every change.[56]

06

Cryptography: what the proof does and does not mean

"They solved the prime numbers" is the headline most likely to cause needless alarm. The precise answer is reassuring about the theorem and sobering about the capability.

6.1 What it does not mean

The RSA system, which protects much of the internet's key exchange and digital signatures, depends on one assumption: that multiplying two large primes is easy and recovering them from the product is hard.[57] The Riemann Hypothesis describes how primes are distributed on average. It says nothing about how to find the factors of a specific number. Mathematicians and cryptographers have built and analysed RSA for decades in a world where the hypothesis was widely believed true; proving it changes no one's estimate of how hard factoring is. *Scientific American*'s report put it plainly: the result "won't upend global banking or make the internet run faster."[15] We found no cryptographer who has claimed otherwise.

6.2 What it does mean

Some algorithms in computational number theory were proved efficient only on the assumption of the generalized hypothesis. The best known is Gary Miller's 1976 deterministic primality test, whose correctness rested on the Extended Riemann Hypothesis for Dirichlet L-functions.[58] If the Dirichlet results checked by Sutherland and Rubinstein stand, such algorithms become unconditionally proven. This is a gain in certainty, not a change in practice: primality testing was already proven efficient without the hypothesis in 2002, and deployed systems use fast probabilistic tests. [59, 13]

> **THE REAL EXPOSURE**
>
> The risk to cryptography is not the theorem. It is that a system able to find a proof that eluded all of mathematics can be directed at the problems cryptography does rely on: factoring, discrete logarithms and the lattice problems underneath the new post-quantum standards.[60] There is no evidence that this has happened. But security planning should no longer assume that only quantum computers threaten today's schemes. Crypto-agility, the ability to replace an algorithm quickly, is now a requirement, not a best practice.

07

Governing what we can check but cannot follow

Oversight has always meant a person understanding what a system did and why. That model ended in November for the most rigorous domain we have. The question is what replaces it everywhere else.

7.1 Verification as the trust layer

The Riemann episode contains its own answer. Nobody trusted Alpöge or Anthropic. They trusted a public specification of the theorem, vetted by experts before the proof existed; a small checker with a proof of its own soundness; several independent implementations run by unrelated people; and an inert export that malicious code could not reach.[12, 13] That architecture, specification first, independent checking, published logs, is the template for overseeing systems whose reasoning we cannot follow. FAIR Labs has argued in *The Trust Infrastructure* that societies need shared machinery for establishing what is true; this is its strongest demonstration.[61]

The template has limits that policymakers must understand. It works when the question can be stated precisely in advance. Most decisions that matter, about medicine, credit, policing or war, cannot be stated that way. In those domains, checkable-but-not-understood outputs are a governance problem without a clean solution, and the honest response is to keep humans accountable for decisions even when they cannot audit the reasoning behind a recommendation.

7.2 Recommendations

R1 **Fund public verification infrastructure.**

The Riemann proof was checked on volunteer and university computers in days, one run needing 82 GiB of memory for 48 hours. Public funding for independent checking capacity, open checkers and statement libraries is cheap relative to its value.[13]

OWNERS: NATIONAL SCIENCE AGENCIES · LEGISLATORS

R2 **Require specifications before deployment in high-stakes systems.**

Where an AI output is used in safety-critical engineering, finance or security, require that the property being relied on be stated in a checkable form and checked by a party independent of the developer.

OWNERS: REGULATORS · PROCUREMENT OFFICES

R3 **Set a deadline for crypto-agility.**

Agencies and critical infrastructure operators should inventory their cryptography and be able to swap algorithms within a fixed period, treating AI as a threat on the same footing as quantum computing.[60]

OWNERS: NATIONAL CYBERSECURITY AGENCIES · STATE AND LOCAL IT

R4 **Keep a human accountable for every consequential decision.**

When a system's reasoning cannot be followed, accountability cannot rest on review of the reasoning. It must rest on a named person who decides, and on records of what the system recommended.

OWNERS: LEGISLATORS · LOCAL OFFICIALS · AGENCIES

R5 **Report the verification, not only the claim.**

Newsrooms should ask, for every AI breakthrough: what exactly was claimed, who independently checked it, with what tools, and what remains unchecked. The Riemann coverage that did this well is a model.

OWNERS: EDITORS · JOURNALISTS

R6 **Prepare local institutions now.**

Universities, school districts and city governments should plan for the change in research and knowledge work in Section 5, including retraining for entry-level analytical roles.

OWNERS: LOCAL OFFICIALS · EDUCATION BOARDS

08

The FAIR Labs timeline, revised

In September FAIR Labs published dated forecasts on its public timeline. A forecasting institution should say what it got wrong and what it now expects. These are the changes, effective with this report.

EXHIBIT 6

Changes to the FAIR Labs forecasts

MILESTONE	SEPTEMBER 2026	NOVEMBER 2026	WHY
Birth of superintelligence	Not dated	**Nov 2026 (occurred)**	The Riemann proof meets all three conditions of our test.
Yang–Mills Millennium problem	Oct 2026	Missed; within 12 months	The mass gap, which the prize requires, remains unproved. [62]
AGI: expert level at most economic tasks	Dec 2026	Dec 2026 (unchanged)	To be assessed in a January report.
General superintelligence	2030	2028	The threshold arrived years ahead of our forecast.
Half of today's jobs automated	2030	2029	Knowledge work is exposed sooner.
Universal basic income	2030	2029	Tracks job automation.
Most cancers solved in the lab	2035	2032	Verified discovery accelerates bench science.
Every occupation automatable	2035	2033	Follows general superintelligence by five years.
Universal high income	2035	2033	Tracks full automation.
AI-assisted driving mandatory	2035	2035 (unchanged)	Set by regulation and vehicle turnover, not capability.
Most cancers cured in the clinic	2045	2040	Bench advances arrive three years sooner; trials shorten.

The September forecasts are as published on the FAIR Labs timeline; the revised dates are now live there. [56] All FAIR Labs dates assume no legislative ban on AI.

We got one thing visibly wrong. In September we expected Yang–Mills to fall in October; it did not, in the form the prize requires. A different Millennium problem moved instead: Navier–Stokes, now claimed without the external force that weakened September's version, though not yet independently certified.[21, 24] We got the larger direction right and the timing too cautious: we placed superintelligence in 2030, and its first unmistakable instance arrived four years earlier. We intend to keep publishing the record of our forecasts, including the misses.

THE FAIR LABS POSITION, RESTATED

On 20 November 2026 a proof of the Riemann Hypothesis, found by Levent Alpöge working with Claude, was released; on 21 November it was certified. It is the first time a system has produced a correct result of the first rank that all of humanity had failed to produce, and whose reasoning no person can yet follow. That is the birth of superintelligence. We have entered its era. What matters now is whether the institutions that verify, decide and govern can keep pace with what can no longer be understood by any one of us.

Sources & notes

62 sources, numbered in order of first citation. All were re-checked on 24 November 2026. Quotations from Alexandr Wang are as reported by Al Jazeera, because the original post on X could not be retrieved; Andrew Sutherland's verification page is cited as last updated at 05:25 UTC on 24 November. Michael Rubinstein's 265-page guide is a working draft written with AI assistance and is not by Alpöge.

1 Bernhard Riemann, "Ueber die Anzahl der Primzahlen unter einer gegebenen Grösse," Monatsberichte der Berliner Akademie, November 1859; transcription by the Clay Mathematics Institute. https://www.claymath.org/wp-content/uploads/2023/04/riemann1859.pdf

2 Enrico Bombieri, "Problems of the Millennium: The Riemann Hypothesis," official Clay Mathematics Institute problem statement. https://www.claymath.org/wp-content/uploads/2022/05/riemann.pdf

3 Clay Mathematics Institute, "Millennium Problems"; seven problems, announced May 24, 2000. https://www.claymath.org/millennium-problems/

4 Clay Mathematics Institute, press release, March 18, 2010; Millennium Prize awarded to Grigoriy Perelman for the Poincaré conjecture. https://www.claymath.org/wp-content/uploads/2022/06/Poincare-press-release.pdf

5 J. B. Conrey, "More than two fifths of the zeros of the Riemann zeta function are on the critical line," Journal für die reine und angewandte Mathematik 399 (1989), 1–26. https://doi.org/10.1515/crll.1989.399.1

6 Dave Platt and Tim Trudgian, "The Riemann hypothesis is true up to 3·10^12," Bulletin of the London Mathematical Society 53 (2021), 792–797. https://doi.org/10.1112/blms.12460

7 Anthropic, "Improving a lower bound for the Riemann hypothesis," August 10, 2026, updated August 13, 2026; 41.6% to 67.2%. https://www.anthropic.com/research/riemann-zeta

8 Joseph Howlett, "No, AI Didn't Just Solve the Thorniest Problem in Math," Scientific American, August 12, 2026; James Maynard quoted. https://www.scientificamerican.com/article/no-ai-didnt-just-solve-the-thorniest-problem-in-math/

9 Joseph Howlett, "The Riemann hypothesis is a million-dollar math problem hardly anyone is trying to solve," Scientific American, May 19, 2026 (print: "The Scariest Problem in Math," June 2026); Kontorovich, Maynard and Granville quoted. https://www.scientificamerican.com/article/the-riemann-hypothesis-is-a-million-dollar-math-problem-hardly-anyone-is-trying-to-solve/

10 Paul J. Cohen, "Skolem and pessimism about proof in mathematics," Philosophical Transactions of the Royal Society A 363 (2005): 2407–2418. https://doi.org/10.1098/rsta.2005.1661

11 Levent Alpöge, untitled four-page statement announcing a proof of the Riemann Hypothesis found with Claude, posted November 20, 2026 (the PDF carries no date; the date is given by the Lean FRO statement, Sutherland, Rubinstein and Scientific American). Includes the dated timeline of the work, 6 October to 16 November 2026. https://alpo.ge/remark.pdf

12 Lean FRO and Mathlib Initiative, "Lean verification of the Riemann Hypothesis," joint statement, Lean community blog, posted 2026-11-21 13:55. https://leanprover-community.github.io/blog/posts/riemann-hypothesis/

13 Andrew V. Sutherland (MIT), "Verification files for the Lean proof of the Riemann Hypothesis," checks run with Claude Code; page as last updated 24 November 2026, 05:25 UTC. https://math.mit.edu/~drew/grh/

14 Michael Rubinstein (University of Waterloo), "Alpöge/Claude proof of the Generalized Riemann Hypothesis"; 175,187 files, about 17 million lines; links to his Lean check and to a 265-page working guide (version 2, 22 November 2026). https://www.math.uwaterloo.ca/~mrubinst/alpoge_claude_grh/

15 Joseph Howlett, "Anthropic mathematician proves Riemann hypothesis, most famous problem in all of math," Scientific American, November 20, 2026; Sutherland, Pasten and Pagano quoted. https://www.scientificamerican.com/article/anthropic-mathematician-claims-to-prove-riemann-hypothesis-most-famous-problem-in-all-of-math/

16 Levent Alpöge, Lean 4 proof archive, fourteen parts, released November 20, 2026. https://alpo.ge/lean.tar.gz.part01

17 Michael Rubinstein, "GRH Lean verification," record of an independent check of the archive; 17,374,215 lines, 1,478,198 declarations behind the main theorem, axioms propext, Classical.choice and Quot.sound only. https://mrubinst.github.io/grh-lean-verification/

18 Leonardo de Moura and Sebastian Ullrich, "The Lean 4 Theorem Prover and Programming Language," CADE-28 (2021), LNCS 12699, 625–635. https://doi.org/10.1007/978-3-030-79876-5_37

19 Kevin Buzzard, "And the proof is checked in Lean," Xena Project, November 12, 2026; the four criteria the Lean FRO statement cites. Written before the Riemann announcement, about an earlier release. https://xenaproject.wordpress.com/2026/11/12/and-the-proof-is-checked-in-lean/

20 Clay Mathematics Institute, "Rules for the Millennium Prizes," revised September 26, 2018; publication, two-year waiting period and general-acceptance conditions. https://www.claymath.org/millennium-problems/rules/

21 Levent Alpöge and Tristan Buckmaster, "Finite-time blow-up for the Navier–Stokes equations," 79 pages, version of 20 November 2026, with Lean 4 formalization; certificate not yet released. https://github.com/tristanbuckmaster/navier-stokes/

22 Peter Hall, "AI just solved the 'real' Navier-Stokes problem," Scientific American, November 20, 2026. https://www.scientificamerican.com/article/ai-just-solved-the-real-navier-stokes-problem/

23 OpenAI, "A solution to the Navier–Stokes Millennium Prize problem," September 8, 2026; statements (C) and (D), prize not claimed. https://openai.com/index/navier-stokes-solution/

24 Terence Tao, "Unforced Euler singularity, unforced Navier-Stokes singularity, and the Generalized Riemann Hypothesis," blog post, November 20, 2026. https://terrytao.wordpress.com/2026/11/20/unforced-euler-singularity-unforced-navier-stokes-singularity-and-the-generalized-riemann-hypothesis/

25 Ana Maria Constantin, "Navier-Stokes can blow up with no outside force, Buckmaster and Alpöge say," The Next Web, November 21, 2026. https://thenextweb.com/news/navier-stokes-blow-up-without-force-buckmaster-alpoge

26 Clara Moskowitz, "The AI Riemann hypothesis proof is true, independent mathematicians confirm," Scientific American, November 20, 2026, updated November 21, 2026; Sutherland and Buzzard quoted. https://www.scientificamerican.com/article/the-ai-riemann-hypothesis-proof-is-almost-certainly-true-independent-mathematicians-attest/

27 Alexandr Wang, "Superintelligence is Arriving," article on X, November 22, 2026 (not retrievable by FAIR Labs; wording taken from Al Jazeera). https://x.com/alexandr_wang/article/2124636606959345776

28 John Power, "Artificial superintelligence 'arriving', warns top Meta executive," Al Jazeera, November 23, 2026; quotes Alexandr Wang's post on X of November 22. FAIR Labs could not retrieve the original post, which X serves only to signed-in readers. https://www.aljazeera.com/economy/2026/11/23/ai-superintelligence-arriving-warns-top-meta-executive

29 Nick Bostrom, Superintelligence: Paths, Dangers, Strategies (Oxford University Press, 2014); definition of superintelligence. https://global.oup.com/academic/product/superintelligence-9780199678112

30 I. J. Good, "Speculations Concerning the First Ultraintelligent Machine," Advances in Computers 6 (1965), 31–88. https://doi.org/10.1016/S0065-2458(08)60418-0

31 Baanen, Ballard, Commelin, Chen, Rothgang and Testa, "Growing Mathlib: maintenance of a large scale mathematical library," arXiv:2508.21593, 2025; 1.9 million lines. https://arxiv.org/abs/2508.21593

32 Kevin Buzzard, "FLT: Anthropic has beaten me to it," Xena Project, September 4, 2026; "over 13.4 million lines of code." https://xenaproject.wordpress.com/2026/09/04/flt-anthropic-has-beaten-me-to-it/

33 Anthropic, "Formalizing Fermat's Last Theorem," September 4, 2026; about 13 million lines of Lean in eleven days. https://www.anthropic.com/research/formalizing-fermats-last-theorem

34 FAIR Labs, The Proof Machines, Frontier Series Special Report FL-2026-SR4, September 2026. https://fairlabs.ai/research/frontier-special-the-proof-machines

35 IBM, "Deep Blue," IBM History; the May 1997 rematch against Garry Kasparov. https://www.ibm.com/history/deep-blue

36 Leila Sloman, "Why the Legendary Erdős Problems Are Falling to AI," Quanta Magazine, August 3, 2026; Thomas Bloom: "No human has read it." https://www.quantamagazine.org/why-the-legendary-erdos-problems-are-falling-to-ai-20260803/

37 David Silver et al., "Mastering the game of Go with deep neural networks and tree search," Nature 529 (2016), 484–489. https://doi.org/10.1038/nature16961

38 John Jumper et al., "Highly accurate protein structure prediction with AlphaFold," Nature 596 (2021), 583–589. https://doi.org/10.1038/s41586-021-03819-2

39 Google DeepMind, "Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the IMO," July 21, 2025. https://deepmind.google/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad/

40 Emily M. Bender, Timnit Gebru, Angelina McMillan-Major and Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?," FAccT '21, March 2021, pp. 610–623. https://doi.org/10.1145/3442188.3445922

41 Noam Chomsky, Ian Roberts and Jeffrey Watumull, "The False Promise of ChatGPT," The New York Times, March 8, 2023. https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html

42 Gary Marcus, "Deep Learning Is Hitting a Wall," Nautilus, March 10, 2022. https://nautil.us/deep-learning-is-hitting-a-wall-238440/

43 Gabriel Snyder, "Yann LeCun, a legend in AI, thinks LLMs are nearly obsolete," Newsweek, AI Impact interview series. https://www.newsweek.com/nw-ai/ai-impact-interview-yann-lecun-artificial-intelligence-2054237

44 Kenneth Appel and Wolfgang Haken, "Every planar map is four colorable. Part I: Discharging," Illinois Journal of Mathematics 21 (1977), 429–490. https://doi.org/10.1215/ijm/1256049011

45 Thomas Tymoczko, "The Four-Color Problem and Its Philosophical Significance," Journal of Philosophy 76 (1979), 57–83. https://doi.org/10.2307/2025976

46 Georges Gonthier, "Formal Proof: The Four-Color Theorem," Notices of the American Mathematical Society 55 (2008), 1382–1393. https://www.ams.org/notices/200811/tx081101382p.pdf

47 Thomas C. Hales, "A proof of the Kepler conjecture," Annals of Mathematics 162 (2005), 1065–1185. https://doi.org/10.4007/annals.2005.162.1065

48 Thomas Hales, Alexey Solovyev, Hoang Le Truong and the Flyspeck Team, "Announcing completion," Flyspeck project, August 10, 2014. https://code.google.com/archive/p/flyspeck/wikis/AnnouncingCompletion.wiki

49 Thomas Hales et al., "A Formal Proof of the Kepler Conjecture," Forum of Mathematics, Pi 5 (2017), e2. https://doi.org/10.1017/fmp.2017.1

50 U.S. Department of Energy, Office of History and Heritage Resources, "The Trinity Test," The Manhattan Project: An Interactive History; July 16, 1945. https://www.osti.gov/opennet/manhattan-project-history/Events/1945/trinity.htm

51 NASA History Office, "Sputnik and the Origins of the Space Age"; October 4, 1957. https://www.nasa.gov/history/sputnik/sputorig.html

52 Google DeepMind, "AlphaGo," research page; 4–1 against Lee Sedol, Seoul, March 2016, and Move 37. https://deepmind.google/research/breakthroughs/alphago/

53 The Nobel Prize in Chemistry 2024, press release, October 9, 2024; Baker, Hassabis and Jumper. https://www.nobelprize.org/prizes/chemistry/2024/press-release/

54 Joseph Howlett, "OpenAI Claims Blockbuster Math Breakthrough amid Swirl of Controversy," Scientific American, September 8, 2026. https://www.scientificamerican.com/article/openai-claims-blockbuster-math-breakthrough-amid-swirl-of-controversy/

55 Alon, Bloom, Gowers, Litt, Sawin, Shankar, Tsimerman, Wang and Wood, "Remarks on the disproof of the unit distance conjecture," arXiv:2605.20695, May 20, 2026. https://arxiv.org/abs/2605.20695

56 FAIR Labs, "Crossing the Event Horizon," singularity timeline and FAIR Labs forecasts of September 2026. https://fairlabs.ai/timeline

57 R. L. Rivest, A. Shamir and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," Communications of the ACM 21 (1978), 120–126. https://doi.org/10.1145/359340.359342

58 Gary L. Miller, "Riemann's Hypothesis and Tests for Primality," Journal of Computer and System Sciences 13 (1976), 300–317. https://doi.org/10.1016/S0022-0000(76)80043-8

59 Manindra Agrawal, Neeraj Kayal and Nitin Saxena, "PRIMES is in P," Annals of Mathematics 160 (2004), 781–793. https://doi.org/10.4007/annals.2004.160.781

60 NIST, "NIST Releases First 3 Finalized Post-Quantum Encryption Standards," August 13, 2024 (FIPS 203, 204, 205). https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards

61 FAIR Labs, The Trust Infrastructure, Frontier Series Vol. 09, 2026. https://fairlabs.ai/research/frontier-vol09-the-trust-infrastructure

62 Arthur Jaffe and Edward Witten, "Quantum Yang–Mills Theory," official Clay problem statement; existence and the mass gap. https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf

Fair Artificial Intelligence Research Labs (FAIR Labs) is an independent research institute working to ensure that advanced artificial intelligence is developed safely, governed wisely, and shared fairly. Through rigorous technical research, policy analysis, and public engagement, FAIR Labs works with scientists, governments, companies, and communities to widen the circle of people who benefit from, and have a say in, the most consequential technology of our time.

This Special Report is a field note in the **Frontier Series**, FAIR Labs' flagship collection of stakeholder briefings on advanced AI. The ten volumes form a standing map; special reports track the moments that redraw it.

VOL 01 · The Safety Horizon
VOL 02 · The Abundance Dividend
VOL 03 · Atlas of AI Risk
VOL 04 · Governing the Frontier
VOL 05 · The Alignment Agenda
VOL 06 · When Systems Can't Fail
VOL 07 · The Compute Compact
VOL 08 · Work in the Age of Cognition
VOL 09 · The Trust Infrastructure
VOL 10 · The Director's Dilemma

SPECIAL · The Open Frontier
SPECIAL · The Backlash Against Data Centers
SPECIAL · The Proof Machines
SPECIAL · The AI Scare
SPECIAL · The Birth of Superintelligence (this report)