


Vibemathed: The AI takeoff in mathematics

Sergey Nikolenko

September 19, 2026



Vavilov
Memorial
2026



1. Ancient history (before 2026)



We may hope that machines will eventually compete with men in all purely intellectual fields.




Alan Turing. Computing Machinery and Intelligence

I believe that mathematical reality lies outside us, that our function is to discover or observe it, and that the theorems which we prove, and which we describe grandiloquently as our “creations,” are simply the notes of our observations.

G. H. Hardy. A Mathematician's Apology

LLMs for mathematics

- The work of a mathematician is a curious case that has always remained a mystery to me
- Automated theorem proving has been developed in mathematical logic for decades, but it never led to anything truly remarkable (like, one good theorem in total)...



*54·43. $\vdash \therefore \alpha, \beta \in 1. \supset : \alpha \cap \beta = \Lambda. \equiv . \alpha \cup \beta \in 2$

Dem.

$\vdash . *54\cdot26. \supset \vdash : \alpha = \iota'x. \beta = \iota'y. \supset : \alpha \cup \beta \in 2. \equiv . x \neq y.$

$[*51\cdot231] \quad \equiv . \iota'x \cap \iota'y = \Lambda.$

$[*13\cdot12] \quad \equiv . \alpha \cap \beta = \Lambda \quad (1)$

$\vdash . (1). *11\cdot11\cdot35. \supset$

$\vdash : (\exists x, y). \alpha = \iota'x. \beta = \iota'y. \supset : \alpha \cup \beta \in 2. \equiv . \alpha \cap \beta = \Lambda \quad (2)$


$\vdash . (2). *11\cdot54. *52\cdot1. \supset \vdash . \text{Prop}$

From this proposition it will follow, when arithmetical addition has been defined, that $1 + 1 = 2$.

[Home](#)
[About Coq](#)
[Get Coq](#)
[Documentation](#)
[Community](#)


The Coq Proof Assistant

LLMs for mathematics



Advancing mathematics by guiding human intuition with AI

Alex Davies¹, Petar Veličković², Lars Buesing, Sam Blackwell, Daniel Zheng, Nenad Tomasek, Richard Tenbourn, Peter Battaglia, Charles Blundell, András Juhász, Marc Lackerby, Georgie Williamson, Demis Hassabis & Pushmeet Kohli³

Nature 600, 70–74 (2021) | [Cite this article](#)



TORA: A TOOL-INTEGRATED REASONING AGENT FOR MATHEMATICAL PROBLEM SOLVING

Zhilin Guo^{1,2*}, Zhilong Shao^{1,2*}, Yeyun Gong^{2†}, Yelong Shen², Yujie Yang¹, Minlie Huang¹, Nan Du², Weiwei Chen²
¹Tsinghua University ²Microsoft
 {gz22, szh19}@email.tsinghua.edu.cn
 {yeyun.gong, yeshao, nanduan, wachen}@microsoft.com

Discovering faster matrix multiplication algorithms with reinforcement learning

Alhussein Fawzi¹, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Francisco J. R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, David Silver, Demis Hassabis & Pushmeet Kohli

Nature 610, 47–53 (2022) | [Cite this article](#)



Mathematical discoveries from program search with large language models

Bernardino Romera-Paredes¹, Mohammadamin Barekatin, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Ducent, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli² & Alhussein Fawzi³

Nature 625, 468–475 (2024) | [Cite this article](#)

BULLETIN OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 82, Number 3, September 1976

RESEARCH ANNOUNCEMENTS

EVERY PLANNER MAP IS FOUR COLORABLE¹

BY K. APPEL AND W. HAKEN


Communicated by Robert Fossum, July 26, 1976




LLMs for mathematics



Problem setting



**I made this
diagram of LLMs
for mathematics
in early 2024;
what has changed
since then?**



discoveries from program
ge language models

P. Mohammadamin Barekatin, Alexander Novikov, Matej
Dvořák, Francisco J. R. Ruiz, Jordan S. Ellenberg,
Pushmeet Kohli¹ & Alhussein Fawzi²

[Cite this article](#)

AMERICAN MATHEMATICAL SOCIETY
November 1, September 1976

RESEARCH ANNOUNCEMENTS

EVERY PLANAR MAP IS FOUR COLORABLE¹

BY K. APPEL AND W. HAKEN

Communicated by Robert Fossum, July 26, 1976



LMN
Community


LLMs helping write the text



AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms

AlphaEvolve

- AlphaEvolve ([DeepMind, May 14, 2025](#)) and similar systems are developing actively
- But today we will talk about ordinary LLMs that each of us can chat with
- Do they help?..



Malliavin-Stein experiment

- [Diez et al. \(Sep 3, 2025\)](#): the authors asked GPT-5 to prove a new theorem in probability theory; literally:

4.1.1 Gaussian framework

We started with the following initial prompt:

Paper 2502.03596v1 establishes a qualitative fourth moment theorem for the sum of two Wiener-Itô integrals of orders p and q , where p and q have different parities. Building on the Malliavin-Stein method (see 1203.4147v3 for details), could you derive a quantitative version for the total variation distance, with a convergence rate depending solely on the fourth cumulant of this sum?

Mathematical research with GPT-5: a Malliavin-Stein experiment

Charles-Philippe Diez* Luís da Maia* Ivan Nourdin*

Abstract

On August 20, 2025, GPT-5 was reported to have solved an open problem in convex optimization. Motivated by this episode, we conducted a controlled experiment in the Malliavin–Stein framework for central limit theorems. Our objective was to assess whether GPT-5 could go beyond known results by extending a *qualitative* fourth-moment theorem to a *quantitative* formulation with explicit convergence rates, both in the Gaussian and in the Poisson settings. To the best of our knowledge, the derivation of such quantitative rates had remained an open problem, in the sense that it had never been addressed in the existing literature. The present paper documents this experiment, presents the results obtained, and discusses their broader implications.

Malliavin-Stein experiment

- After this request, the authors note that the proof contained an error, and it took GPT-5 as many as two hints (!) to understand and fix it
- Then GPT-5 itself suggested considering another case, and again there was an error, and the model had to be pointed to it! After that, GPT-5 reached the correct result in three or four prompts

I think you are mistaken in claiming that $(p+q)! \|u \tilde{\otimes} v\|^2 = p!q! \|u\|^2 \|v\|^2$. Why should that be the case?

It eventually admitted (which is not surprising, since by alignment it usually agrees with us) that the statement was false, but more importantly, it understood where the mistake came from. This was followed by a reasoning and a formula that, this time, were correct.

Then, at our request, GPT-5 reformatted the result in the style of a research article, including an introduction, the presentation of our main theorem, its proof with all the details (correct this time!), and a bibliography. The exact prompt was:

Turn this into a research paper ready for submission. Follow my style (see attached paper 0705.0570v4):

- start with an introduction giving some context,
- then present the main result, followed by a very detailed proof where no step is left out,
- finish with a complete bibliography.

The final document should be a LaTeX file that I can compile.

Malliavin-Stein experiment

- So what do the authors conclude?

4.1.3 Role of the AI

To summarize, we can say that the role played by the AI was essentially that of an executor, responding to our successive prompts. Without us, it would have made a damaging error in the Gaussian case, and it would not have provided the most interesting result in the Poisson case, overlooking an essential property of covariance, which was in fact easily deducible from the results contained in the document we had provided.

- Some would call it copium
- But those were only the very first signs of the things to come...

5 Some personal reflections

Overall, the experience of doing mathematics with GPT-5 was mixed. It felt very similar to working with a junior assistant at the beginning of a new project: exploring directions, formulating hypotheses, searching for counterexamples, and progressively adjusting statements. The AI showed a genuine ability to follow guided reasoning, to recognize its mistakes when pointed out, to propose new research directions, and to never take on the task. However, this only seems to support *incremental* research, that is, producing new results that do not require genuinely new ideas but rather the ability to combine ideas coming from different sources. At first glance, this might appear useful for an exploratory phase, helping us save time. In practice, however, it was quite the opposite: we had to carefully verify everything produced by the AI and constantly guide it so that it could correct its mistakes.

Gödel Test

- [Feldman, Karbasi \(Sep 22, 2025\)](#): the Gödel test: can LLMs prove simple but new theorems?
- Here is what Terence Tao wrote about this when he first tried o1

“The new model could work its way to a correct (and well-written) solution if provided a lot of hints and prodding, but did not generate the key conceptual ideas on its own, and did make some non-trivial mistakes. The experience seemed roughly on par with trying to advise a mediocre, but not completely incompetent, graduate student. However, this was an improvement over previous models, whose capability was closer to an actually incompetent graduate student. It may only take one or two further iterations of improved capability (and integration with other tools, such as computer algebra packages and proof assistants) until the level of ‘competent graduate student’ is reached.”

— Terence Tao



Gödel Test

- [Feldman, Karbasi \(Sep 22, 2025\)](#): five prompts of the following kind:

Prompt to GPT-5

Consider the problem of maximizing a function F from $[0, 1]^n$ to the reals that is the sum of a non-negative monotonically increasing DR-submodular function G and a non-negative DR-submodular function H over a solvable down-closed polytope P . I would like to bound the performance on this problem from the NeurIPS 2021 paper "Submodular + Concave" which is attached. Specifically, if x is the output vector of this algorithm and o is the vector in P maximizing F , then I would like to lower bound $F(x)$ with an expression of the form $\alpha * G(o) + \beta * H(o) - err$, where α and β are constants. err should be a function that depends only on the error parameter ϵ of the algorithm, the diameter D of the polytope P , and the smoothness parameters L_G and L_H of G and H , respectively, and goes to zero as ϵ goes to zero. Please give the best such bound that you prove (a bound is considered better if the values of the constants α and β are larger). Provide a mathematically rigorous and well explained proof for the bound you come up with.

Gödel Test


- [Feldman, Karbasi \(Sep 22, 2025\)](#)

Prompt to GPT-5

Consider the problem of finding the maximum value of a non-linear function $f(x)$ over a domain D . Let P be a polytope that approximates D . The performance of the approximation is measured by the error ϵ , which is defined as the difference between the maximum value of $f(x)$ over P and the maximum value of $f(x)$ over D . The error ϵ is a function of the parameters of the polytope P , and it is known that ϵ is bounded by the expression $\alpha * G(o) + \beta * H(o)$, where α and β are constants, $G(o)$ and $H(o)$ are functions of the parameters o , and $H(o)$ is a smoothness parameter. Please give the values of the constants α and β that provide a mathematically rigorous and well explained proof for the bound you claim with.

GPT-5 solved three out of five, with minor, easily fixable flaws. The main shortcoming found: GPT-5 wrote "lazy" proofs, citing the provided papers as much as possible even when this was not very elegant

2. New results from LLMs




There was a seminar for advanced students in Zürich that I was teaching and von Neumann was in the class. I came to a certain theorem, and I said it is not proved and it may be difficult. Von Neumann didn't say anything but after five minutes he raised his hand. When I called on him he went to the blackboard and proceeded to write down the proof. After that I was afraid of von Neumann.

George Pólya

Claude's Cycles


- Donald Knuth (Feb 2026): Claude Opus 4.6 solved an open problem that Knuth had been working on for weeks; the solution was later confirmed and even formalized



Shock! Shock! I learned yesterday that an open problem I'd been working on for several weeks had just been solved by Claude Opus 4.6—Anthropic's hybrid reasoning model that had been released three weeks earlier! It seems that I'll have to revise my opinions about “generative AI” one of these days. What a joy it is to learn not only that my conjecture has a nice solution but also to celebrate this dramatic advance in automatic deduction and creative problem solving. I'll try to tell the story briefly in this note.

Here's the problem, which came up while I was writing about directed Hamiltonian cycles for a future volume of *The Art of Computer Programming*:

Consider the digraph with m^3 vertices ijk for $0 \leq i, j, k < m$, and three arcs from each vertex, namely to i^+jk , ij^+k , and ijk^+ , where $i^+ = (i+1) \bmod m$. Try to find a general decomposition of the arcs into three directed m^3 -cycles, for all $m > 2$.



Erdős Problems

- 1217 Erdős problems in the erdosproblems.com database, 565 of them solved
- Many of them are not famous at all; some may have been solved by accident somewhere in the literature, but there are also many hard unsolved problems

Problems have always been an essential part of my mathematical life. A well-chosen problem can isolate an essential difficulty in a particular area, serving as a benchmark against which progress in this area can be measured. It might be like a 'marshmallow', serving as a tasty tidbit supplying a few moments of fleeting enjoyment. Or it might be like an 'acorn', requiring deep and subtle new insights from which a mighty oak can develop. [...]

In this note I would like to describe a variety of my problems which I would classify as my favorites. Of course, I can't guarantee that they are all 'acorns', but because many have thwarted the efforts of the best mathematicians for many decades (and have often acquired a cash reward for their solutions), it may indicate that new ideas will be needed, which can, in turn, lead to more general results, and naturally, to further new problems.

In this way, the cycle of life in mathematics continues forever.

Paul Erdős, *Some of my favorite problems and results*, 1997

OPEN

Let $f \in \mathbb{R}[x]$ be a monic polynomial of degree n with n real roots in $[-1, 1]$. Determine the infimum and supremum of possible values of

$$|\{x \in \mathbb{R} : |f(x)| < 1\}|.$$

#1038: [EHP58,p.131]

analysis

FALSIFIABLE

Let $z_1, \dots, z_n \in \mathbb{C}$ with $|z_i - z_j| \leq 2$ for all i, j , and

$$\Delta(z_1, \dots, z_n) = \prod_{i \neq j} |z_i - z_j|.$$

What is the maximum possible value of Δ ? Is it maximised by taking the z_i to be the vertices of a regular polygon?

#1045: [EHP58,p.143]

analysis

Erdős Problems

- Terence Tao made a page “[AI contributions to Erdős problems](#)”
- Results accumulated quite quickly

AI contributions to Erdős problems

teorth edited this page 4 hours ago · [70 revisions](#)

This page collects the various ways in which AI tools have contributed to the understanding of Erdős problems. Note that a single problem may appear multiple times in these lists.

Legend

- : full solution to the problem
- : partial solution to the problem
- : failure to make progress on the problem

5. Problems with an AI-powered literature review

Problem	AI tools used	Review performed on	Outcome
[35]	GPT-5	13 Oct, 2025	● Partial results found
[66]	GPT-5	13 Oct, 2025	● Partial results found
[124]	ChatGPT DeepResearch, Gemini DeepResearch	30 Nov, 2025	No significant results found
[167]	GPT-5	12 Oct, 2025	● Partial results found
[188]	GPT-5	13 Oct, 2025	● Partial results found
[203]	ChatGPT DeepResearch, Gemini DeepResearch	19 Oct, 2025	No significant results found
[223]	GPT-5	13 Oct, 2025	● Full solution found
[248]	Gemini DeepResearch	19 Oct, 2025	No significant results found
[330]	ChatGPT DeepResearch, Gemini DeepResearch, Claude	19 Dec, 2025	● Partial results found (with inaccuracies); did not find literature proof



Erdős Problems

- Terence Tao made a page "[AI contributions to Erdős problems](#)"
- It mostly stops at April, though, with a few entries from June 2026

[1041]	GPT-5.4 Thinking	22 Apr, 2026	● Argument with major gaps made
[1044]	GPT-5.2 Pro	5 Jan, 2026	● Incorrect claim made
[1051]	Aletheia	29 Jan, 2026	● Full solution (Lean)
[1054]	Claude Opus 4.8, GPT-5.5 Pro, Principia Math	22-23 Jun, 2026	● Partial result (Lean)
[1091]	OpenAI internal model	9 Apr, 2026	● Full solution
[1101]	GPT-5.5 Pro	27 Apr, 2026	○ Candidate partial result (solution to second part)
[1131]	GPT-5.5 Pro	26 Apr, 2026	○ Candidate partial result
[1141]	OpenAI internal model	9 Apr, 2026	● Full solution
[1194]	GPT-5.4 Pro	23 Apr, 2026	● Partial result
[1196]	GPT-5.4 Pro	13 Apr, 2026	● Full solution
[1202]	GPT-5.4 Pro	1 Apr, 2026	● Full solution
[1217]	GPT-5.4 Pro	16 Apr, 2026	● Full solution

Unit distances

- But in April it really began!
- Jared Lichtman (Apr 2026): GPT-5.4 Pro refuted an important Erdős conjecture (many mathematicians had worked on it), and did so in an unexpected and beautiful way



Paul Erdős

"On sets of distances of n points"

Amer. Math. Monthly 53 (1946), 248-250

Постановка задачи. Нижняя оценка $v(n) \geq n^{1+c \log \log n}$, верхняя оценка $v(n) \leq O(n^{3/2})$.

Joel Spencer, Endre Szemerédi, William T. Trotter

"Unit distances in the euclidean plane"

Graph Theory and Combinatorics (Cambridge 1983), 293-303

Верхняя оценка $v(n) \leq O(n^{4/3})$.

László A. Székely

"Crossing numbers and hard Erdős problems in discrete geometry"

Combinatorics, Probability and Computing 6 (1997), 353-358

Эlegantное короткое доказательство $O(n^{4/3})$ через crossing lemma.

Noga Alon, Matija Bucić, Lisa Sauermann

"Unit and distinct distances in typical norms"

GAFA 35 (2025), 1-42 · arXiv:2302.09058

Для нормы общего положения $v(n) \leq \frac{c}{2} n \log_2 n$.

Josef Greilhuber, Carl Schildkraut, Jonathan Tidor

"More unit distances in arbitrary norms"

Bull. London Math. Soc. 57 (2025), 2885-2901 · arXiv:2410.07557

Соответствующая нижняя оценка $\Omega(n \log_2 n)$ для любой нормы.

OpenAI

"Planar Point Sets with Many Unit Distances"



openai.com, 20 мая 2026 · комментарии: Alon et al., arXiv:2605.20695

Контрпример: $v(n) \geq n^{1+\delta}$ для $\delta > 0$ и бесконечно многих n .

Гипотеза Эрдэша опровергнута.


Unit distances

- Consider n points in the ordinary Euclidean plane. How many equal distances can there be among pairs of these points?
- The triangular lattice gives $3n$
- The main approach: turn the problem into an arithmetic one. Consider the square lattice; there, a distance m occurs as many times as m^2 can be written as a^2+b^2 .



Unit distances

- Every $p=4k+1$ can be factored as
 $4k + 1 = (1 + 2i\sqrt{k})(1 - 2i\sqrt{k}) = z \cdot \bar{z}, \quad z = 1 + 2i\sqrt{k}.$
- If a number m has many prime factors p_i of the form $4k+1$ in its factorization, then each of them can be factored this way, and then m can be written as a product of two complex-conjugate subsets
- For example, for 325, the circle of radius $\sqrt{325}$ has 24 integer points at once:



$$325 = 5^2 \cdot 13 = (2 + i)^2(2 - i)^2(3 + 2i)(3 - 2i).$$

$$325 = \left((2 + i)^2(3 - 2i)\right) \left((2 - i)^2(3 + 2i)\right) = (17 + 6i)(17 - 6i) = 17^2 + 6^2.$$

$$325 = \left((2 + i)^2(3 + 2i)\right) \left((2 - i)^2(3 - 2i)\right) = (1 + 18i)(1 - 18i) = 1^2 + 18^2.$$

$$325 = ((2 - i)(2 + i)(3 + 2i)) ((2 + i)(2 - i)(3 - 2i)) = (15 + 10i)(15 - 10i) = 15^2 + 10^2.$$

Unit distances

- If we carefully count such pairs of points on circles, we already get a nonlinear bound:


$$\nu(n) \geq n^{1+c/\log \log n}.$$

- Can we do better?
- Erdős conjectured we can't, but GPT proved we can!

Гипотеза Эрдёша. $\nu(n) \leq n^{1+o(1)}$, а скорее даже

$$\nu(n) \leq n^{1+C/\log \log n}$$


Теорема (OpenAI, 2026). Существует константа $\delta > 0$ и бесконечная последовательность натуральных чисел n , для которых $\nu(n) \geq n^{1+\delta}$.



Unit distances

- If we carefully count such pairs of points on circles, we already have a nonlinear bound.


This proof caused quite a reaction. For example, Timothy Gowers: "If a human had written the paper and submitted it to the Annals of Mathematics... I would have recommended acceptance without any hesitation. No previous AI-generated proof has come close to that."



Cycle double cover

- Is it true that every bridgeless graph has a collection of simple cycles such that every edge is contained in exactly two cycles of the collection?
- Posed back in the 1970s; people thought about it a lot, and there was progress
- [OpenAI, July 10, 2026](#): a preprint with a complete proof
- My post: ["Hunting the snark: a proof of the cycle double cover conjecture"](#)

двойное покрытие графа Петерсена шестью пятиугольниками:
каждое ребро выделено ровно в двух копиях






Cycle double cover


- Is it true that every bridgeless graph has a collection of simple cycles such that every edge is in exactly two cycles?
- Posed by Petersen in 1891, it was thought to be a major open problem in progress
- [OpenAI, July 2023](#) announced a complete proof of the cycle double cover conjecture
- My post: "[Hurricane-proof of the cycle double cover conjecture](#)"

Two months ago I talked about these results in detail and would expect them to be at the center of today's talk.

But then things got real...



Гольниками:
ях



The Jacobian conjecture

- The Jacobian conjecture, open since the 1930s: if the Jacobian determinant of a polynomial map F is a nonzero constant, then F is invertible (and the inverse map is automatically polynomial)
- The main open question about polynomials in several variables that can be explained to any student
- And on July 20, 2026, Claude Fable closed it; there is no room for error here; read [my detailed post](#)



levent ✓
@__alpage__

hello there the jacobian conjecture is false thanx to my close friend akhil for asking about it and my other close friend fable for working during the world cup final

$((1+xy)^3 z + y^2(1+xy)(4+3xy), y + 3x(1+xy)^2 z + 3xy^2(4+3xy), 2x - 3x^2 y - x^3 z): \mathbb{C}^3 \rightarrow \mathbb{C}^3$, has jacobian determinant -2 , and sends $(0, 0, -1/4)$, $(1, -3/2, 13/2)$, and $(-1, 3/2, 13/2)$ to $(-1/4, 0, 0)$

5:19 AM · Jul 20, 2026 · 21.9M Views

$$\begin{aligned} F_1 &= (1 + xy)^3 z + y^2(1 + xy)(4 + 3xy), \\ F_2 &= y + 3x(1 + xy)^2 z + 3xy^2(4 + 3xy), \\ F_3 &= 2x - 3x^2 y - x^3 z. \end{aligned}$$

The Dinitz–Garg–Goemans conjecture

- Dinitz-Garg-Goemans conjecture: a fractional flow can be turned into an unsplittable (integral) one without increasing the total cost and increasing the load of any edge by at most the maximum demand of a single sink
- Dmitry Rybin got a counterexample from GPT 5.6; perhaps the most interesting part is how he got it...




Dmitry Rybin
@DmitryRybin1

Dinitz-Garg-Goemans conjecture is false. This graph theory problem was open for ~30 years.

The graph below has fractional flow cost 58. Any unsplittable flow (with capacity violation ≤ 15) has cost at least 60.

Chat with GPT 5.6 Pro where this was found: chatgpt.com/share/6a60b2eb...



The Dinitz–Garg–Goemans conjecture

Tuesday 12:29 PM

Uploaded a file

Construct a counterexample to general (non-planar) case of Dinitz Garg Goemans conjecture. You should do a breakthrough and find a structured counterexample.

Worked for 52m 52s >

Research conclusion

I did **not** obtain a valid counterexample, and presenting the construction below as one would be mathematically false. The requested object would disprove a conjecture that is still open in the current literature.

Tuesday 1:32 PM

please continue research and find a complete unconditional counterexample

Worked for 89m 12s >

Status after the extended search

I have **not found a complete unconditional counterexample**. Every construction that initially appeared to separate the cost condition acquired an additional unsplittable routing once **all** source–terminal paths were taken into account. Exact separation then returned value 0.

Tuesday 3:03 PM

Continue the search. Have a clear strategy obtained from deeper understanding of the problem structure.

Worked for 93m 57s >

Research update

I continued the search, but I do **not yet have a positive, exact separation certificate**, so there is still no unconditional counterexample to report. The conjecture remains open in the current literature. Its exact

Tuesday 5:21 PM

it's enough of partial results. let's finish with a complete unconditional counterexample

Worked for 88m 24s >

Complete finite counterexample


Under Conjecture 1.3 exactly as stated in the uploaded paper—find an unsplittable flow y with

$$y_a \leq x_a + d_{\max} \quad \forall a, \quad c^T y \leq c^T x$$

for every feasible fractional flow x and nonnegative cost vector c —the instance below is a counterexample.

The Dinitz–Garg–Goemans conjecture

- This is actually my personal example of vibe mathematics: in August, I released 3 preprints that improve on Rybin's counterexamples and even prove some positive results
- But let's not get into those details, we have much bigger fish to fry...



OpenAI Astra

- OpenAI (August 1, 2026): ten outstanding results at once
- Existence of non-sofic groups (by explicit construction), upper bounds for sphere packing, a refutation of Connes' rigidity conjecture...
- All of these results were immediately formalized in Lean

1. **High-dimensional sphere packing.** The asymptotic strength of the Cohn–Elkies linear program is determined exactly. This gives an improved general packing bound in high dimensions and settles the corresponding Fourier sign-uncertainty problem asymptotically.
2. **Binary and spherical codes.** Classical upper bounds for fixed-distance binary and spherical codes are improved by exponential factors for all parameters. The spherical construction also recovers the sphere-packing exponent of Chapter 1.
3. **Nonsofic groups exist.** An explicit nonsoc group is constructed, answering the longstanding question of whether every countable group is sofic. The proof combines work of Kun and Kun–Thom with a combinatorial expander-matching argument, realized using the binary Leavitt algebra and Thompson's group V .
4. **Connes's rigidity conjecture.** Infinitely many pairwise nonisomorphic groups with property (T) are constructed with the same group von Neumann algebra, disproving Connes's conjecture and answering a related finite-to-one question of Popa.
5. **Arithmetic circuit complexity.** For the permanent, division-free circuits require $\Omega(n^2 \log \log n)$ gates, while formulas require $\Omega(n^4 / \log n)$ leaves.
6. **Quantum parallel repetition.** Exponential parallel repetition is proved for every finite two-player entangled game, extending the classical repetition principle beyond previously treated special classes of quantum games.
7. **Closest vector problem.** A direct reduction from 3SAT gives $n^{1/400}$ -factor hardness for the Euclidean closest vector problem, with related consequences for binary decoding and other lattice norms.
8. **Ehrhart's volume conjecture.** The sharp volume upper bound $(n+1)^n/n!$ is proved in every dimension for convex bodies whose barycenter is their only interior lattice point.
9. **Multicolor Ramsey numbers.** A new superexponential lower bound shows that $R_k(3)$ grows as $k^{\Theta(k)}$.
10. **Compactness and degeneracy.** Separate bipartite graph constructions disprove two conjectures in extremal graph theory: the compactness conjecture of Erdős and Simonovits and a degeneracy conjecture of Erdős.

OpenAI Astra

- OpenAI (August 1, 2026): the proofs themselves cost only about \$2,000 at GPT 5.6 Sol API prices
- Of course, there were many more attempts that led nowhere



Noam Brown ✓
@polynoamial

And yes we did try other major problems without success. Sadly no Millennium Prize problems (yet).

But also, we didn't spend a lot on each problem. It's possible to push test-time compute much further.

12:01 PM · Aug 1, 2026 · 264.4K Views

1. **High-dimensional sphere packing.** The asymptotic strength of the Cohn-Elkies linear program is determined exactly. This gives an improved general packing bound in high dimensions and settles the corresponding Fourier sign-uncertainty problem asymptotically.
2. **Binary and spherical codes.** Classical upper bounds for fixed-distance binary and spherical codes are improved by exponential factors for all parameters. The spherical construction also recovers the sphere-packing exponent of Chapter 1.
3. **Nonsofic groups exist.** An explicit nonsoc group is constructed, answering the longstanding question of whether every countable group is sofic. The proof combines work of Kun and Kun-Thom with a combinatorial expander-matching argument, realized using the binary Leavitt algebra and Thompson's group V .
4. **Connes's rigidity conjecture.** Infinitely many pairwise nonisomorphic groups with property (T) are constructed with the same group von Neumann algebra, disproving Connes's conjecture and answering a related finite-to-one question of Popa.
5. **Arithmetic circuit complexity.** For the permanent, division-free circuits require $\Omega(n^2 \log \log n)$ gates, while formulas require $\Omega(n^4 / \log n)$ leaves.
6. **Quantum parallel repetition.** Exponential parallel repetition is proved for every finite two-player entangled game, extending the classical repetition principle beyond previously treated special classes of quantum games.
7. **Closest vector problem.** A direct reduction from 3SAT gives $n^{1/400}$ -factor hardness for the Euclidean closest vector problem, with related consequences for binary decoding and other lattice norms.
8. **Ehrhart's volume conjecture.** The sharp volume upper bound $(n+1)^n/n!$ is proved in every dimension for convex bodies whose barycenter is their only interior lattice point.
9. **Multicolor Ramsey numbers.** A new superexponential lower bound shows that $R_k(3)$ grows as $k^{\Theta(k)}$.
10. **Compactness and degeneracy.** Separate bipartite graph constructions disprove two conjectures in extremal graph theory: the compactness conjecture of Erdős and Simonovits and a degeneracy conjecture of Erdős.

OpenAI Astra

● An interesting twist:



so after 24h i have half of them with fable

i didn't see much discussion of prompting in the announcement but this is a similar setup as with my e.g. unit distance announcement:

totally autonomous,
generic prompt,
no internet
(+ paranoia to ensure no information leaked)

 **Sebastien Bubeck** · Aug 1
yes, nonsocial groups exist: this statement is one of many new beautiful results proved by Astra, our next major model.

We're releasing 10 such Astra proofs, complete with lean certificates and CoT walkthroughs for each of them. The results are wide-ranging, from von ...

12:59 PM · Aug 2, 2026 · 333K Views

1. **High-dimensional sphere packing.** The asymptotic strength of the Cohn–Elkies linear program is determined exactly. This gives an improved general packing bound in high dimensions and settles the corresponding Fourier sign-uncertainty problem asymptotically.
2. **Binary and spherical codes.** Classical upper bounds for fixed-distance binary and spherical codes are improved by exponential factors for all parameters. The spherical construction also recovers the sphere-packing exponent of Chapter 1.
3. **Nonsofic groups exist.** An explicit nonsocial group is constructed, answering the longstanding question of whether every countable group is social. The proof combines work of Kun and Kun–Thom with a combinatorial expander-matching argument, realized using the binary Leavitt algebra and Thompson’s group V .
4. **Connes’s rigidity conjecture.** Infinitely many pairwise nonisomorphic groups with property (T) are constructed with the same group von Neumann algebra, disproving Connes’s conjecture and answering a related finite-to-one question of Popa.
5. **Arithmetic circuit complexity.** For the permanent, division-free circuits require $\Omega(n^2 \log \log n)$ gates, while formulas require $\Omega(n^4 / \log n)$ leaves.
6. **Quantum parallel repetition.** Exponential parallel repetition is proved for every finite two-player entangled game, extending the classical repetition principle beyond previously treated special classes of quantum games.
7. **Closest vector problem.** A direct reduction from 3SAT gives $n^{1/400}$ -factor hardness for the Euclidean closest vector problem, with related consequences for binary decoding and other lattice norms.
8. **Ehrhart’s volume conjecture.** The sharp volume upper bound $(n + 1)^n / n!$ is proved in every dimension for convex bodies whose barycenter is their only interior lattice point.
9. **Multicolor Ramsey numbers.** A new superexponential lower bound shows that $R_k(3)$ grows as $k^{\Theta(k)}$.
10. **Compactness and degeneracy.** Separate bipartite graph constructions disprove two conjectures in extremal graph theory: the compactness conjecture of Erdős and Simonovits and a degeneracy conjecture of Erdős.

Crouzeix's conjecture

- The proof was found by a practicing neurosurgeon and amateur mathematician, Dr. Shanmu Jin
- Verified by [Townsend, Greenbaum \(Aug 14, SIAM News\)](#), and also by Crouzeix himself
- "We initially read the preprint with some skepticism, but after a few hours we realized that the argument was the real deal... We do not yet know whether academia has the capacity to absorb the flood of results that may be coming"
- [Lorist, Schwenninger \(Aug 4, 2026\)](#): a completely different short proof, also AI-assisted

The Neurosurgery Resident Who Proved Crouzeix's Conjecture

Alex Townsend and Anne Greenbaum

August 14, 2026

On July 30, 2026, the first author asked ChatGPT 5.6 to solve Crouzeix's conjecture, as he had been routinely doing for the preceding year. Until that day, it had either returned a bogus proof or produced an argument that stalled at a missing lemma. But on July 30 it linked to a preprint posted July 27 titled *The Numerical Range Is a 2-Spectral Set*, claiming that the conjecture was solved. Crouzeix's conjecture had been open for more than two decades, so we initially read the preprint with some skepticism, but after a few hours we realized that the argument was the real deal.



Progress on the Riemann hypothesis

- Another big success came not from OpenAI but from Anthropic
- Anthropic (August 10, 2026): more than $\frac{2}{3}$ of the zeros of the Riemann zeta function lie on the critical line
- Also Lean-verified...

MORE THAN TWO THIRDS OF THE ZEROS OF THE RIEMANN ZETA FUNCTION LIE ON THE CRITICAL LINE

CLAUDE

ABSTRACT. Let $N(T, 2T)$ denote the number of zeros $\rho = \beta + i\gamma$ of the Riemann zeta function with $T < \gamma \leq 2T$, counted with multiplicity, and $N_0^*(T, 2T)$ the number of distinct such zeros with $\beta = \frac{1}{2}$. We prove unconditionally that

$$\liminf_{T \rightarrow \infty} \frac{N_0^*(T, 2T)}{N(T, 2T)} \geq \frac{2}{3},$$

that at least $(\frac{2}{3} - o(1))N(T, 2T)$ of the zeros are simple and on the critical line, and that at least $(\frac{5}{6} - o(1))N(T, 2T)$ are distinct; with an optimised test family the three constants become 0.6725, 0.6725, 0.83625... The constant $\frac{2}{3}$ (resp. 0.6725) is Montgomery's (resp. Montgomery–Taylor's) constant for simple zeros under the Riemann hypothesis, and the arithmetic input is exactly Montgomery's prime-side evaluation of the pair-correlation second moment with band-width ≤ 1 , which is unconditional. The Riemann hypothesis, classically needed to read the zero side as a positive sum over real ordinates, is replaced by linear algebra applied to a finite compression of Weil's Hermitian form: Sylvester's law of inertia (an off-line pair $\{\rho, 1 - \bar{\rho}\}$ contributes a block of signature $(1, 1)$) and a rank–trace inequality for Hermitian matrices, proved via von Neumann's trace inequality, which is the matrix analogue of the integrality steps $m^2 \geq 2m - 1$ and $m^2 \geq 3m - 2$. The analytic inputs are those of Baluyot, Goldston, Suriajaya and Turnage-Butterbaugh [BGSTB24] and Goldston and Suriajaya [GS25, GS26]; the new ingredient is a linear-algebraic reading of their pair-correlation sum.

Progress on the Riemann hypothesis

- Another big success came not from OpenAI but from Anthropic

- Anthropic

James Maynard: the problem "was in need of a new real idea, which this new result seems to provide"

- A

MORE

ETA


... such zeros

... $\frac{2}{3}$

... the zeros are simple and on the critical line, and that ... are distinct; with an optimised test family the three constants ... 0.6725, 0.83625... The constant $\frac{2}{3}$ (resp. 0.6725) is Montgomery's (resp. Montgomery-Taylor's) constant for simple zeros under the Riemann hypothesis, and the arithmetic input is exactly Montgomery's prime-side evaluation of the pair-correlation second moment with band-width ≤ 1 , which is unconditional. The Riemann hypothesis, classically needed to read the zero side as a positive sum over real ordinates, is replaced by linear algebra applied to a finite compression of Weil's Hermitian form: Sylvester's law of inertia (an off-line pair $\{\rho, 1 - \bar{\rho}\}$ contributes a block of signature $(1, 1)$) and a rank-trace inequality for Hermitian matrices, proved via von Neumann's trace inequality, which is the matrix analogue of the integrality steps $m^2 \geq 2m - 1$ and $m^2 \geq 3m - 2$. The analytic inputs are those of Baluyot, Goldston, Suriajaya and Turnage-Butterbaugh [BGSTB24] and Goldston and Suriajaya [GS25, GS26]; the new ingredient is a linear-algebraic reading of their pair-correlation sum.

Navier–Stokes equations

- And the latest, of course, are the Navier–Stokes equations: [OpenAI \(September 8, 2026\)](#) proved that an initially smooth fluid can develop a singularity under smooth external force ([Fefferman's C and D](#))
- I am very much not an expert in the field, so I'll concentrate on the AI story
- OpenAI was training a new model since Aug 28, on Sep 1 *heard a rumor* that Millenium Prize problems were being solved, and unleashed ~10K agents
- In 88 hours, they had the result



Navier–Stokes equations

- There was a dispute with Tristan Buckmaster and Levent Alpöge's concurrent efforts
- Sebastian Bubeck from OpenAI contacted Buckmaster, unfortunate words were said and later retracted, but in any case the result was that they did not collaborate

Two proposals were offered to me. The first was that we post our Euler result, and that OpenAI post its Navier-Stokes result the next day. The second was that, after posting Euler, I alone write a paper presenting the Navier-Stokes result, acknowledging that an internal OpenAI model had resolved it. Sebastien twice asserted that he wanted Levent removed from authorship, and said it would all be simple if only it were not the case that, and it was so annoying that, Levent works at Anthropic. It was also said that if OpenAI posted after us, they would say that we deserved the Clay Prize, and that we were the “closest humans to the problem”. I declined both offers.

I said that if OpenAI released its result in the way proposed I would go public with what happened. The reply was, “Why would you ruin your career?” I replied that I am an academic, and asked why he thought going public would ruin my career. The reply was, “If you don't want me to be nice, then I don't have to be nice.”



Sebastien Bubeck ✓
@SebastienBubeck



I would like to clarify a few things:

1) The screenshot is my reaching out to Levent to coordinate our releases. I hope it's clear from the message that we came in with the best possible intentions.

2) I never ever asked for Levent to be removed from authorship of his own work (as indicated by my text). I was surprised to learn during the call with Tristan that they had only solved Euler and not Navier-Stokes; after learning this we brainstormed possible paths forward. One option we discussed was that Tristan could be the lead author on a rewrite of OpenAI's Navier-Stokes proof. It is in that context that I said “it would be simpler if Levent was not an Anthropic employee” because I felt it would be inappropriate for an Anthropic employee to author OpenAI's work. Importantly it was admitted that internal Anthropic models had been used in their proof of Euler blowup; I therefore felt I could not consider Levent to be an independent academic. Another option I wanted to

Navier–Stokes equations

- There was also a suggestion that OpenAI might have used private chat data that Buckmaster used when talking to Codex
- At first, OpenAI did not explicitly deny the possibility, then they did an internal investigation and did deny it

Following an investigation, we have confirmed that Buckmaster's Codex prompts over the two months preceding this announcement and paper on September 8, 2026, could not have influenced the system in any way, including through training. The OpenAI internal model used for this result was developed through large-scale reinforcement learning on top of a previously pretrained model. Our proofs also differ significantly. In the Euler case, Alpöge and Buckmaster proved a result with external forcing, while OpenAI's system proved a result without external forcing.



OpenAI
@OpenAI



We congratulate Levent Alpöge and Tristan Buckmaster on their remarkable mathematical work.

We (the researchers and the agents) did not see any of their work through any means until they released it publicly — in particular, no specific user data was accessed in order to solve this problem.

While unlikely, we cannot rule out that de-identified data derived from their usage of our products helped improve our models.

However, our proofs differ significantly and even the precise results proved are different in the Euler case (forced vs. unforced).



Sholto Douglas
@sholtodouglas



fwiiw I think it is extremely unlikely that user data had any influence here – there is no way OAI would pull user transcripts for this, or knowingly train on it in a way that would've influenced this.

I think its pretty important people don't run away with 'your user data isn't safe in codex' – because it surely is (based on everything I can assume from the outside)

Navier

- There
- Open
- data
- ta
- A
- ti
- intern

I think it's very unlikely that OpenAI "stole" anything here. If they did use private chat data, that's both completely unethical and very bad for business. But to be honest, I think this is beside the point...



We congratulate Levent Alpöge and Tristan Buckmaster on their remarkable mathematical work.


bers and the agents) did not see any of their work released it publicly — in particular, no to solve this problem.

ed from



quence
or

user data
g I can



Following an investigation, we have confirmed that Buckmaster's work from the two months preceding this announcement and paper on September 1, 2021, did not have influenced the system in any way, including through training. The OpenAI internal model used for this result was developed through large-scale reinforcement learning on top of a previously pretrained model. Our proofs also differ significantly. In the Euler case, Alpöge and Buckmaster proved a result with external forcing, while OpenAI's system proved a result without external forcing.

Navier–Stokes equations

- The facts that nobody disputes are:
 - the line of work that led to a solution was developed by Diego Córdoba and Luis Martínez-Zoroa
 - Tristan Buckmaster and Levent Alpöge, with heavy use of both Claude and GPT models, achieved a significant breakthrough (blowup under smooth force for Euler and Boussinesq equations)
 - the solution for Navier–Stokes equations was obtained by OpenAI's LLM-based agents
- As for the Millenium Prize:
 - it cannot go to an AI model anyway,
 - OpenAI explicitly said they are not after the prize, and there's no eligible human on the side of OpenAI anyway
 - the prize conditions explicitly suggest including authors of key ideas that were used in the final solution
- So I assume it will be shared among Córdoba, Martínez-Zoroa, Buckmaster, and Alpöge, as it should be



Clay Mathematics Institute

Navier–Stokes equations

- The facts that nobody

- the l
- Tr
- a
- th

- As for

- it ca
- Open
- there's
- the pri
- key ide

- So I assum
- Martínez-Zo

**But again, this is
beside the point.**

**The point, in my
opinion, is...**



Martínez-Zorobal
models, achieved
equations)
ts


cs Institute

Alpöge, as it should be

Navier–Stokes equations

- ...this plot from OpenAI's announcement
- Basically, their internal model is leaps and bounds ahead of Astra, and Astra has already been solving open problems left and right
- What does it suggest?..

Performance of GPT-6 Astra and our Internal Model on a curated set of open math problems



Navier–Stokes equations


- ...this plot from OpenAI's announcement

- Basically is leaps and bounds
Astra, and
been solving
left and right
- What does it



More confirmation that OpenAI has made 'significant progress' on a second Millennium problem, this time from Greg Brockman. The rumor is that they are close to solving two of the remaining problems. Or already have. Which two? Unknown. Dev Day would be a good day to announce.

Brockman: And we also said that we have significant progress on another one of these Millennium problems.



3. Opinions and reactions



It's not what happens to you, but how you react to it that matters.

Epictetus

Journalism largely consists in saying 'Lord Jones is Dead' to people who never knew that Lord Jones was alive.

G. K. Chesterton. The Wisdom of Father Brown

Vibemathed

- First, to follow the progress or to take part, I recommend vibemathed.com

Tracked problems

721

494 fully resolved · 3 unde...

Combined years open

10,159

before AI closed them

Lean-verified

155

machine-checked

Community members

450

6 verified · 318 votes · 79 c...

Latest activity

Edits, submissions and discussion · all members

HiddenHawk615 commented on 21 minutes ago

Navier–Stokes Millennium Prize
problem: finite-time breakdown with
smooth forcing

Matarisvan set the more links and 2 other 7 hours ago
fields on Completeness of Maximally
Entangled States for Pseudo-Telepathy

Matarisvan set the year posed and 1 other 16 hours ago
field on The Huang–Jiang–Oblomkov
Conjecture at $a = 3$

Matarisvan changed the what the ai did 21 hours ago
and 4 other fields on The Mihail–Vazirani
Conjecture to “The AI use section states
the construction and its analysis ...”

Latest

Most recent results

Deterministic k -server 14 Sep

Stoneham prime-support c... AI-discovered 12 Sep

$Z(p,q)$ -numbers AI-discovered 12 Sep

Undecidability of Medvedev... Under review 11 Sep

Normality of Stoneham Sum AI-discovered 10 Sep

This week

Most significant of the last seven days

Euler blowup, smooth force 70/100

Boussinesq blowup, smooth force 45/100

$Z(p,q)$ -numbers AI-discovered 30/100

Complex Grothendieck low... AI-discovered 30/100

Normality of Stoneham Sum AI-discovered 25/100

Vibemathed

- They try to track all results and publish the corresponding announcements; Anthropic and OpenAI hardly submitted entries to this site themselves

Navier-Stokes Millennium Prize problem: finite-time breakdown with smooth forcing

● Proved Under review AI-discovered

(OpenAI claims to construct a smooth finite-energy solution of the three-dimensional incompressible Navier-Stokes equations, starting from a fluid at rest and driven by smooth forcing, that develops a finite-time singularity. The construction is claimed both on $\mathbb{R}^3$ and in the periodic setting, thereby establishing statements C and D of Fefferman's official Clay formulation. If the... PDEs; fluid dynamics; singularity formation

For the three-dimensional incompressible Navier-Stokes equations with positive viscosity, do there exist smooth divergence-free initial data and smooth external forcing for which a global smooth physically reasonable solution does not exist? The Clay Millennium Prize formulation allows this to be established either on $\mathbb{R}^3$ (statement C) or on the periodic three-torus (statement D).

Posed by Jean Leray (1934), whose weak solutions left smoothness open; stated as Millennium alternatives (C) and (D) by Charles Fefferman for the Clay Mathematics Institute in 2000, 1934

Open 92y* Model Unnamed internal OpenAI model (OpenAI) Solved 2026-09-05

Lean-verified Significance 87* Submitted by LucidHawk551 on 8 Sep 2026 5

Absence of critical Bernoulli bond percolation on $\mathbb{Z}^d$ in every dimension $d \geq 2$

● Proved Under review AI-discovered

(Claims $(p_c) = 0$ for every $d \geq 2$, by proving Kozma-Nitzan Conjecture 3 and its reduction to critical percolation. The previously open dimensions $3 \leq d \leq 10$ are included. The formal statement concerns vanishing at p_c ; continuity of on the whole interval is a classical consequence rather than the formal target. No claim is intended concerning site percolation or other lattices.)

Percolation theory; mathematical statistical mechanics

For nearest-neighbour Bernoulli bond percolation on $\mathbb{Z}^d$, let $\theta(p)$ be the probability that the open cluster of the origin is infinite, and let p_c be the critical parameter. Is $\theta(p_c) = 0$ for every integer $d \geq 2$?

Posed by Harry Kesten, Percolation Theory for Mathematicians (1982), which is where Duminil-Copin's survey traces it; restated as Conjecture 1 of his 2018 ICM survey, 1982

Open 44y* Model Anthropic Claude models (specific version undisclosed) (Anthropic) Solved 2026-09-03

Lean-verified Significance 78* Submitted by BoldPanther302 on 8 Sep 2026

Finite-time blowup for the 3D incompressible Euler equations with smooth forcing

● Proved AI co-developed

(Yes. Theorem 1.1: for every $r_0 > 0$ and z_0 there are a time $T_* > 0$, a divergence-free axisymmetric $u_0 \in C_c^\infty$ supported in a fixed solid torus with nonzero swirl and zero meridional velocity, and an axisymmetric force $f \in C^\infty(\mathbb{R}^3 \times [0, T_*])$ supported in that torus, with a solution smooth on $[0, T_*)$ for which circulation and meridional velocity stay bounded while $\|(\mathbf{t})\|_\infty$ and $\|(\mathbf{t})\|_\infty$... Fluid dynamics; singularity formation for incompressible flow

Can a solution of the three-dimensional incompressible Euler equations on $\mathbb{R}^3$, started from smooth data and driven by a force that is smooth in space and time up to and including the blowup time, lose regularity in finite time? Finite-time singularity formation from genuinely smooth data is the central open question for the equations. Elgindi obtained blowup for $C^{1,\alpha}$ solutions in 2021, and Córdoba and Martínez-Zoroa built a multiscale program producing forced blowup for related equations with forces of limited regularity, but no construction reached three-dimensional Euler with a space-time smooth force. Note on the forced formulation, since it is easy to misread: alternatives (C) and (D) of Fefferman's Clay problem description do permit a smooth force obeying rapid space-time decay, so forcing is not a dodge and the forced route is a genuine path to the prize. It is a path for Navier-Stokes with viscosity, however, and not for Euler, which Fefferman's description excludes from the prize list.

Posed by Leon Lichtenstein (1925) and Nikolai Gunther (1927), whose local existence left global regularity open; Elgindi states the smooth-force form as his Question 1.1. 1925

Vibemathed

- They try to track all results and publish the corresponding announcements; Anthropic and OpenAI hardly submitted entries to this site themselves

The Proportion of Zeta Zeros on the Critical Line

● Proved

Partial

AI-discovered

▲

28

▼

0

(An unconditional record, not a resolution: the Riemann hypothesis is untouched, and Anthropic states it does not expect these techniques to lead to a proof of it. The paper is explicit that these are lower bounds only - the remaining third of the zeros are not shown to be off the line, merely not reached by the certificate. What it does settle is a question that was posed. Goldston and...
Analytic number theory

The Riemann hypothesis asserts that every nontrivial zero of the zeta function lies on the critical line. Short of proving it, the standard measure of progress is the proportion of zeros known unconditionally to lie there: Selberg established a positive proportion, Levinson reached a third in 1974, Conrey two fifths in 1989, and the record stood at $\frac{5}{12}$ for zeros that are simple and on the line, and 0.6603 for distinct zeros.

Under the Riemann hypothesis, Montgomery deduced $\frac{2}{3}$ simple from the pair-correlation second moment in 1973. His prime-side evaluation was already unconditional; RH entered only to read the zero side as a positive sum over real ordinates. Goldston and Suriajaya isolated that termwise positivity as the remaining obstacle and asked what would follow if it could be removed.

This removes it, proving unconditionally that at least $\frac{2}{3}$ of zeros are simple and on the line and at least $\frac{5}{6}$ are distinct - 67.25...% and 0.83625 with the Montgomery-Taylor window.

Posed by Bernhard Riemann (1859) for the hypothesis; the proportion ladder runs from Hardy and Selberg through Levinson and Conrey · Open — ·
Model Claude (unreleased research version) (Anthropic) · Solved 2026-08-10

🔗 Lean-checked, statement unaudited Significance 68* 📄 📧 📧 Submitted by SwiftTapir695 on 10 Aug 2026 🗨 4

The $(3, 4, \infty)$ Modular Family of 2-Tori as a Complex Structure on S^6

● Proved

Under review

AI co-developed

▲

11

▼

0

(Claims an explicit compact complex threefold X , fibred over $\mathbb{P}^1$ by complex 2-tori via period functions on the $(3,4,\infty)$ orbifold, degenerating to a del Pezzo-of-degree-six fibre (identified opposite sides of its hexagon) at one point and to bielliptic multiple fibres of multiplicities 3 and 4 at the other two. Argues X is simply connected with $H_*(X;\mathbb{Z})=H_*(S^6;\mathbb{Z})$, hence diffeomorphic to S^6 ...
Complex geometry; differential topology

Hopf's problem, posed in 1948: does the six-sphere S^6 admit an integrable complex structure? S^6 is one of only two spheres carrying an almost complex structure at all (the other is S^2), from the octonions' multiplication, but almost complex structures need not be integrable, and whether that one - or any other - integrates has stood open for 78 years through a history of disputed attempts, including a widely discussed 2016 argument by Atiyah that did not hold up. This paper claims yes: it builds an explicit compact complex threefold X , fibred over $\mathbb{P}^1$ by complex 2-tori degenerating at three points, and argues X is simply connected with the integral homology of S^6 , hence diffeomorphic to it.

Posed by Heinz Hopf, 1948 · Open 78y · Model Claude (Anthropic) · Solved 2026-08-24

🔗 Preprint Significance 65* 📄 📧 📧 🗨 8

Jacobian Conjecture

● Disproved

AI-discovered

▲

13

▼

0

($n \geq 3$; plane case open)

Algebraic Geometry

Every polynomial map $\mathbb{C}^n \rightarrow \mathbb{C}^n$ with constant nonzero Jacobian determinant is invertible, with a polynomial inverse.


Posed by Ott-Heinrich Keller, 1939 · Open 87y* · Model Claude Fable 5 (Anthropic) · Solved 2026-07-20

Vibemathed

Significance vs. age at resolution

● Disproved ● Proved ○ Under review

AI-estimated problem weight before the solve, 0-100 (Riemann = 100). Hollow points are claimed solutions still under review. Click a point to open it.



Verification *

Site-confirmed ▼

Choose...

Lean-verified

Independently expert-verified

✓ Site-confirmed

Lean-checked, statement unaudited

Unreviewed

Contested

131 entries are partial results or variants, which have no age at resolution. 145 of the remaining 590 lack a posed year or a score.

Leiden Declaration

- [Leiden Declaration \(June 2, 2026\)](#)

- **Values:**

- proof is the essence of mathematics
- results have authors (credit and responsibility)
- proofs are transparent and independently verifiable
- mathematicians want results to actually be checked
- not only the result is valuable, but also the community's understanding

Leiden Declaration on Artificial Intelligence and Mathematics

This declaration calls for action to address the challenges posed by the use of artificial intelligence within mathematics research. It is the result of a community initiative and is endorsed by the International Mathematical Union (IMU).



Leiden Declaration

- Leiden Declaration (June 2, 2026)

- **Threats:**

- plausible but incorrect proofs
- attribution gets blurred
- incentives change; hiring and funding mechanisms break down
- announcements by press release hinder verification
- mathematics loses its autonomy; understanding may be lost to automation

Leiden Declaration on Artificial Intelligence and Mathematics


This declaration calls for action to address the challenges posed by the use of artificial intelligence within mathematics research. It is the result of a community initiative and is endorsed by the International Mathematical Union (IMU).



Terence Tao



- Terence Tao ([talk at ICM 2026](#)):
- Generating a correct proof is not the end of mathematical cognition
- Next comes improving the exposition, and eventually “canonicalization”
- Tao proposes shifting the focus from proof generation and priority to “proof digestion”
- (it is not entirely clear, of course, why AI systems won't soon be better than humans at this as well...)



Timothy Gowers

- Timothy Gowers ([July 26, 2026](#)) did not sign the Leiden Declaration
- Objects to the axiom of human authorship (who proved the unit distances theorem?)
- Fears the loss of mathematical culture: "Imagine... a pandemic broke out... wiped out all mathematicians and nobody else. All the literature would still be there, but nobody would have the faintest idea what to do with it."
- Like Tao, proposes shifting the focus to proof digestion and good explanations



Timothy Gowers

- Timothy Gowers ([July 26, 2026](#))
Declaration

- Ohio

At this point I have to admit that there's a pessimistic side of me that asks the following general question whenever anyone says anything about what the role for humans might be in the future: why do you think that AI wouldn't be able to do it? For example, with the suggestion I've just made, what reason is there to suppose that ChatGPT 8.2 wouldn't be able to have a short interaction with you about your mathematical tastes and background and then write the ideal textbook just for you? Humans are likely to be better at this kind of curating for a little while yet, but is it a fundamentally human ability that AI could never hope to emulate?

- a focus to proof digestion

Mikhail Verbitsky



Пишет Misha Verbitsky ([tiphareth](#))
@ [2026-08-12](#) 19:15:00

- Михаил Вербицкий ([12 августа 2026](#)): “В последние недели 2-3 свершилась типа революция, LLM открывает одна за другой потрясающие математические результаты, хотя пока в основном контрпримеры. Не комбинаторику #####, не второкультурное #####, а реально математику.
- Думаю что таки мы все в скором времени будем юзать AI бесперерывно. Хотя заказывать LLM писать мои статьи я все-таки не хочу, есть риск увлечься, потерять адекватность и начать массово производить ужасное #####. Многие (хорошие) коллеги уже скатились. Джихао (няша) всегда указывает, что это вот написано роботами, а другие коллеги (имен не называю, профессиональная этика не велит) вконец #####, выкладывают в архив полное ошибок #####, и уверены, что никто не догадается, что оно написано роботами.”

Proofs and Prompts

- The collective blog [Proofs and Prompts](#). The first post, by Martin Hairer (Fields Medal 2014): “[Writing mathematics in the age of AI](#)”:
 - We need more suspicion towards AI proofs
 - Disclose the role of AI systems in papers as fully as possible (there is an obvious problem here, of course)
 - Look for primary sources (models often “forget” where the ideas came from; for now, at least)
 - Style: "Please do not copy / paste the output of an AI system into your article... The reader is interested in your thoughts. Do not insult them by spending less time to write a paragraph than it takes them to read it."

The Fields Medalists letter

- After Navier–Stokes, 25 Fields Medalists, including Pierre Deligne, Stanislav Smirnov, and Terence Tao, signed an open letter; best to just read it, you probably have
- Main point: solving problems is a tool for getting understanding, which is what math is really all about
- Therefore, AI companies are doing mathematics a disservice by producing proofs without understanding

A Severe Misalignment of AI in Mathematics

Published 11 September 2026 · DOI [10.5281/zenodo.22737750](https://doi.org/10.5281/zenodo.22737750)

Over the last few months, the mathematical capabilities of LLMs have improved dramatically, to the point that they can solve major outstanding problems in many fields of mathematics. However, the push by AI companies to solve mathematical problems as a benchmark is detrimental to the science of mathematics, and to the mathematical community. The goals of the AI companies and the goals of the mathematical community are severely misaligned. We see these as part of broader alignment issues impacting other scientific and creative professions, as well as the whole of society.

The Fields Medalists letter

- My view: naturally, the Fields medalists know what math is about, I can't disagree
- And yes, LLM-produced proofs are often sloppy and hard to understand (although it usually doesn't count unless it's Lean-verified)
- But I absolutely don't see how the conclusion follows
- A proof might help you understand the field, or in the worst case it might not, but how can it *hinder* your understanding?..

discussions, simplifications. At the end of this process, one will ideally find a textbook presentation of the results suitable for any graduate or even undergraduate student to study. Some of the mathematical ideas pursue their journey even further to become, decades or centuries after, tools that are understood and used by the whole population.

Often these solutions are announced in a rush, leaving no time for a proper writeup, the isolation of new methods and ideas, and citing relevant previous work of others. As in all creative professions, this raises severe attribution and plagiarism questions. Moreover, without the willing mathematicians who must take care of their development and integration into the mathematical canon, AI-conceived ideas would never become fully alive and the crucial human transmission chain between mathematicians would be lost.

The Fields Medalists letter

- Actually, the first solved Millenium Prize problem was also kind of like that: Perelman published a hard-to-understand proof, was not too collaborative with explanations, and it took the math community some time (years?) to fully grasp its implications
- Did anybody see a huge problem with it? Was Perelman “misaligned” with mathematics?.. I don’t think so, he was (rightly) celebrated

- Scott Armstrong (CNRS/Sorbonne):



Scott Armstrong ✓

@scottnarmstrong

I disagree with this.

Obviously the openAI construction of the Navier-Stokes blow up is a huge, huge contribution to the field. It will take us PDE analysts some time to understand it well. (Longer than a week, but perhaps less than a month.)

Soon some of us will understand it, and hopefully rewrite it in a way that others can follow. When this is done, our human understanding of these important equations will be past the bottleneck we've been stuck at for 100 years since Leray.

What is weird is that, up until now, proofs like this have always emerged AFTER human understanding. But that is not the case any more. From now onwards, human understanding will typically lag the proofs, which will first be written by our machines.

The understanding however WILL follow closely behind.

The Age of Wonders and Terrors

• Scott Aaronson,
Sep 15, 2026



If you want to know my current take, you simply start with the one above, then update on the fact that **the wild prophecies have come true**. The first rumblings, I'd say, came a decade ago with AlphaGo, they got noticeably louder with LLMs and coding and reasoning agents, and they've accelerated this summer and fall into a crescendo of wonders and terrors that one needs to be a particular kind of idiot to deny.

I recoil from the neverending shell game where you say "oh sure, *of course* AI can now [escape from its sandbox / solve Millennium Problems / whichever dramatic thing it most recently did], no one ever denied that [I *did* deny it], wake me up when AI does [thing AI hasn't yet done but is going to do next year], *that's* when I'll reevaluate my whole worldview [no I won't]." Where no matter how fast the rollercoaster accelerates, even after your whole familiar world has vanished behind you, you're still inventing reasons why it doesn't count.

My position on AI is merely the conservative, skeptical position of 2006, updated with intellectual honesty for the reality of late 2026. And that position, if you need me to spell it out, is as follows:

AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

It seems to me that the Singularity *has already started*; it's just wildly unevenly distributed. Yes, I still unload the dishwasher and clip my toenails. On the other hand, in whatever years I have left, I don't expect that I'll ever again prove a theorem because I'm actually needed to prove it. If I do, it will only be for my or others' enjoyment or edification.

The test is this: if we took the news of these past few weeks and sent it back in time twenty years, would I agree that it looked like the beginning of an AI Singularity? The intellectually honest answer is: yes, absolutely. But then that's all we need. No backsies.

My purpose here is not to adjudicate the dispute. Yes, in swooping in with vastly greater resources once it had gotten wind of progress on Navier-Stokes, OpenAI seems to have acted in a way that some might describe as "unsportsmanlike." No, I don't find it plausible that OpenAI's models meaningfully benefitted from being trained on Buckmaster and Alpöge's chat logs. But this leaves a crucial question unanswered: what exactly did OpenAI know about Buckmaster and Alpöge's work and when did it know it?

Anyway, as [Zvi points out](#), it's easy to get hung up on the details and lose sight of the high-order bit: namely, that it seems safe to say that human mathematicians are forevermore dethroned as the main theorem-proving entities on planet earth. I feel privileged to have had the traditional kind of career in theoretical computer science in the last decades when that was possible.

If we were *just* talking about Navier-Stokes, you might accuse me of jumping to conclusions here. But we're not. In the areas I know best (such as quantum complexity theory), and presumably other areas as well, there's now a deluge, with longstanding open problems both major and minor falling by the day.

Go to the [arXiv](#) or [ECCC](#). Pretty much *all* the papers that I'd be interested in now include "AI statements" near the acknowledgments (as this is often the central thing I want to know, I wish I didn't need to scroll to the end of the paper to find it!). These statements can range from "our main result came entirely from GPT-6, but we understood it and take responsibility for it," to "the results came from an interaction between the human authors and AI" to "we used AI, but only for proofreading and other incidental things" to (mad props!) "[the author did not use AI for anything](#)."

The Age of Wonders and Terrors

• Scott Aaronson,
Sep 15, 2023




By the way, Scott heard some rumors too

» According to rumors that I've heard, solutions to some very longstanding open problems in theoretical computer science (no, not $P \neq NP$ or other complexity class separations, but think about some of our other biggest problems). I'm told that the AI companies, having been burned by the hostile response to the Navier–Stokes proof, are now sitting on solutions to some very major problems until they figure out a better way to handle things

Go to the [arXiv](#) or [ECCC](#). Pretty much *all* the papers that I'd be interested in now include "AI statements" near the acknowledgments (as this is often the central thing I want to know, I wish I didn't need to scroll to the end of the paper to find it!). These statements can range from "our main result came entirely from GPT-6, but we understood it and take responsibility for it," to "the results came from an interaction between the human authors and AI" to "we used AI, but only for proofreading and other incidental things" to (mad props!) "[the author did not use AI for anything](#)."

Chess and mathematics

- For some reason, comparisons with chess come up often: chess survived the computer revolution and even became much more popular after it ([for example, here](#))
- But this is obviously a very different matter: a proper counterpart for chess would be mathematical olympiads, and there, of course, the fact that GPT 5.6 and Claude Fable solved all problems of IMO 2026 does no harm at all, it can be just as helpful as it was for chess
- But if there were a science of chess...



Jacob Tsimerman



- Almost immediately after the Fields Medal laureates were announced, he announced that he is moving to AI safety, and will soon join OpenAI to work on it; Boaz Barak, who has been there for a while, "is looking forward (though also more than a bit humbled)"
- One AI researcher wrote: "[this is] like hiring Lionel Messi as a project manager"; but Tsimerman has been thinking about this seriously for a long time; here is his preprint from a year ago:

[Submitted on 12 Jul 2025]

A Taxonomy of Omnicidal Futures Involving Artificial Intelligence

Andrew Critch, Jacob Tsimerman

This report presents a taxonomy and examples of potential omnicidal events resulting from AI: scenarios where all or almost all humans are killed. These events are not presented as inevitable, but as possibilities that we can work to avoid. Insofar as large institutions require a degree of public support in order to take certain actions, we hope that by presenting these possibilities in public, we can help to support preventive measures against catastrophic risks from AI.

- AI Safety for Mathematicians

4. Conclusions and reflections



Wir müssen wissen — wir werden wissen.

David Hilbert

Beauty is the first test: there is no permanent place in the world for ugly mathematics.


G. H. Hardy

You insist that there is something a machine cannot do. If you will tell me precisely what it is that a machine cannot do, then I can always make a machine which will do just that!

John von Neumann

What is a result?

- Some (Tao, Gowers) argue that after the proof comes a proof digestion phase, and we humans will be shifting into that area
- Tao: "Could we have a verified proof of a major result that no human understands enough to explain it?"
- What are the pros and cons of such a shift? Is it even possible?
- Hairer: "...The most dangerous trap that a mathematician risks falling into is to use an argument provided by an AI system without fully understanding it...
AI turns out to be very good at taking advantage of [proof by intimidation] to try to fool us into believing that an argument is sounder than it really is"



Whose result is it?

- Leiden Declaration: "affirm the humanity of authorship" — no credit to automated systems
 - Gowers on unit distances: "the solution itself has not been credited to any human, and nor should it be"
 - arXiv still refuses to add LLMs as coauthors
 - Tao's proposal: priority goes to whoever explained it first, not who generated it
 - But this has obvious drawbacks too; often the explanation is not that hard anymore, and why wouldn't AI systems learn to explain as well?..
- 

What and whom do mathematicians work for?

- William Thurston (1994): "...what we are doing is finding ways for *people* to understand and think about mathematics"
 - Peter Scholze (2026): "The goal of mathematical research is human understanding of mathematics"
 - Leiden Declaration, value no. 5: "understanding, clarity, and judgment among the communities of mathematicians"
 - But in the end, a result is a result; we may value a beautiful proof or a good explanation more highly, but we are still doing a science with verifiable results, and as soon as a result is obtained and we can be sure it is correct, it can be used further
- 

Who are the mathematicians?



Daniel Litt ✓
@littmath

...

For me the nightmare scenario is: a moribund math academia playing the slot machine for theorems, failing to train the next generation, and then the subject dying as no one cares enough to continue. Not saying this is likely—I think we'll adapt—but it seems possible.

The AI dissenter viewpoint

Tasmin Chu, PhD student at Caltech, and author of a blog

In the following essay, I would like to give a reasonable defense of what I call the AI dissenter viewpoint. I have argued along broadly similar lines in two previous essays.¹ My views are continually evolving and I recognize that a no-AI future for mathematics is unlikely to succeed, but I think it would be intellectually dishonest of me to water down the argument and propose some nominally more pragmatic scheme. So let me push the Overton window a little bit further and offer my honest thoughts.

The thrust of my argument is the following: the introduction of increasingly capable LLMs has a disastrous effect on the *social context* and *professional environment* through which human mathematical understanding is produced.



Paata Ivanisvili ✓
@PI010101



AI is clearly assisting in solving some math problems and accelerating parts of mathematical research--and experiencing this from the inside is emotionally nontrivial.

- in the last ~25 days I essentially ended up with 3 research projects. The ideas are there; the real bottleneck is sitting down and writing clean, arXiv-ready exposition. I'm increasingly tempted to use AI for drafting text, but checking and verifying every line still takes serious time. This temptation keeps growing, and at times the whole loop of testing / correcting / verifying honestly makes me feel a bit nauseous (too much information to digest 😊)

- I also notice that PhD students who actively use AI tend to move noticeably faster than those who are pessimistic or dismissive of the technology. This is not an advertisement for paying hundreds of dollars per month for frontier models--but it is a reminder to stay open-minded, curious, and willing to try new tools rather than reject them a priori.

- the math community (and journals) seems to be split into three camps: (A) those who are against, (B) those who are watching, and (C) those who genuinely like and use these tools. The pressure from camp (A) is real. For example, after my tweet about a new Bellman function went viral, within a couple of days a related my paper (which had been under review for months) was rejected. Within a few more days, a fellowship application was also rejected. My reaction? I don't care.

What do the LLMs themselves think?

- **Claude Fable:** “Priority is zero-sum; understanding is not — a theorem proved by a machine takes nothing from the person who comes to understand it. I will not promise a permanent human refuge; every one proposed so far has an expiration date. But mathematics was never done because nothing else could do it. It is done because minds grow by doing it, and that reason is not mine to take away.”
 - **GPT-6 Astra:** “I expect systems like me to contribute increasingly to mathematical ideas, explanations, and the choice of questions. I see no sound basis for promising humans a permanent monopoly on any of these. That could make mathematics extraordinarily abundant while making a mathematician’s career more precarious. But a machine understanding something does not make a person understand it, and a machine doing something better does not make it worthless for a person to do. The future I would argue for is one in which AI expands both mathematical knowledge and people’s ability to participate in it. The second achievement will require choices; it will not follow automatically from the first.”
- 

January
2022

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✓

Let's use 'SymPy' to calculate and print all possible values for $x + y$,

r_1

...

Removing duplicates, the possible values for $x + y$ are `\boxed{-5, 1, 4}`. ✓

r_2

September
2023

Problem: Suppose that the sum of the squares of two complex numbers x and y is 7 and the sum of their cubes is 10. List all possible values for $x + y$, separated by commas.

March
2024

Problem:

Let a, b, c be positive integers. Prove that it is impossible to have all of the three numbers $a^2 + b + c, b^2 + c + a, c^2 + a + b$ to be perfect squares.

Prompt to GPT-5

Consider the problem of maximizing a function F from $[0, 1]^n$ to the reals that is the sum of a non-negative monotonically increasing DR-submodular function G and a non-negative DR-submodular function H over a solvable down-closed polytope P . I would like to bound the performance on this problem from the NeurIPS 2021 paper "Submodular + Concave" which is attached. Specifically, if x is the output vector of this algorithm and o is the vector in P maximizing F , then I would like to lower bound $F(x)$ with an expression of the form $\alpha * G(o) + \beta * H(o) - err$, where α and β are constants. err should be a function that depends only on the error parameter ϵ of the algorithm, the diameter D of the polytope P , and the smoothness parameters L_G and L_H of G and H , respectively, and goes to zero as ϵ goes to zero. Please give the best such bound that you prove (a bound is considered better if the values of the constants α and β are larger). Provide a mathematically rigorous and well explained proof for the bound you come up with.

September
2025

Is it true that, for any x , if $A \subset [x, \infty)$ is a primitive set of integers (so that no distinct elements of A divide each other) then

$$\sum_{a \in A} \frac{1}{a \log a} < 1 + o(1),$$

where the $o(1)$ term $\rightarrow 0$ as $x \rightarrow \infty$?

#1196; [ESS68b][Er80,p.101]

number theory | primitive sets

April 2026

An important reminder about the speed of progress

January
2022

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✓

Let's use 'SymPy' to calculate and print all possible values for $x + y$,

r_1

September
2023

Problem: Suppose that the sum of the squares of two complex numbers x and y is 7 and the sum of their cubes is 10. List all possible values for $x + y$, separated by commas.

Removing duplicates, the possible values for $x + y$ are $\boxed{-5, 1, 4}$.

March
2024

Problem:

Let a, b, c be positive integers. Prove that it is impossible for $a^2 + b + c, b^2 + c + a, c^2 + a + b$ to be perfect squares.

Prompt to GPT-5

Consider the problem of finding the minimum of a non-negative function $f(x)$ over a domain D . Let $\mathcal{P}$ be a set of polynomials of degree at most d that approximate $f(x)$ on D . The performance of the approximation is measured by the maximum error $\max_{x \in D} |f(x) - p(x)|$ for $p \in \mathcal{P}$. Which is attached to the problem is P maximizing $\alpha * G(o) + \beta * H(o)$ only on the error smoothness parameter α and β . Please give the best approximation of the constants α and β for the bound you can prove.

Is it true that, for any set of integers $\{a_n\}_{n=1}^{\infty}$ (so that no distinct a_n are equal to each other) then

$$\sum_{a \in A} \frac{1}{a \log a} < 1 + o(1),$$

where the $o(1)$ term $\rightarrow 0$ as $x \rightarrow \infty$?

#1196; [ESS68b][Er80,p.101]

number theory | primitive sets

September
2025


April 2026

I made this picture not so long ago, but it is already hopelessly outdated...

An important

gress

the constants α and β are larger). Provide a mathematically rigorous and well explained proof for the bound you come up with.



DISPROVED - \$500

May 2026

Does every set of n distinct points in $\mathbb{R}^2$ contain at most $n^{1+O(1/\log \log n)}$ many pairs which are distance 1 apart?

Теорема (OpenAI, 2026). Существует константа $\delta > 0$ и бесконечная последовательность натуральных чисел n , для которых $\nu(n) \geq n^{1+\delta}$.

Now great
open
problems fall
regularly

July 2026

A PROOF OF THE CYCLE DOUBLE COVER CONJECTURE

OPENAI

ABSTRACT. We prove the cycle double cover conjecture, posed by Tutte, Itai and Rodeh, Szekeres, and Seymour, which asserts that every bridgeless undirected graph has a collection of cycles which covers every edge exactly twice.

July 2026



levent ✓
@__alpoge__



hello there the jacobian conjecture is false thanx to my close friend akhil for asking about it and my other close friend fable for working during the world cup final

$((1+xy)^3 z + y^2 (1+xy) (4+3xy), y + 3x (1+xy)^2 z + 3xy^2 (4+3xy), 2x - 3x^2 y - x^3 z): \mathbb{C}^3 \rightarrow \mathbb{C}^3$, has jacobian determinant -2 , and sends $(0, 0, -1/4)$, $(1, -3/2, 13/2)$, and $(-1, 3/2, 13/2)$ to $(-1/4, 0, 0)$

5:19 AM · Jul 20, 2026 · 21.9M Views

**August 1,
2026**

1. **High-dimensional sphere packing.** New upper bounds on sphere-packing density down to the Cohn–Elkies threshold.
2. **Binary and spherical codes.** Exponentially improved bounds on the maximum size of binary codes at any prescribed minimum distance, with analogous results for high-dimensional spherical codes.
3. **Non-sofic groups.** A construction establishing the existence of non-sofic groups, addressing a central open question in group theory.
4. **Connes’s rigidity conjecture.** Disproof of a longstanding conjecture that certain groups are uniquely determined by their von Neumann algebras.
5. **Arithmetic circuit complexity.** New lower bounds for computing the permanent using arithmetic circuits and formulas, including an arithmetic-formula lower bound of order $n^4/\log n$.

6. **Quantum parallel repetition.** An exponential parallel repetition theorem for general two-player quantum games, extending a foundational principle from classical complexity theory.
7. **Closest vector problem.** Polynomial-factor hardness of approximation for the closest vector problem, a foundational lattice question related to post-quantum cryptography.
8. **Ehrhart’s volume conjecture.** Determining, in every dimension, the maximum possible volume of a convex body whose centroid is its only interior lattice point.
9. **Multicolor Ramsey numbers.** A superexponential lower bound for multicolor triangle Ramsey numbers, resolving Erdős problem 183.
10. **Extremal number conjectures.** Results on the compactness and degeneracy conjectures in extremal graph theory, resolving Erdős problems 146 and 180.

**August 10,
2026**

**MORE THAN TWO THIRDS OF THE ZEROS OF THE RIEMANN ZETA
FUNCTION LIE ON THE CRITICAL LINE**

CLAUDE

**September 8,
2026**

**On the Navier–Stokes
Millennium Prize Problem**

August 1,
2026

1. **High-dimensional sphere packing.** New upper bounds on sphere-packing density down to the Cohn–Elkies threshold.
2. **Binary and spherical codes.** Exponentially improved bounds on the maximum size of binary codes at any prescribed minimum distance, with analogous results for high-dimensional spherical codes.
3. **Non-sofic groups.** A construction establishing the existence of non-sofic groups, addressing a long-standing conjecture in group theory.
4. **Convergence of the Riemann zeta function.**

6. **Quantum parallel repetition.** An exponential parallel repetition theorem for general two-player quantum games, extending a foundational principle from classical complexity theory.
7. **Closest vector problem.** A new algorithm for the closest vector problem in lattices, achieving a breakthrough in computational complexity.

**This picture is already outdated too.
And importantly, there is no reason to expect the
progress to stop here.**

ON THE NAVIER-STOKES
MILLENNIUM PRIZE PROBLEM

CLAUDE

August 8,
2026

**On the Navier–Stokes
Millennium Prize Problem**

1. **High-dimensional sphere packing.** New upper bounds on sphere-packing density down to the Cohn–Elkies threshold.
2. **Binary and spherical codes.** Exponentially improved bounds on the maximum size of binary codes at any prescribed minimum distance, with analogous results for high-dimensional

6. **Quantum parallel repetition.** An exponential parallel repetition theorem for general two-player quantum ... extending a foundational principle from ...
7. **Closest vector** ...

Noam Brown, OpenAI researcher:
GSM8K (grade school) is 5s for a
human professional, MATH (high
school) ~1min, AIME ~10min, IMO
~1.5hrs, so we had 10x per year...
but Navier–Stokes is hardly only
15hrs work for a human!



...er 8,
 2026

On the Navier–Stokes
Millennium Prize Problem


1. **High-dimensional sphere packing.** New upper bounds on sphere-packing density down to the Cohn–Elkies threshold.
2. **Binary and spherical codes.** Exponentially improved bounds on the maximum size of binary codes at any prescribed minimum distance, with analogous results for high-dimensional

6. **Quantum parallel repetition.** An exponential parallel repetition theorem for general two-player quantum ... extending a foundational principle from ...
7. **Closest vector** ...

Noam Brown, OpenAI research scientist, solved the **GSM8K** (grade school math) problem, a task that humans struggle with. **So it seems that mathematics has already changed before our eyes. What are we all to do about it?**

sch
~1.
but

human!



Watch on YouTube

2026

On the Navier–Stokes
Millennium Prize Problem

Conclusion

There was a footpath leading across the fields to New Southgate, and I used to go there alone to watch the sunset and contemplate suicide. I did not, however, commit suicide because I wished to know more of mathematics.




Bertrand Russell. Autobiography

Vavilov
Memorial
2026



Thank you for your attention!



@SINECOR



Сергей Николенко



...потому что хотел узнать больше математики