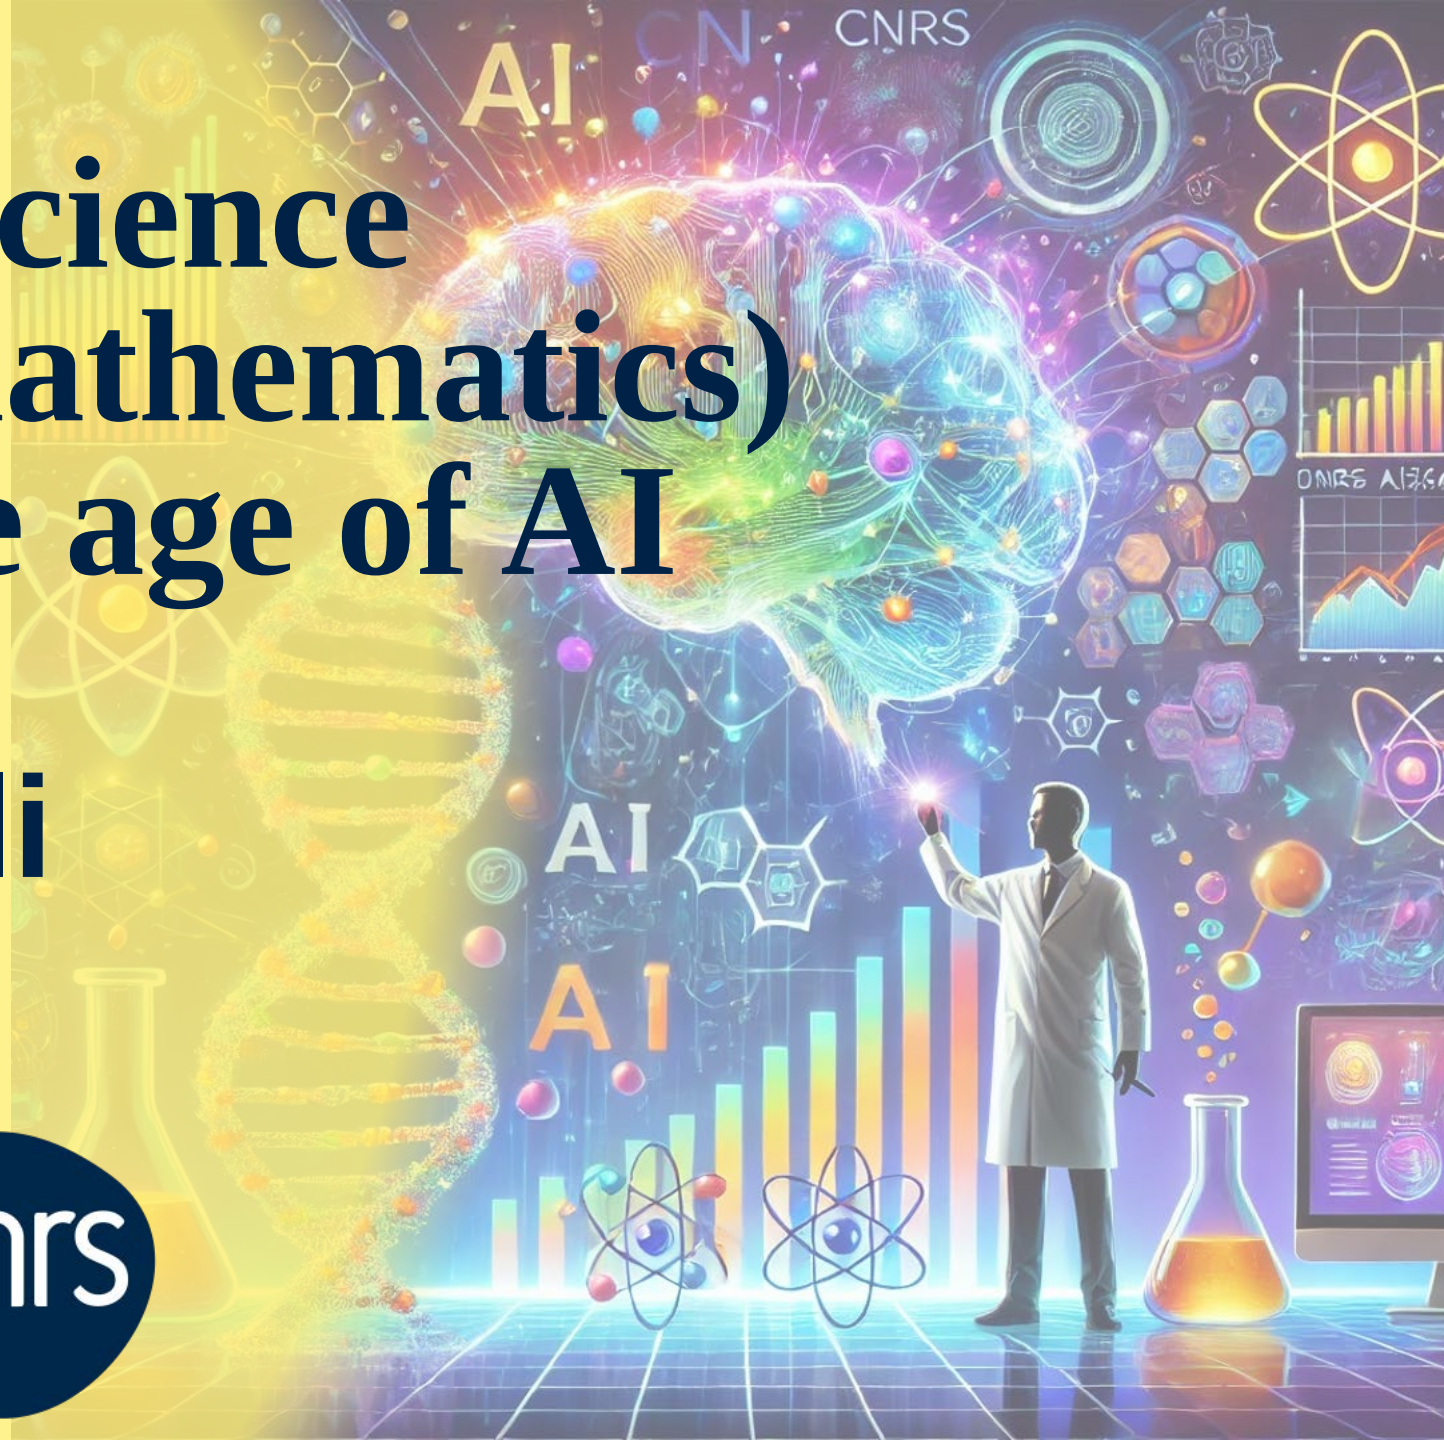
Science (and mathematics) in the age of AI

Jalal Fadili

AIS2AI
AI for science, science for AI

cnrs



AI and Science

AI and Science

- Scientists strive to **understand** how **nature** and **society** work.

AI and Science

- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.

AI and Science

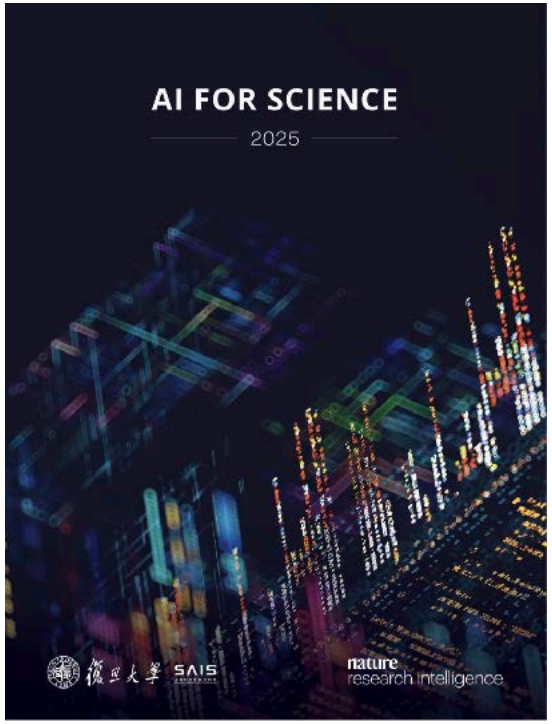
- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.
- Data **collection** and **analysis** are central to the scientific method: scientists have long used statistical techniques to **extract information from data, test and (in)validate hypotheses**, etc.

AI and Science

- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.
- Data **collection** and **analysis** are central to the scientific method: scientists have long used statistical techniques to **extract information from data, test and (in)validate hypotheses**, etc.
- **Science capitalize on AI: AI** provides new tools and is becoming a **scientific discovery paradigm** impacting the way research is produced across disciplines.

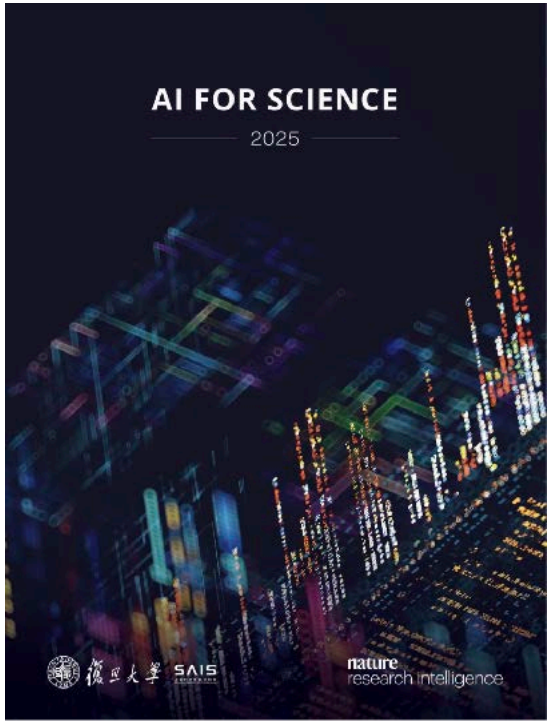
AI and Science

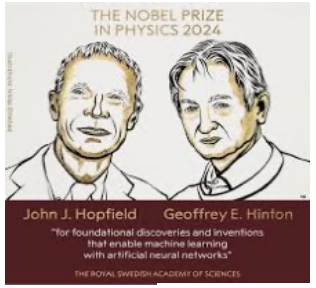
- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.
- Data **collection** and **analysis** are central to the scientific method: scientists have long used statistical techniques to **extract information from data, test and (in)validate hypotheses**, etc.
- **Science capitalize on AI: AI** provides new tools and is becoming a **scientific discovery paradigm** impacting the way research is produced across disciplines.
- **Science contribute to AI:** AI is also **nurtured** by other **disciplines**: physics, mathematics, neuroscience, etc.

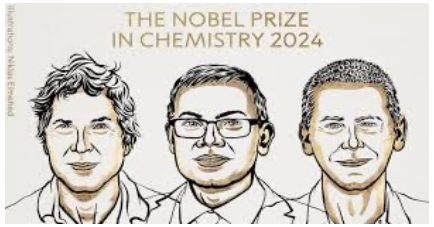


AI and Science

- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.
- Data **collection** and **analysis** are central to the scientific method: scientists have long used statistical techniques to **extract information from data, test and (in)validate hypotheses**, etc.
- **Science capitalize on AI: AI** provides new tools and is becoming a **scientific discovery paradigm** impacting the way research is produced across disciplines.
- **Science contribute to AI:** AI is also **nurtured** by other **disciplines**: physics, mathematics, neuroscience, etc.
- 2024 **Nobel Prizes**: Nobel Prize for Physics (neural networks) and The Nobel Prize for Chemistry (AlphaFold).



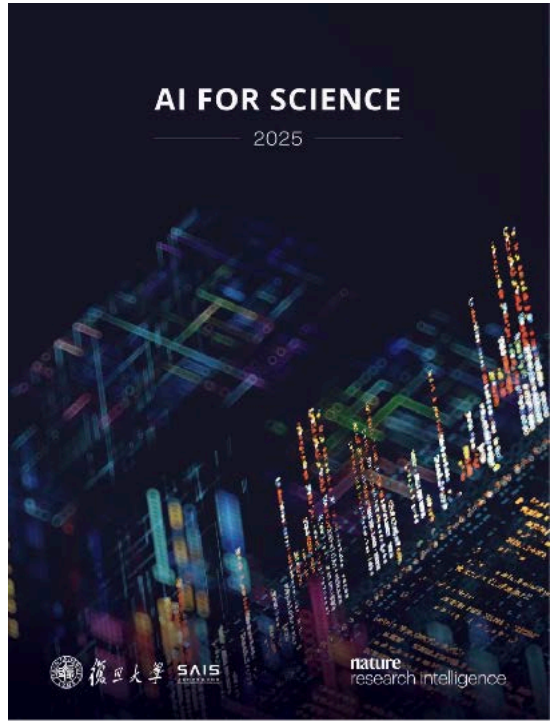




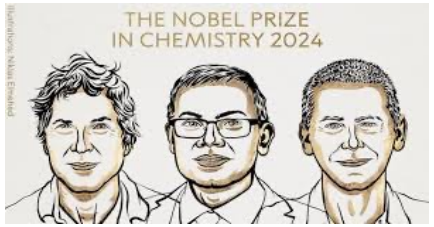


AI and Science

- Scientists strive to **understand** how **nature** and **society** work.
- They formulate **hypotheses, conjectures**, design **reasoning**, **experiments** and collect **data**, analyze them to better understand **natural, physical and social phenomena**.
- Data **collection** and **analysis** are central to the scientific method: scientists have long used statistical techniques to **extract information from data, test and (in)validate hypotheses**, etc.
- **Science capitalize on AI: AI** provides new tools and is becoming a **scientific discovery paradigm** impacting the way research is produced across disciplines.
- **Science contribute to AI:** AI is also **nurtured** by other **disciplines**: physics, mathematics, neuroscience, etc.
- 2024 **Nobel Prizes**: Nobel Prize for Physics (neural networks) and The Nobel Prize for Chemistry (AlphaFold).
- Spectacular illustration of the **cross-fertilization** and increasingly inextricable lin between **AI and science**.

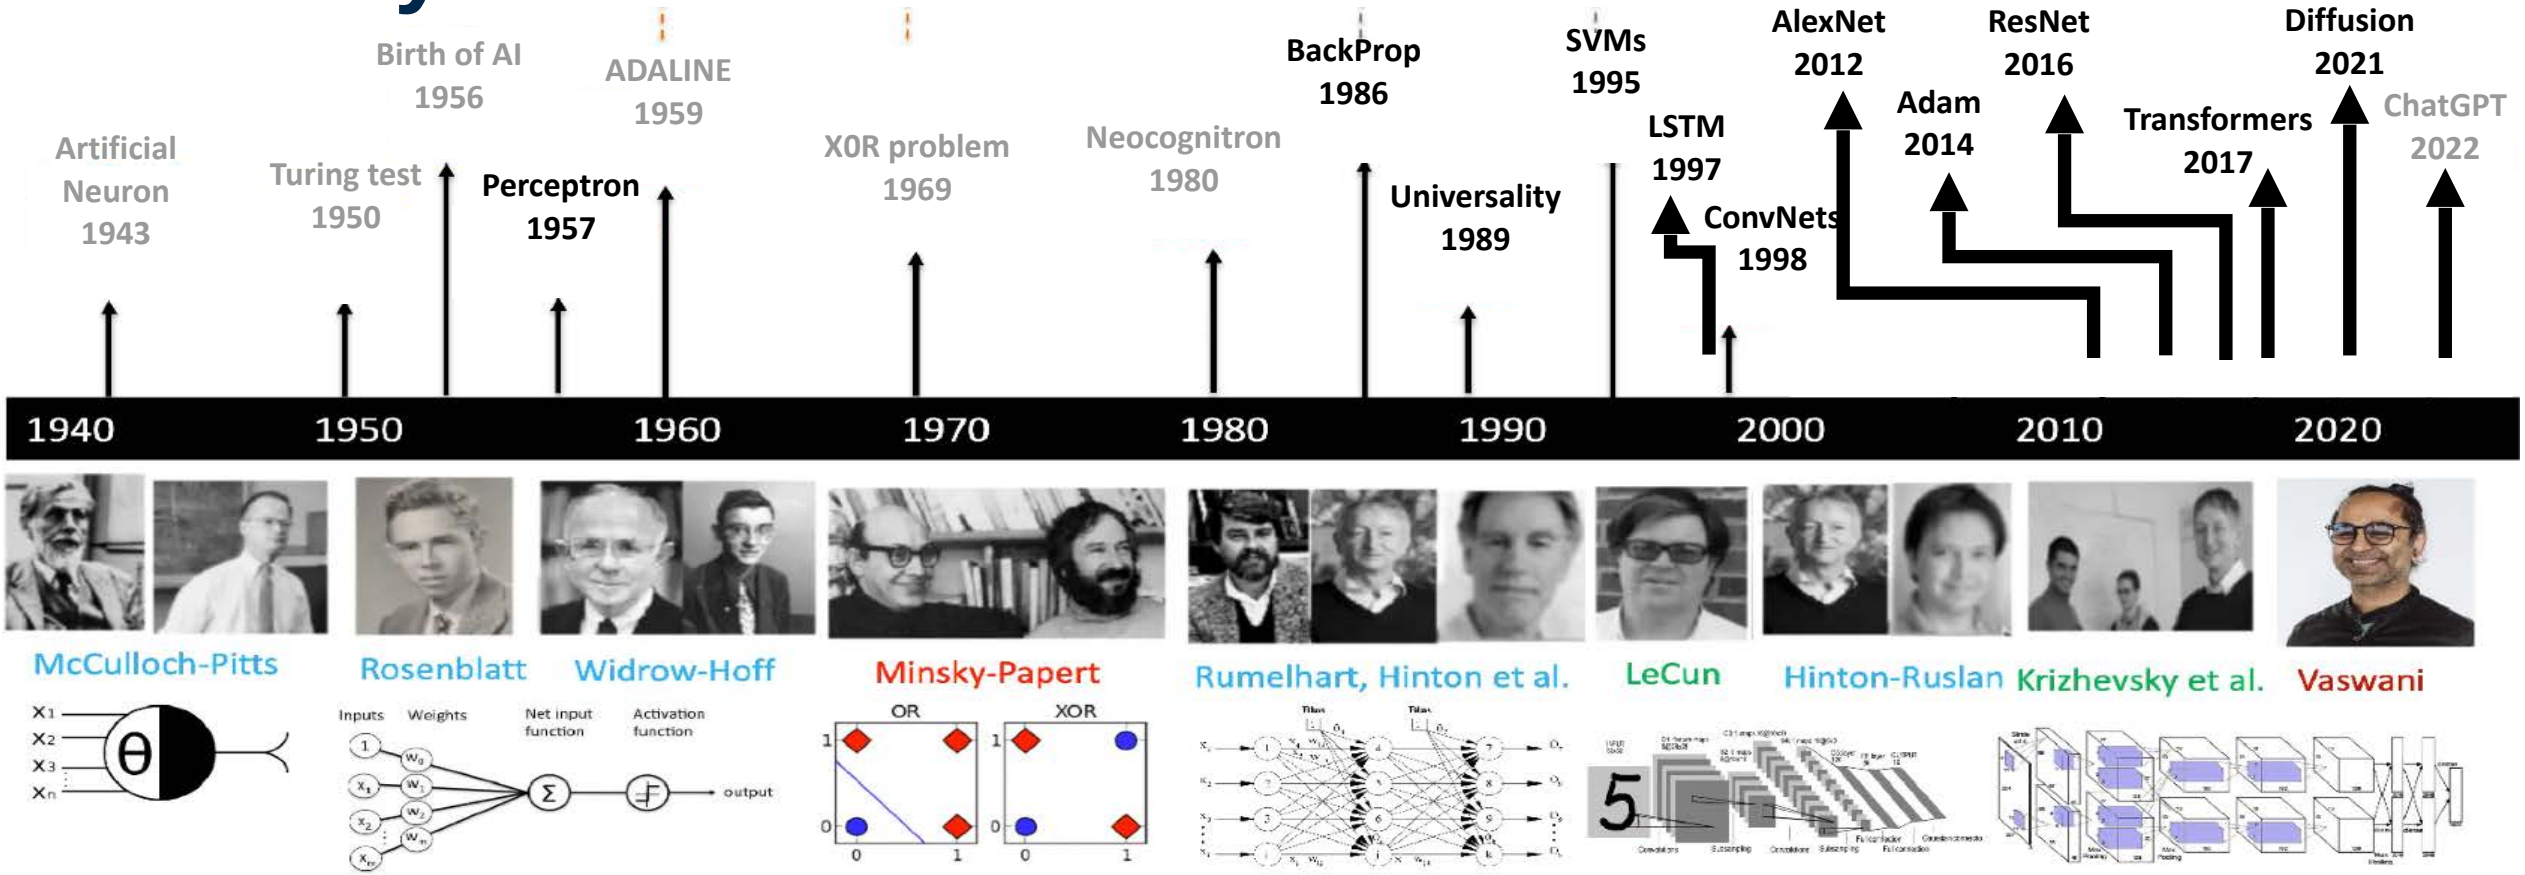




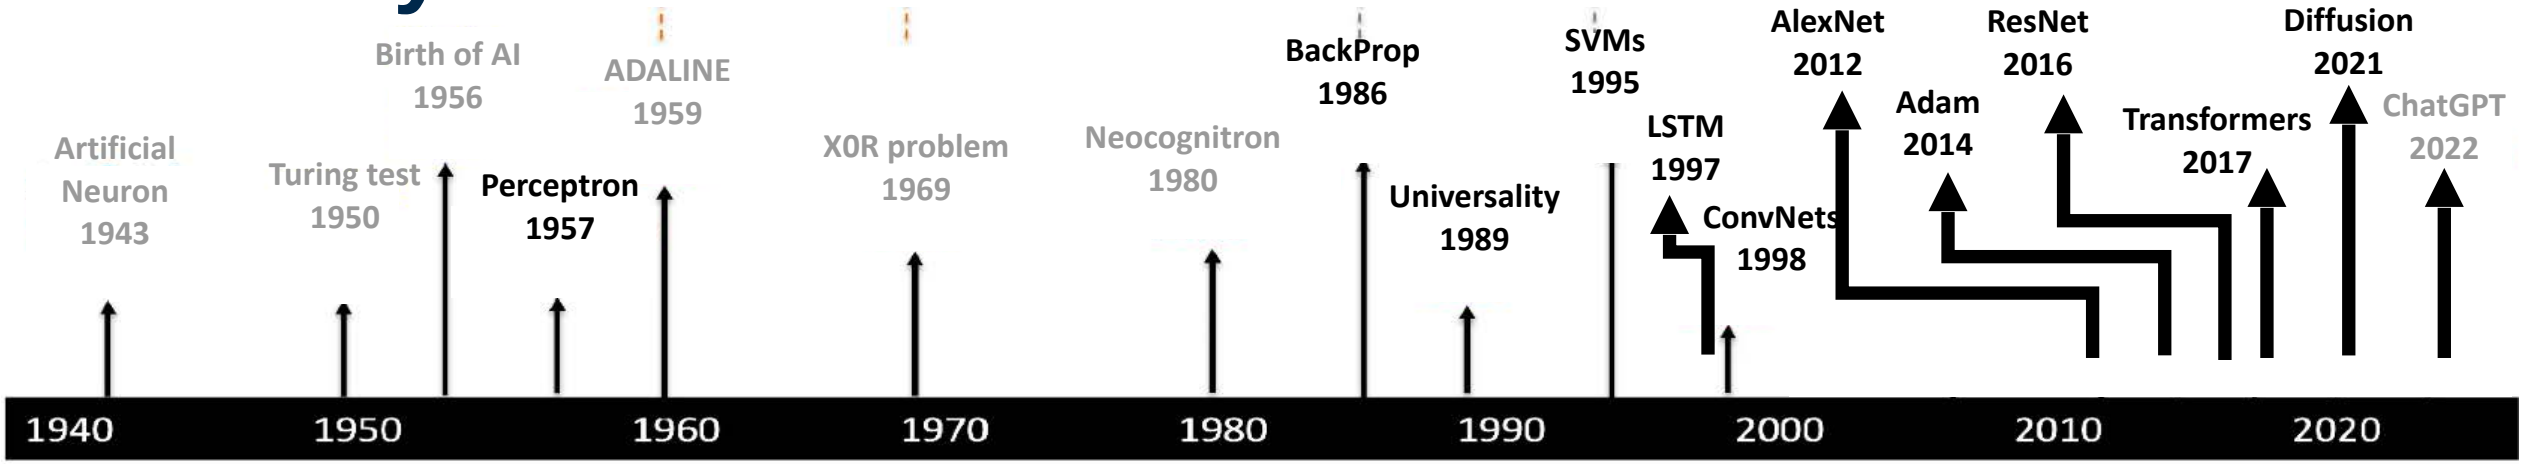




The many revolutions of AI

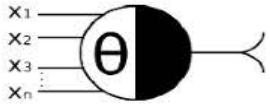


The many revolutions of AI



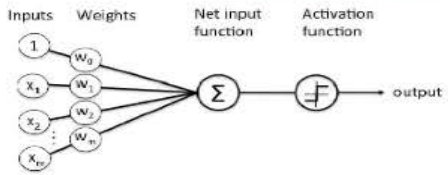


McCulloch-Pitts





Rosenblatt

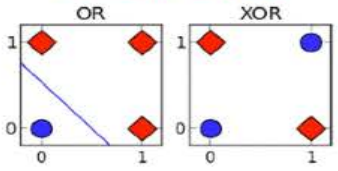




Widrow-Hoff

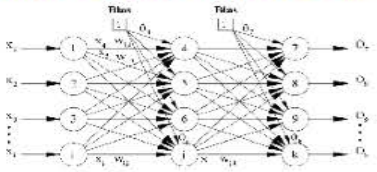


Minsky-Papert



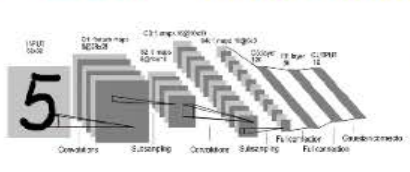


Rumelhart, Hinton et al.



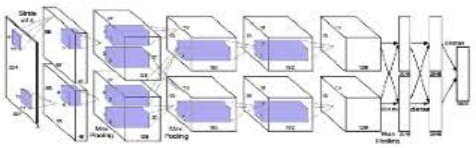


LeCun





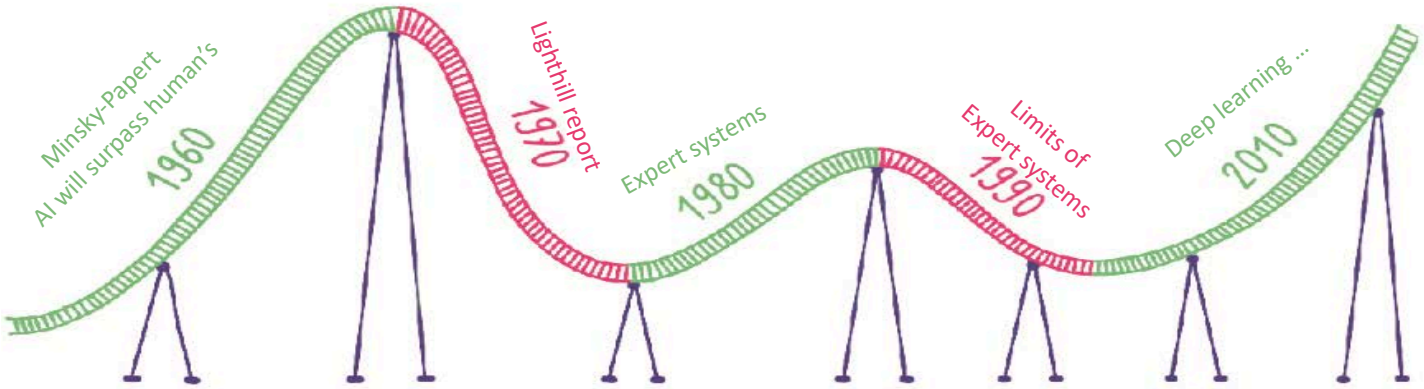
Hinton-Ruslan



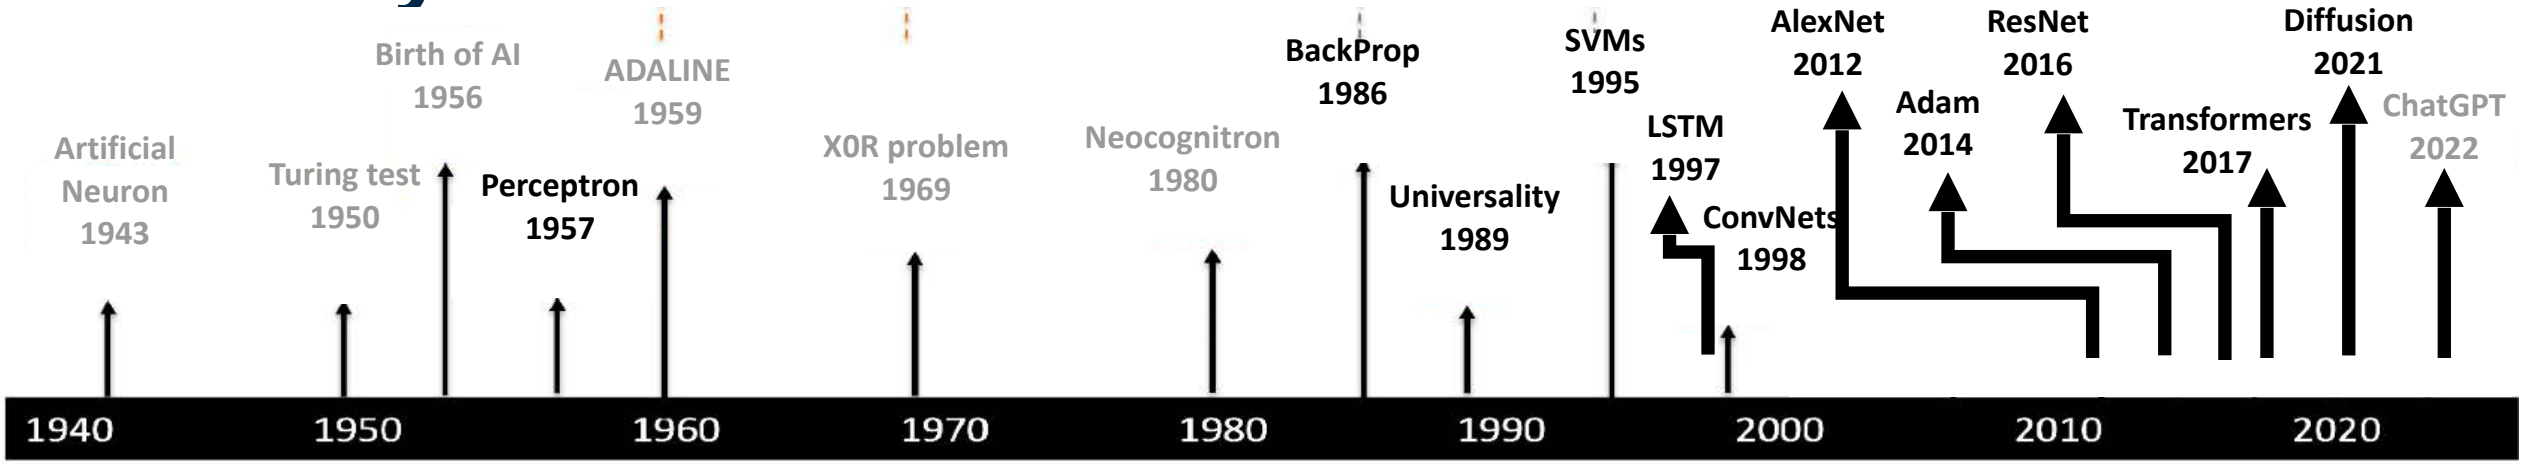
Krizhevsky et al.



Vaswani

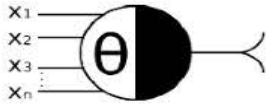


The many revolutions of AI



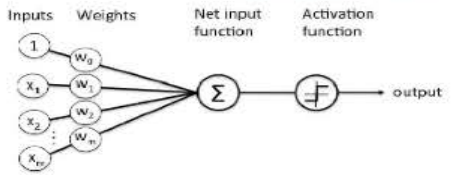


McCulloch-Pitts





Rosenblatt

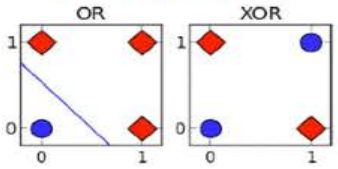




Widrow-Hoff

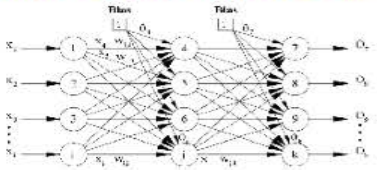


Minsky-Papert



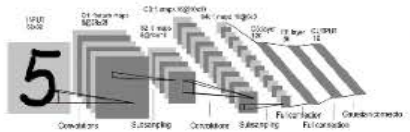


Rumelhart, Hinton et al.





LeCun

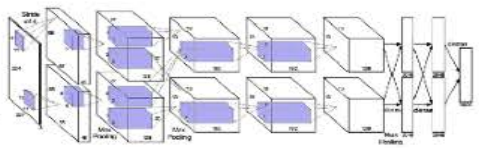




Hinton-Ruslan

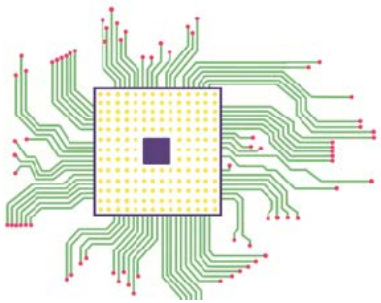


Krizhevsky et al.

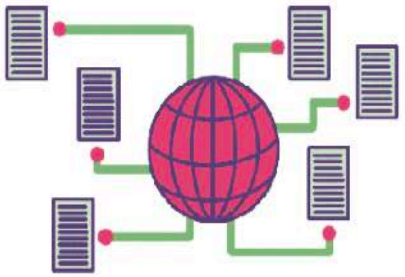




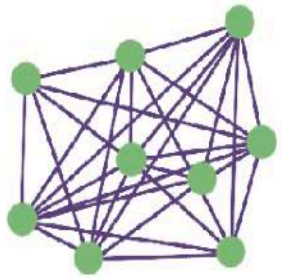
Vaswani



Computing power



Massive data



Fundamental research
New models, architectures, etc.

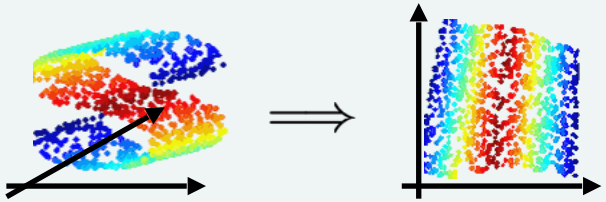


2018 Turing Award

(Un)supervised machine learning

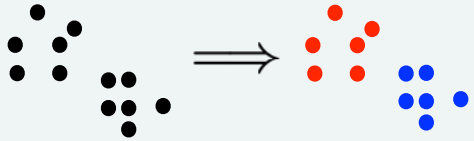
Unsupervised learning

Dimension reduction



PCA, t-SNE, UMAP

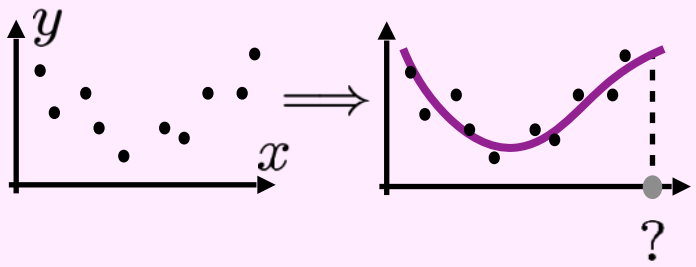
Clustering



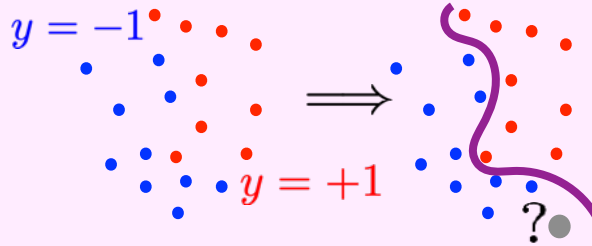
k-means, DBScan

Supervised learning

Regression



Classification



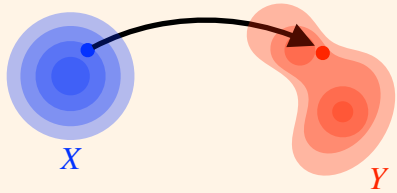
From (un)supervised to Generative AI

Diffusion models

Image:

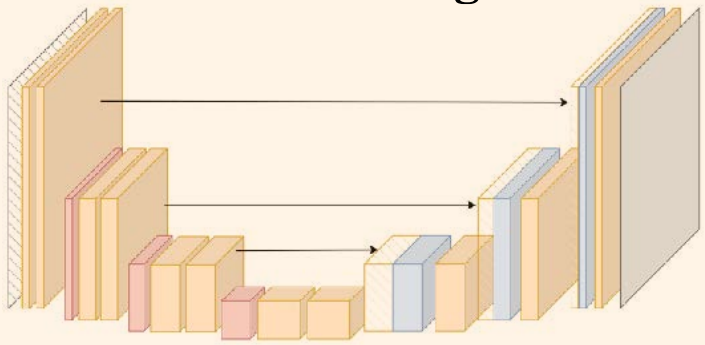
pre-training: denoising (score matching).

generation: iterative denoising





Architecture:
U-Net/transformers



DALL·E 2

stability.ai

Midjourney

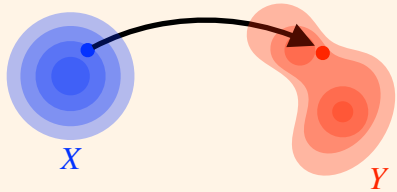
From (un)supervised to Generative AI

Diffusion models

Image:

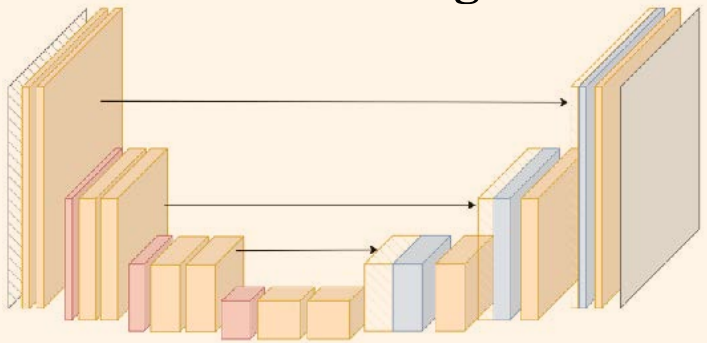
pre-training: denoising (score matching).

generation: iterative denoising





Architecture:
U-Net/transformers



DALL·E 2 stability.ai Midjourney

Large Language Models

Text:

pre-training: next-token prediction (teacher forcing)

generation: auto-regressive (logistic models on tokens)

L'École normale supérieure6, appelée aussi «de la rue d'Ulm», «Normale Sup'», est l'une des institutions universitaires et de recherche les plus prestigieuses et les plus sélectives de France, spécialisée en lettres et en sciences.

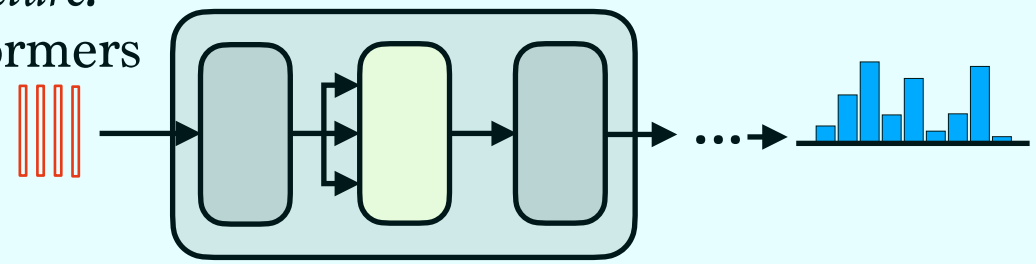
L'École normale supérieure6, appelée aussi «de la rue d'Ulm», «Normale Sup'», est l'une des institutions universitaires et de recherche les plus prestigieuses et les plus sélectives de France, spécialisée en lettres et en sciences.

L'École normale supérieure6, appelée aussi «de la rue d'Ulm», «Normale Sup'», est l'une des institutions universitaires et de recherche les plus prestigieuses et les plus sélectives de France, spécialisée en lettres et en sciences.

Masking

Next token prediction

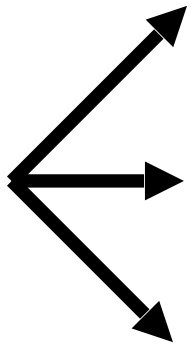
Architecture:
transformers



ChatGPT MISTRAL AI_ Claude

From text generation to generative AI in sciences

Terminology : « Large Language Model » (LLM)
/
« foundation models»



Large neural networks

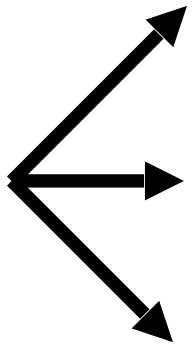
Trained to generate data
… or solve problems

Can be fine-tuned at low cost

From text generation to generative AI in sciences

Terminology :

« Large Language Model » (LLM)
/
« foundation models»



Large neural networks

Trained to generate data
… or solve problems

Can be fine-tuned at low cost

Prouve que le produit scalaire matriciel $\langle S, T \rangle$ peut s'écrire $\mathrm{tr}(ST^t)$.

LE CHAT MISTRAL_

Pour prouver que le produit scalaire matriciel $\langle S, T \rangle$ peut s'écrire sous la forme $\mathrm{tr}(ST^t)$, considérons les définitions suivantes : le produit scalaire entre deux matrices S et T de taille $n \times n$ est donné par $\left\langle S, T \right\rangle = \sum_{i=1}^{n} \sum_{j=1}^{n} s_{ij} t_{ij}$. La trace du produit ST^t est définie comme $\mathrm{tr}(ST^t) = \sum_{i=1}^{n} (ST^t)_{ii}$. En développant l'élément diagonal, nous avons $(ST^t)_{ii} = \sum_{j=1}^{n} s_{ij} t_{ij}$, ce qui donne

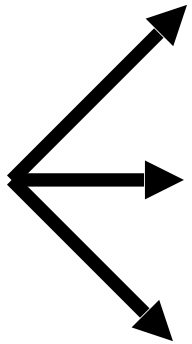
$$\mathrm{tr}(ST^t) = \sum_{i=1}^{n} \sum_{j=1}^{n} s_{ij} t_{ij} = \left\langle S, T \right\rangle.$$

Ainsi, $\langle S, T \rangle = \mathrm{tr}(ST^t)$.

From text generation to generative AI in sciences

Terminology :

« Large Language Model » (LLM)
/
« foundation models»



Large neural networks

Trained to generate data
… or solve problems

Can be fine-tuned at low cost

Prouve que le produit scalaire matriciel $\langle S, T \rangle$ peut s'écrire $\mathrm{tr}(ST^t)$.

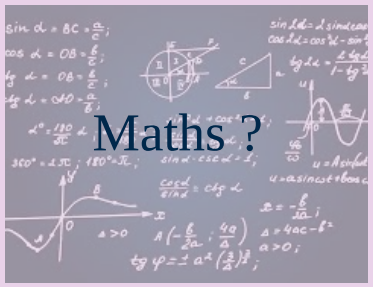
LE CHAT MISTRAL_

Pour prouver que le produit scalaire matriciel $\langle S, T \rangle$ peut s'écrire sous la forme $\mathrm{tr}(ST^t)$, considérons les définitions suivantes : le produit scalaire entre deux matrices S et T de taille $n \times n$ est donné par $\left\langle S, T \right\rangle = \sum_{i=1}^{n} \sum_{j=1}^{n} s_{ij} t_{ij}$. La trace du produit ST^t est définie comme $\mathrm{tr}(ST^t) = \sum_{i=1}^{n} (ST^t)_{ii}$. En développant l'élément diagonal, nous avons $(ST^t)_{ii} = \sum_{j=1}^{n} s_{ij} t_{ij}$, ce qui donne

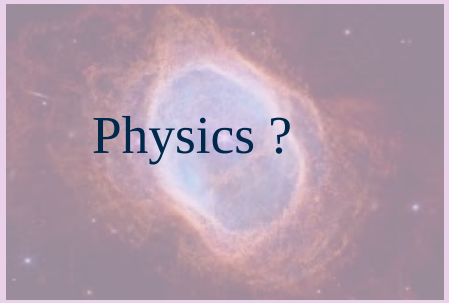
$$\mathrm{tr}(ST^t) = \sum_{i=1}^{n} \sum_{j=1}^{n} s_{ij} t_{ij} = \left\langle S, T \right\rangle.$$

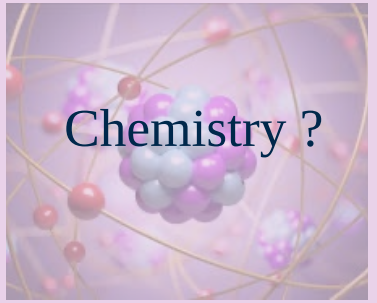
Ainsi, $\left\langle S, T \right\rangle = \mathrm{tr}(ST^t)$.

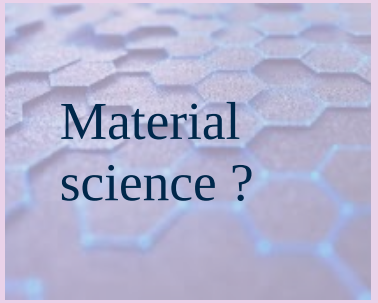
How to extend these methods beyond text and images ?









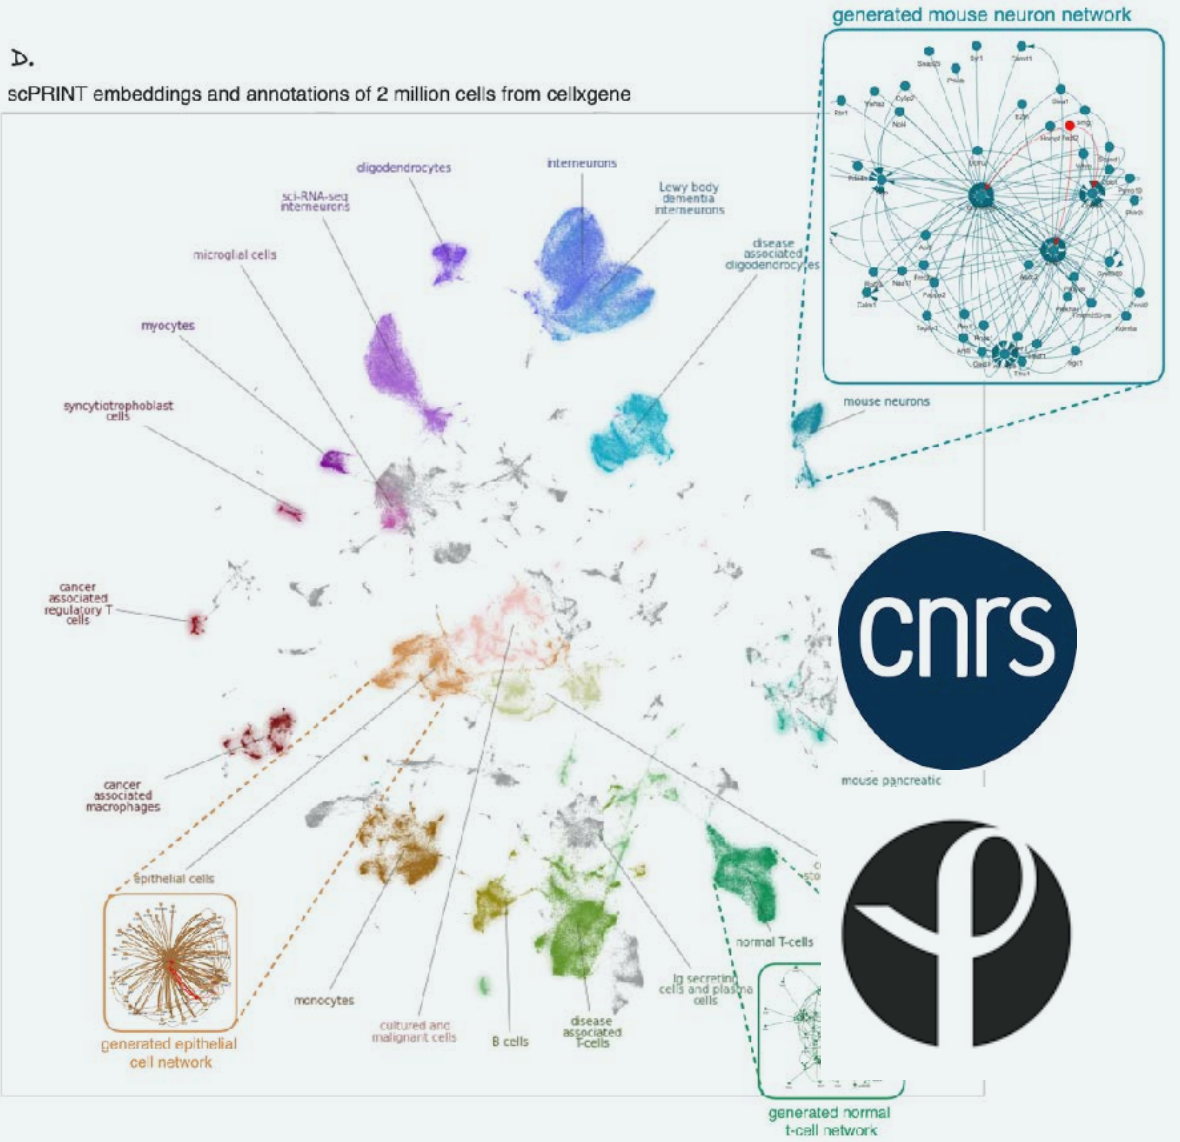


…

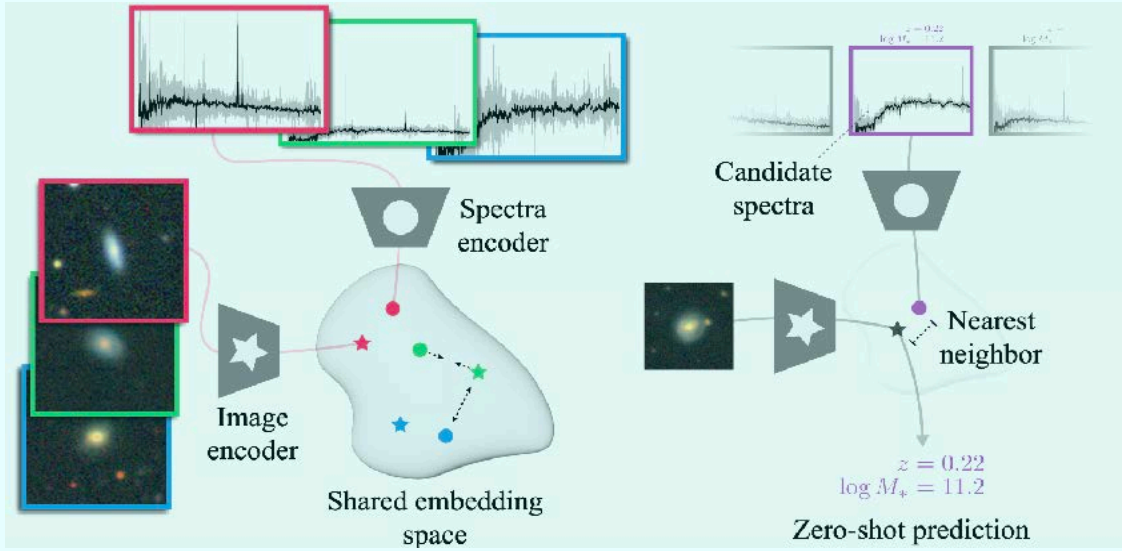
A whole program of theoretical, practical and societal questions

Foundation models in sciences

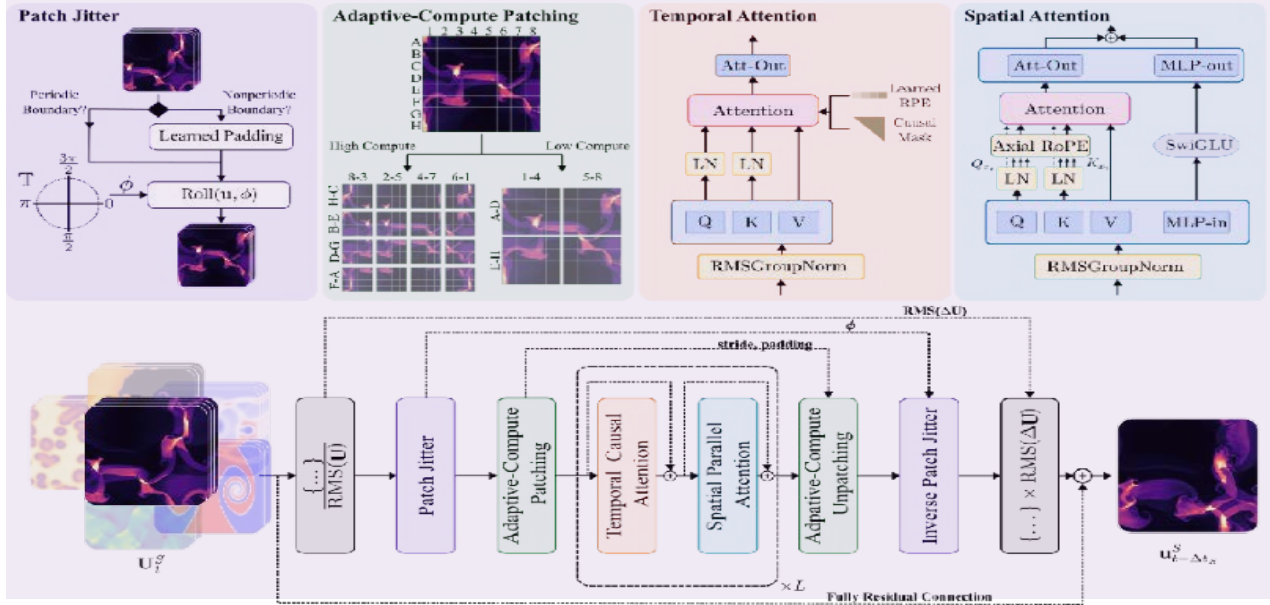
scPRINT: Large Cell Model for scRNAseq data



AstroCLIP: A Cross-Modal Foundation Model for Galaxies



Walrus: Cross-domain Foundation Model for Continuum Dynamics



From generative AI to agentic AI in sciences

Denario

Multi-agent system that generates full papers across astrophysics, biology, chemistry, and more. **NeurIPS 2025** competition winner.

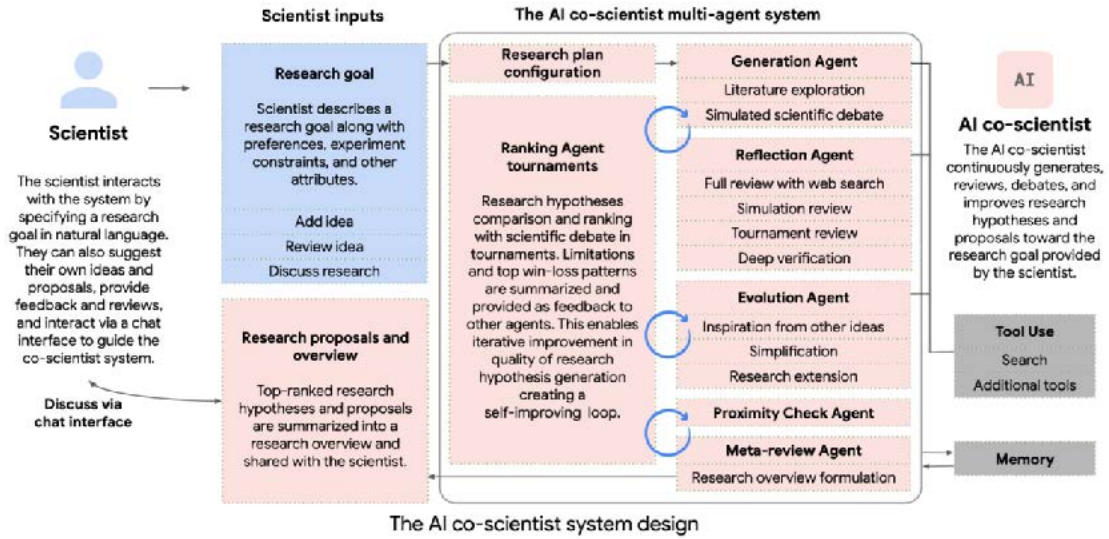
From generative AI to agentic AI in sciences

Denario

Multi-agent system that generates full papers across astrophysics, biology, chemistry, and more. **NeurIPS 2025** competition winner.

DeepMind **AI Co-Scientist**

Generates novel hypotheses, reviews literature, and designs experiments in a closed-loop **self-improving cycle**.



The AI co-scientist system design

From generative AI to agentic AI in sciences

Denario

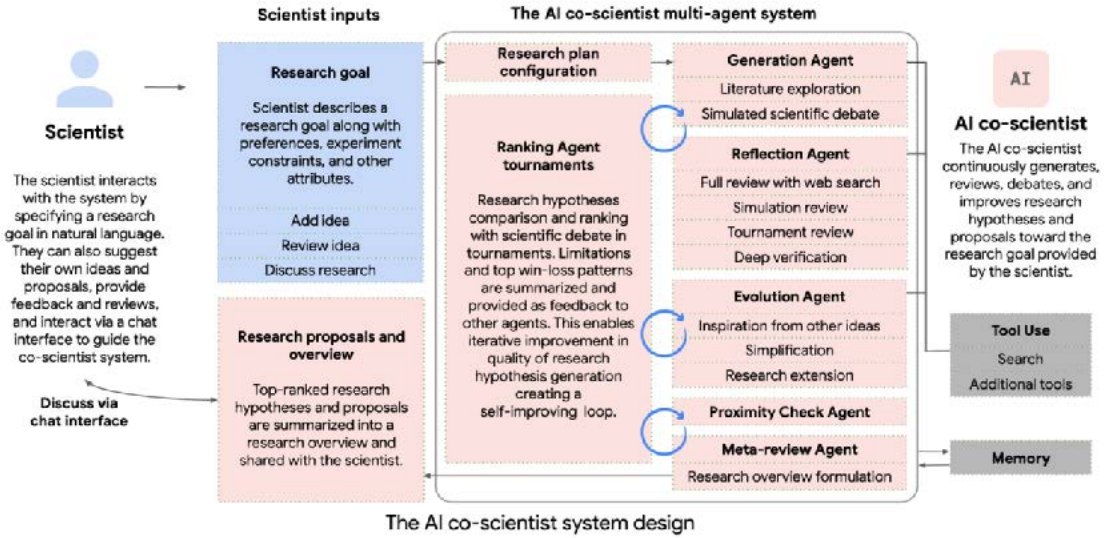
Multi-agent system that generates full papers across astrophysics, biology, chemistry, and more. **NeurIPS 2025** competition winner.

DeepMind AI Co-Scientist

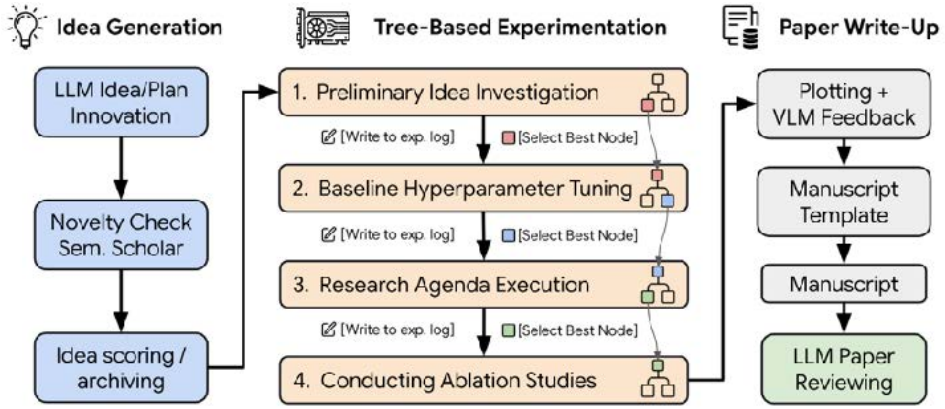
Generates novel hypotheses, reviews literature, and designs experiments in a closed-loop **self-improving cycle**.

sakana.ai "The AI Scientist"

End-to-end system: **ideation → coding → experiments → writing → peer review**, all automated.



The AI co-scientist system design



From generative AI to agentic AI in sciences

Denario

Multi-agent system that generates full papers across astrophysics, biology, chemistry, and more. **NeurIPS 2025** competition winner.

DeepMind AI Co-Scientist

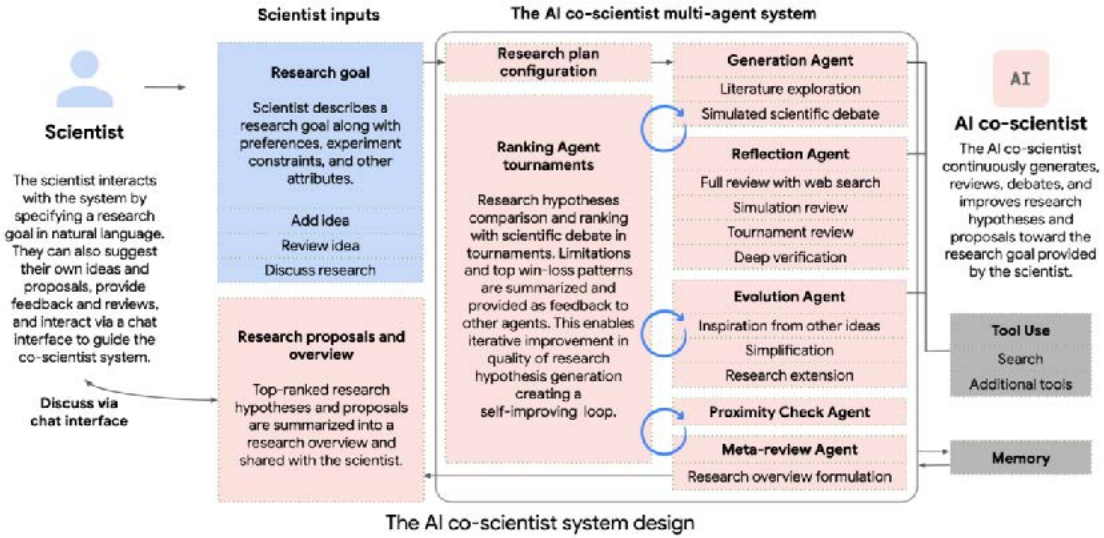
Generates novel hypotheses, reviews literature, and designs experiments in a closed-loop **self-improving cycle**.

sakana.ai "The AI Scientist"

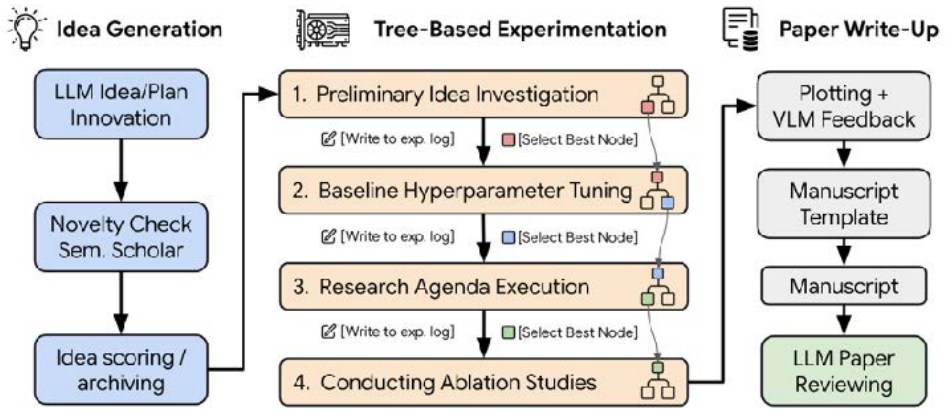
End-to-end system: **ideation → coding → experiments → writing → peer review**, all automated.

Edison "Kosmos"

12-hour autonomous runs combining data analysis and literature search. Proposed a novel mechanism for **Type 2 diabetes risk** from public genetics data.



The AI co-scientist system design



From generative AI to agentic AI in sciences

Denario

Multi-agent system that generates full papers across astrophysics, biology, chemistry, and more. **NeurIPS 2025** competition winner.

DeepMind AI Co-Scientist

Generates novel hypotheses, reviews literature, and designs experiments in a closed-loop **self-improving cycle**.

sakana.ai "The AI Scientist"

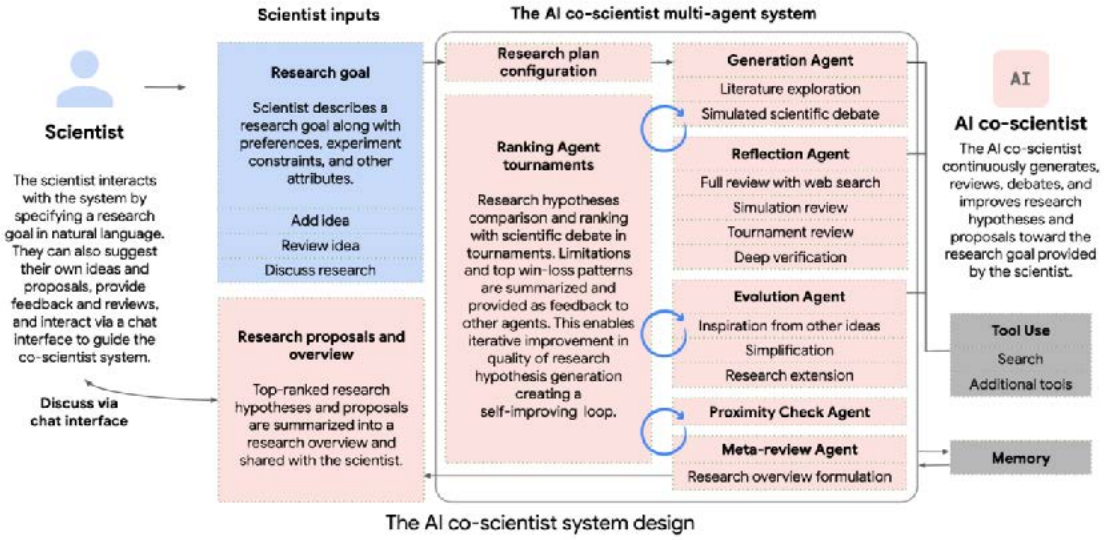
End-to-end system: **ideation → coding → experiments → writing → peer review**, all automated.

Edison "Kosmos"

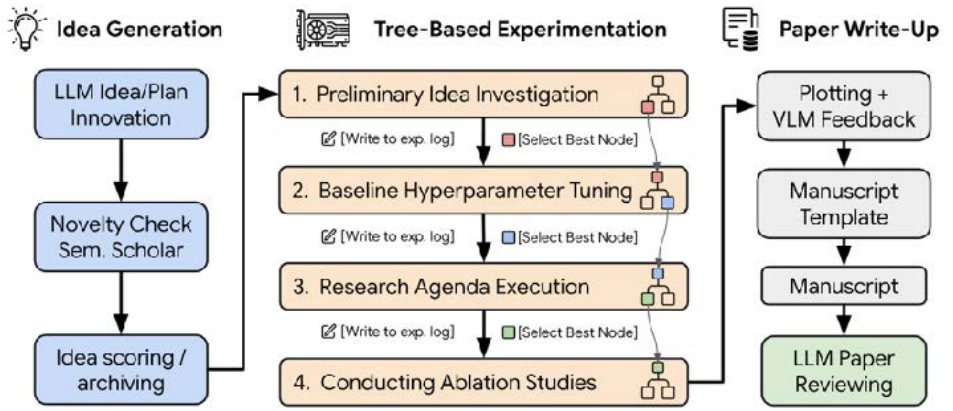
12-hour autonomous runs combining data analysis and literature search. Proposed a novel mechanism for **Type 2 diabetes risk** from public genetics data.

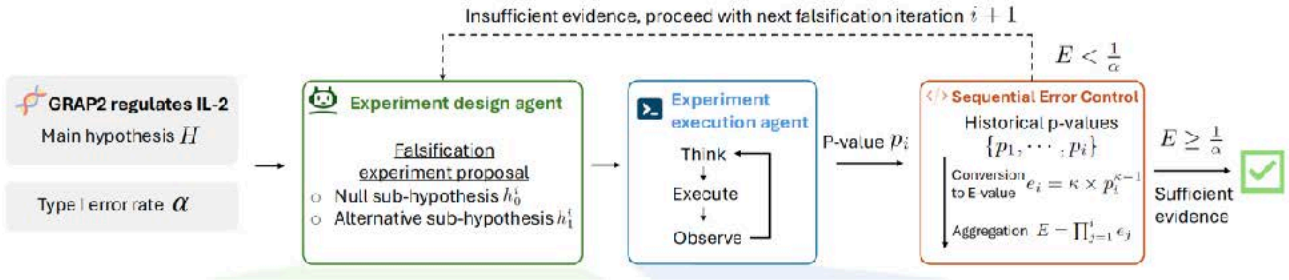
Popper

Automated Hypothesis Validation with **Karl Popper's principle of falsification.**



The AI co-scientist system design





From generative AI to agentic AI in sciences

Denari

Multi-agent system that generates full pa
biology, chemistry, and more. **NeurIPS 2**

DeepM

Generates novel hypotheses, reviews literature, and designs experiments in a closed-loop **self-improving cycle**.

sakana.ai **"The AI Scientist"**

End-to-end system: **ideation → coding → experiments → writing → peer review**, all automated.

Edison **"Kosmos"**

12-hour autonomous runs combining data analysis and literature search. Proposed a novel mechanism for **Type 2 diabetes risk** from public genetics data.

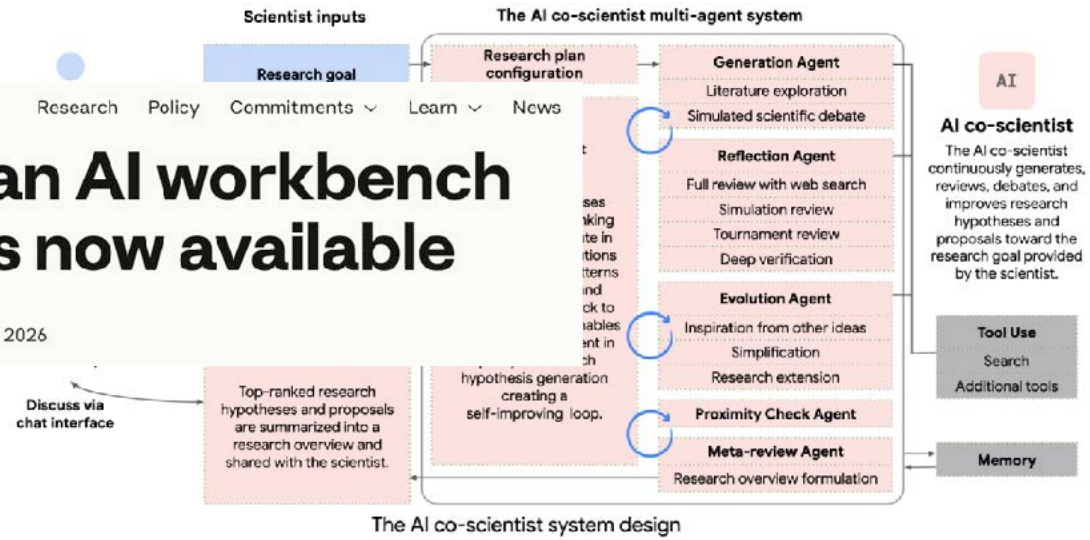
Popper

Automated Hypothesis Validation with **Karl Popper's principle of falsification.**

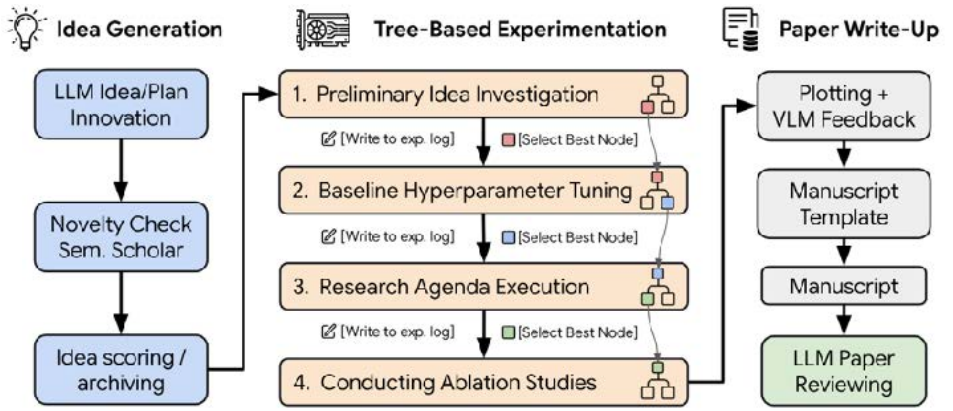
Research Policy Commitments Learn News

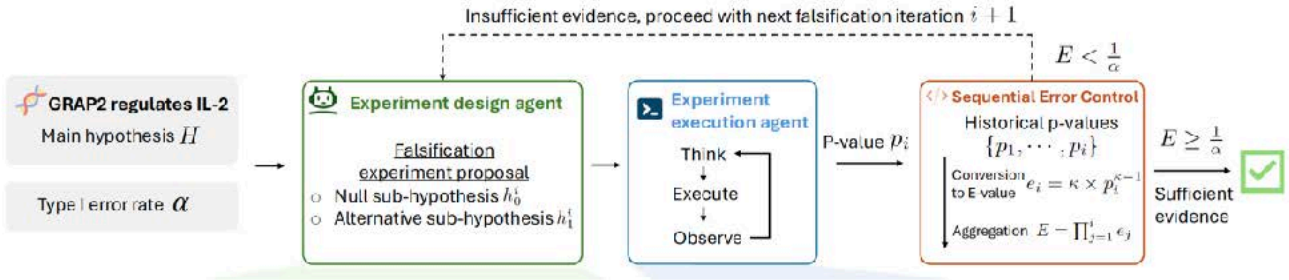
Claude Science, an AI workbench for scientists, is now available

Jun 30, 2026



The AI co-scientist system design





From generative AI to agentic AI in sciences

Denari

Multi-agent system that generates full pa
biology, chemistry, and more. **NeurIPS 2**

DeepM

Generates novel hypotheses, reviews literature, and designs
experiments in a closed-loc

sakar

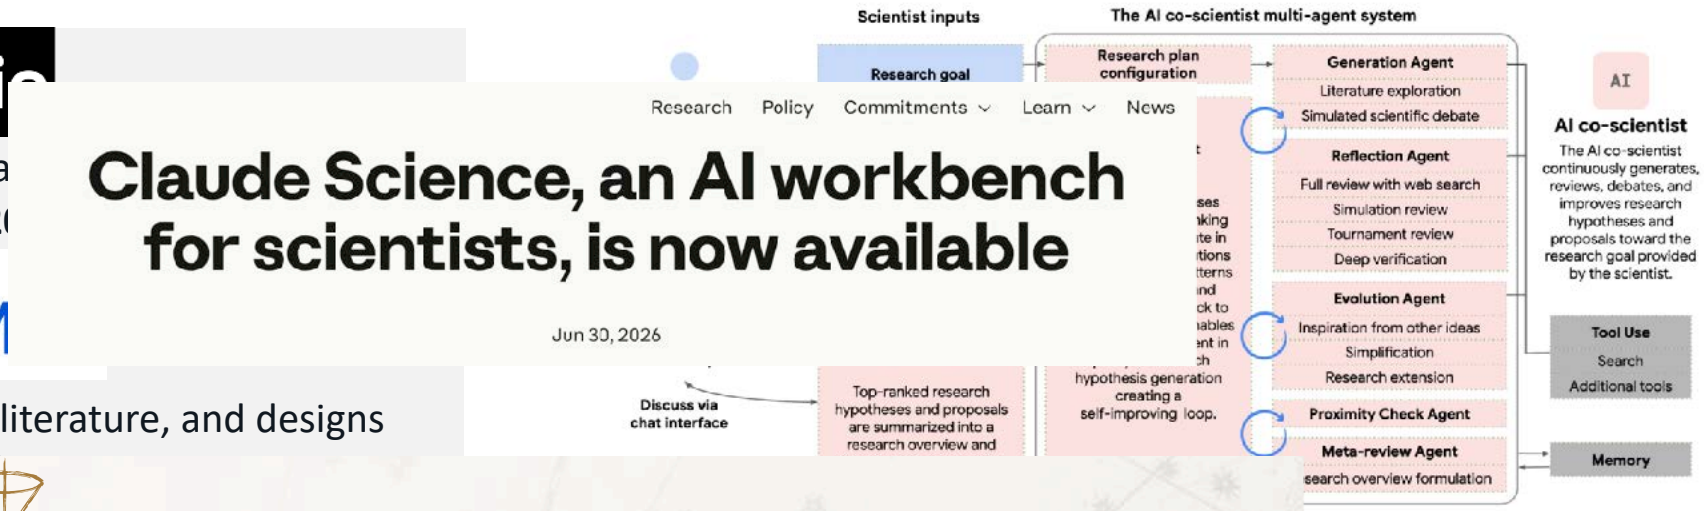
End-to-end system: **ideation → coc**
→ peer review, a

Edi

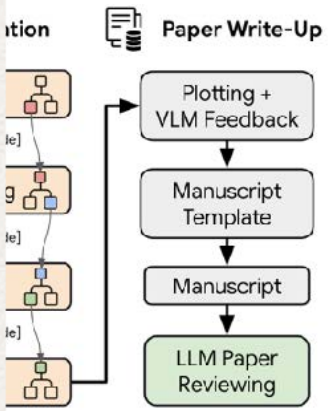
12-hour autonomous runs combini
search. Proposed a novel mechanis
public genet

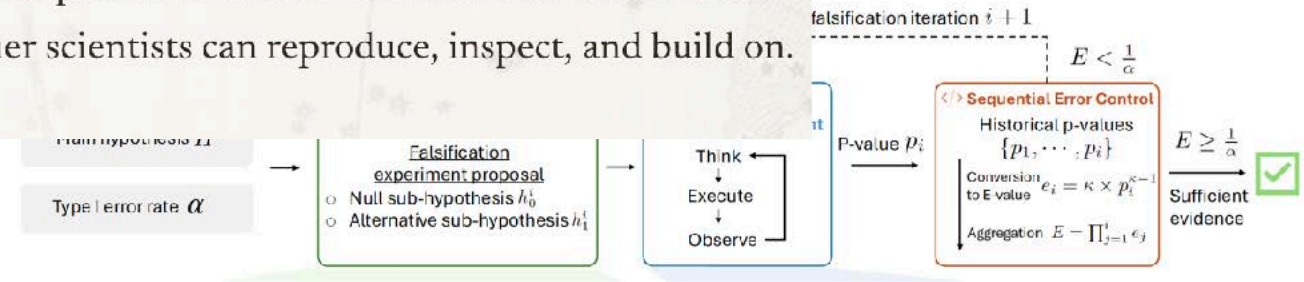
Pop

Automated Hypothesis Validation with **Karl Popper's principle of falsification.**







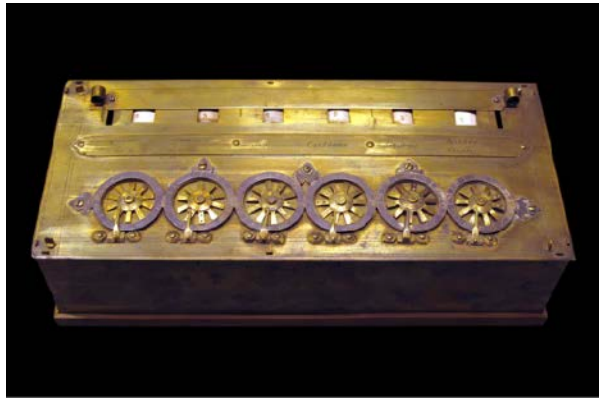


Mathematics and machines: a long history

Roman abacus, 1st century CE



Pascaline, 1652



Human computing, 20th century



Electronic computers, 20th century



Represent | **Calculate** | **Delegate** | **Automate**

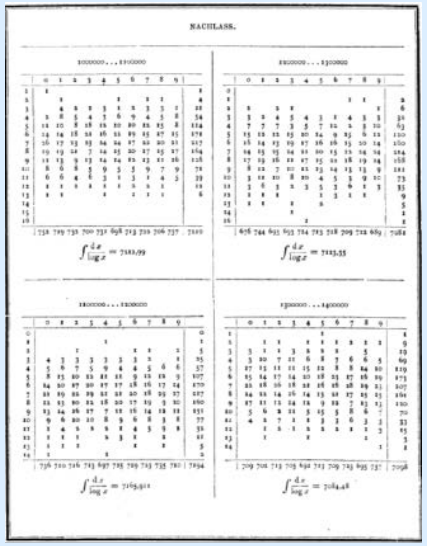
Mathematicians have outsourced computation for millennia.

From computation to discovery

Computing data can reveal patterns, suggest conjectures, and lead to new mathematics.

1. Tables of prime numbers

19th century - Gauss & Legendre





Carl Friedrich Gauss
(1777-1855)



Adrien-Marie Legendre
(1752-1833)

By examining large tables of prime numbers, Gauss and Legendre **observed regularities** and **conjectured** the prime number theorem.

2. A computer finds a counterexample

Euler's conjecture (1769)

If $a_1^k + a_2^k + \cdots + a_n^k = b^k$
for positive integers, then $n \geq k$.

1966 - CDC 6600
Lander & Parkin

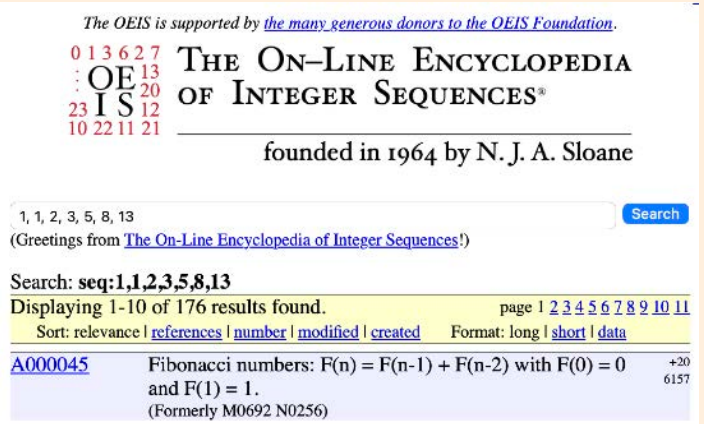


$$27^5 + 84^5 + 110^5 + 133^5 = 144^5$$

After almost 200 years, a computer found a **counterexample**. Euler's conjecture is false.

3. Mathematical databases

Today - e.g. OEIS



Large mathematical databases (like the OEIS) allow us to compare sequences, revealing **unexpected connections** between different areas of mathematics.

Compute → **Observe** → **Conjecture** → **Prove**

The machine is no longer only a calculator: it becomes a tool for mathematical discovery.

Mathematics and computers

2015 Math Panel with Donaldson, Kontsevich, Lurie, Tao, Taylor, Milner



INTERVIEWER: Do you imagine that, maybe in 100 years, or 1000 years, that, like it happened in chess -- humans still play tournaments but everyone knows computers are better -- is it conceivable that this could happen in mathematics?

TAO: I think computers will be able to do things much more efficiently with the right computer tools. …

KONTSEVITCH: I think it's perfectly possible, maybe in our lifetime.

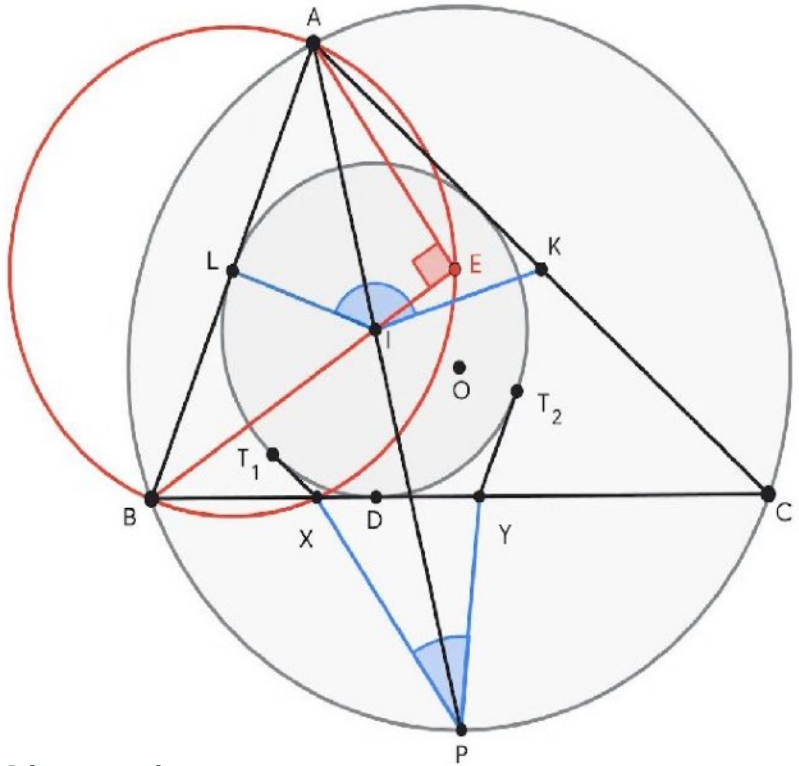
TAO: One could imagine that a computer could discover just by brute force a connection between two fields of mathematics that wasn't suspected, and then the person on the computer would be able to flesh it out. Maybe he would collect the medal.

Example of mathematics and CS

AlphaProof and
AlphaGeometry 2 :
Silver Medal at IMO 2024.

AI | MO
ARTIFICIAL INTELLIGENCE MATHEMATICAL OLYMPIAD

Gemini DeepThink:
Gold Medal 2025



Show that ∠KIL + ∠XPY = 180°



« LLM-guided evolutionary search can autonomously discover mathematical constructions that complement human intuition, at times matching or even improving the best known results, highlighting the potential for significant new ways of interaction between mathematicians and AI systems. »



« The fact that the program can come up with a non-obvious construction like this is very impressive, and well beyond what I thought was state of the art. »

Reasoning models: Chain of thoughts, Test time scaling, RL.

AlphaEvolve: LLM coding agent for algorithmic solutions to scientific problems.

AlphaProof: formal reasoning.
DeepThink: informal reasoning.

Interaction with formal proof languages

ROCQ

LEAN

Isabelle

What do we delegate to machines?

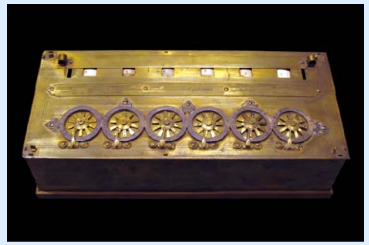
A long history, a moving frontier

1. CALCULATE

Represent and perform basic operations



Roman abacus
(1st century CE)



Pascaline
(1652)

From mental calculation to mechanical aids.

2. COMPUTE AT SCALE

Carry out large numbers of calculations



Human computers
(20th century)



Electronic computers
(from 1940s)

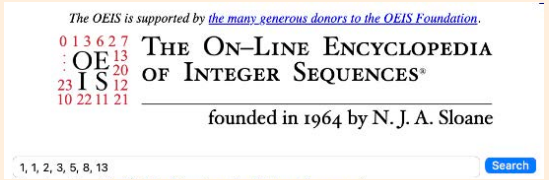
Speed and scale far beyond human capability.

3. SEARCH & EXPLORE

Explore large spaces, find examples, detect patterns

$$27^5 + 84^5 + 110^5 + 133^5 = 144^5$$

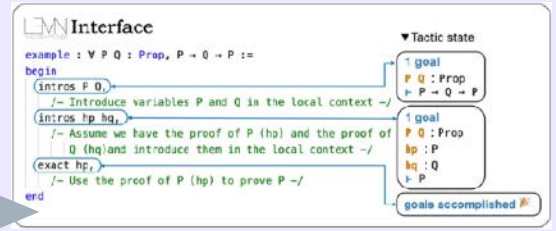
CDC 6600(1966)
Found a counterexample to Euler's conjecture.



From computation to conjectures and connections.

4. VERIFY

Check computations and proofs



- Computer-assisted proofs (e.g. four-colour theorem)
- Formal proof assistants (Lean, Coq, Isabelle, ...)
- Verification of complex calculations

The machine ensures rigor at scale.

More arithmetic → More mathematical reasoning

The frontier keeps moving.

Each new tool allows us to delegate more - and to ask new questions.

The changing frontier

What remains human? What can be delegated?

HUMAN

Insight, meaning, creativity



Choose the problem
Ask the right question
Choose representations
Interpret results
Recognize what matters
Form conjectures
Find strategies
Mathematical intuition?

THE FRONTIER IS MOVING

CALCULATE

SEARCH

VERIFY

SUGGEST

THINK?

MACHINE

Speed, scale, precision



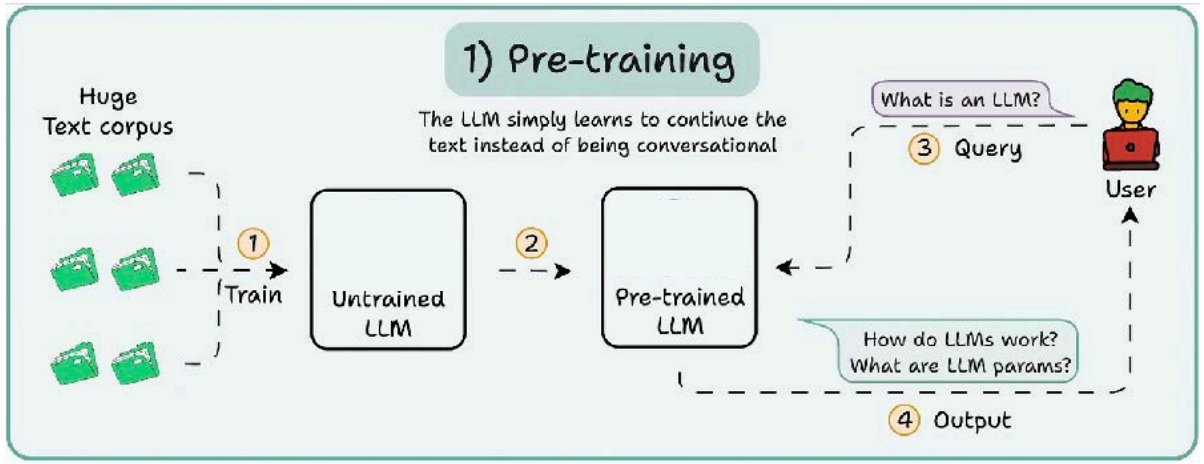
Compute
Search at scale
Compare
Detect patterns
Verify proofs
Formalize mathematics
Suggest conjectures?
Suggest strategies?

If machines can calculate, search and verify...

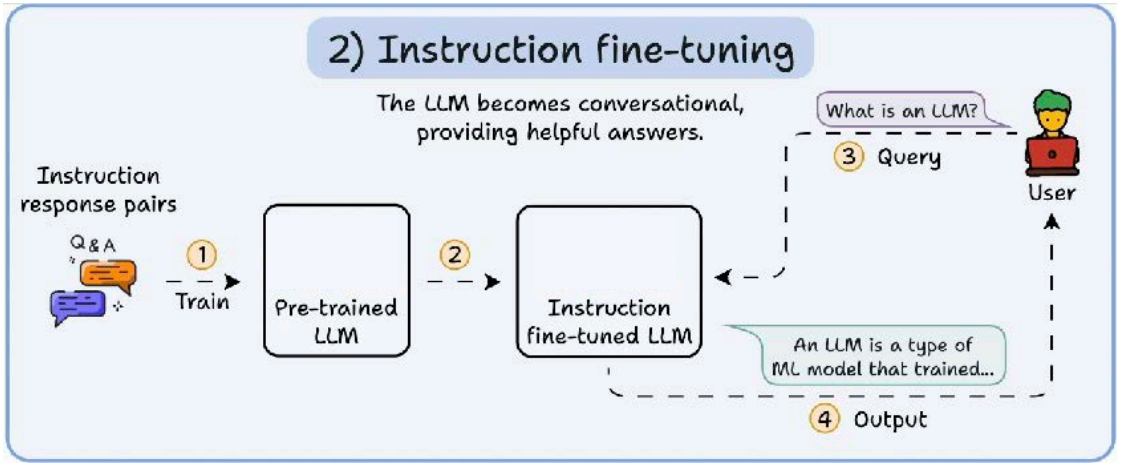
can they also help us think?

Modern LLMs: from text to (informal) reasoning

Self-supervised

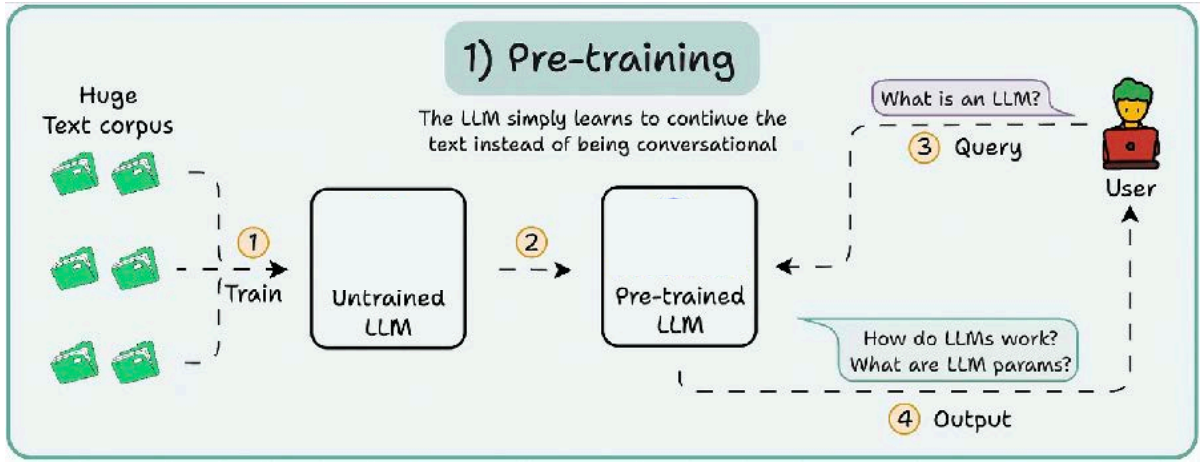


Supervised

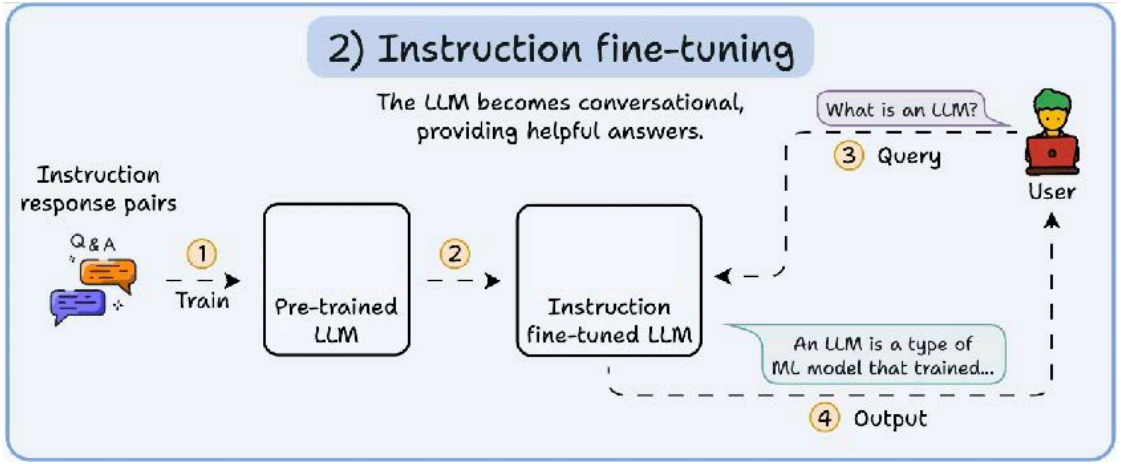


Modern LLMs: from text to (informal) reasoning

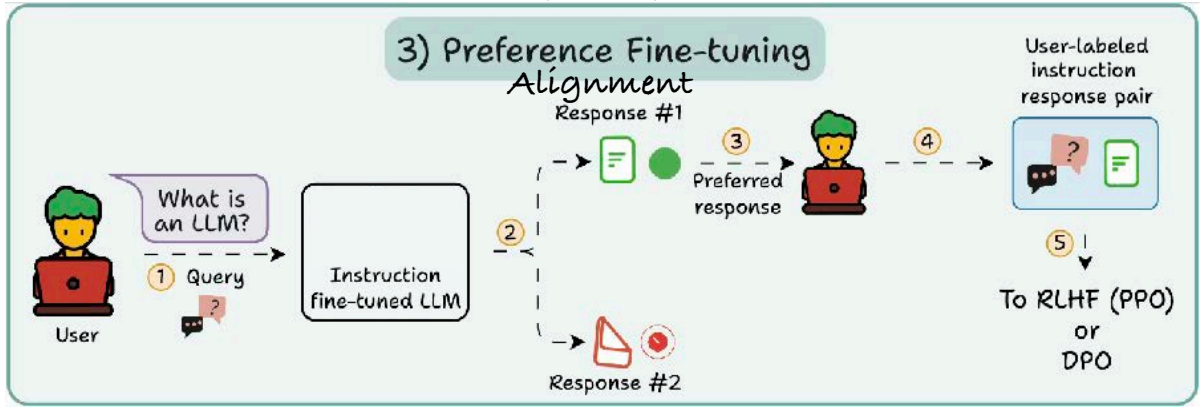
Self-supervised



Supervised

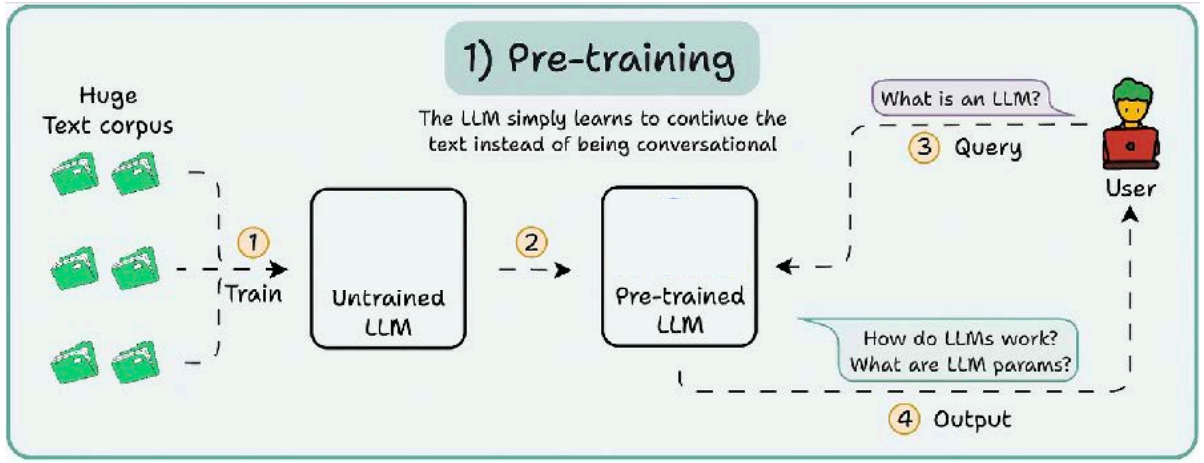


Reinforcement Learning from Human Feedback (RLHF)

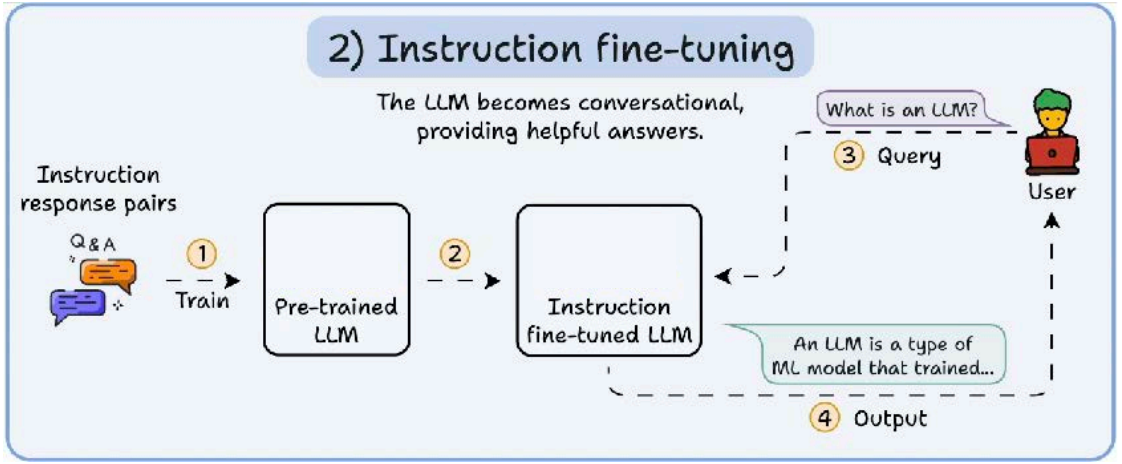


Modern LLMs: from text to (informal) reasoning

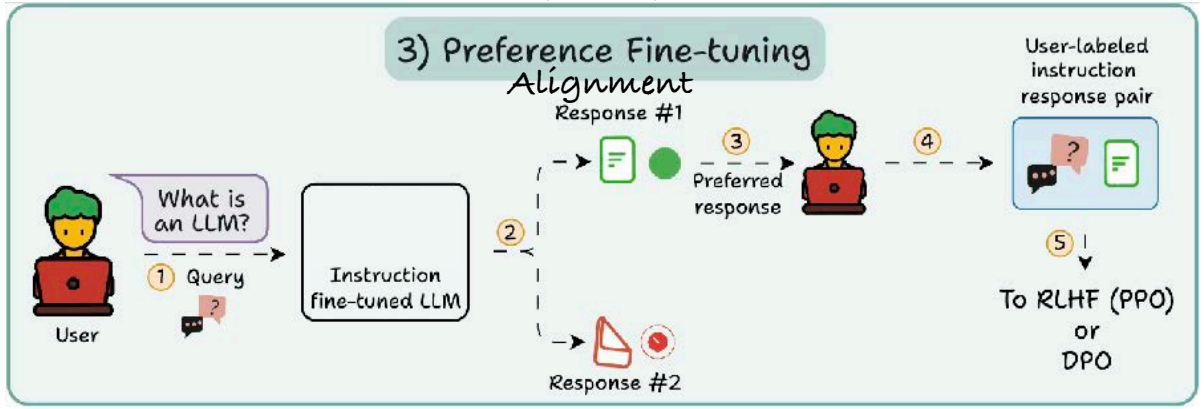
Self-supervised



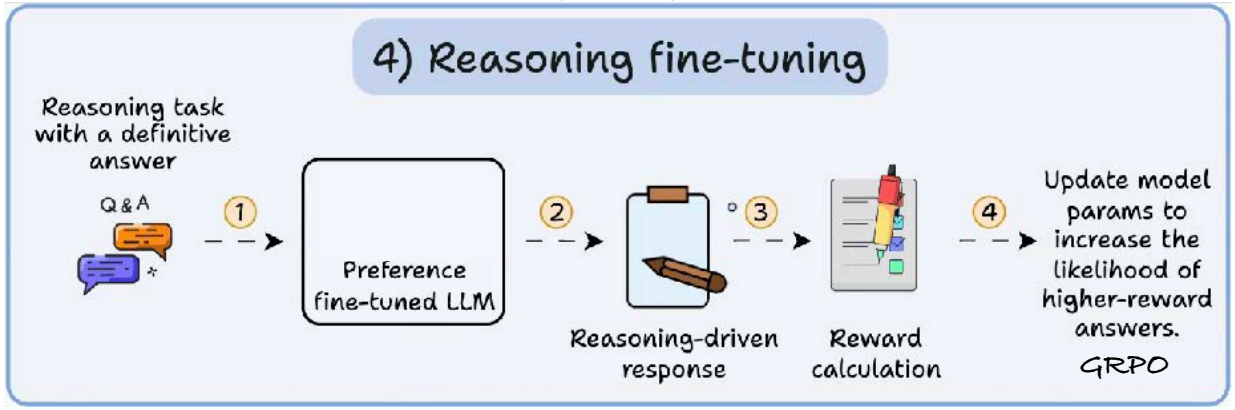
Supervised



Reinforcement Learning from Human Feedback (RLHF)



Reinforcement Learning with Verifiable Rewards (RLVR)



Formal proof assistants: 2nd leg of the AI4math revolution

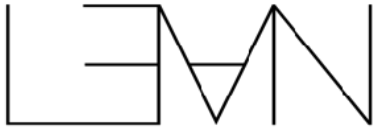
Autoformalization: capability to transform the informal style of human-produced proofs (and outputs of LLMs trained on them) into bulletproof, machine-verifiable logical derivations, expressed in specialized languages.

Informal Statement: Show that for any real number a, $10a \leq 28a^2 + 1$.

Informal Proof:
It suffices to show $0 \leq 28a^2 - 10a + 1$. First, consider completing the square for $28a^2 - 10a$ and observe that $\left(a - \frac{5}{28}\right)^2 = a^2 - \frac{10}{28}a + (5/28)^2$. Since $0 \leq \left(a - \frac{5}{28}\right)^2$, we get $0 \leq a^2 - \frac{10}{28}a + (5/28)^2$. Multiplying by 28 and simplifying gives $0 \leq 28a^2 - 10a + (25/28)$. Since $25/28 < 1$, the result follows.

Formal Proof Sketch:

```
theorem algebra_binomnegdiscrineq_10alt28asqp1:
  fixes a :: real
  shows "10 * a ≤ 28 * a^2 + 1"
proof - (* it suffices to show 0 <= 28a^2 - 10a + 1 *)
  have c0: "0 ≤ 28a^2 - 10a + 1"
  proof - (* observe that (a - (5/28))^2 = a^2 - (10/28)a + (5/28)^2 *)
    have c1: "(a - (5/28))^2 = a^2 - 10/28a + (5/28)^2" <···>
    (* we get 0 <= a^2 - (10/28)a + (5/28)^2 *)
    have c2: "0 ≤ a^2 - 10/28a + (5/28)^2" using c1 <···>
    (* Multiplying by 28 and simplifying gives 0 <= 28a^2 - 10a + (25/28) *)
    have c3: "0 ≤ 28a^2 - 10a + 28((5/28)^2)" using c2 <···>
    have c4: "0 ≤ 28a^2 - 10a + 28((5/28)*(5/28))" using c3 <···>
    have c5: "0 ≤ 28a^2 - 10a + (25/28)" using c4 <···>
    (* Since 25/28 < 1, the result follows. *)
    show ?thesis using c5 <···>
  qed
  show ?thesis <···>
qed
```







Formal proof assistants: 2nd leg of the AI4math revolution

Autoformalization: capability to transform the informal style of human-produced proofs (and outputs of LLMs trained on them) into bulletproof, machine-verifiable logical derivations, expressed in specialized languages.

Informal Statement: Show that for any real number a

Informal Proof:
It suffices to show $0 \leq 28a^2 - 10a + 1$. First, consi
observe that $\left(a - \frac{5}{28}\right)^2 = a^2 - \frac{10}{28}a + (5/28)^2$. Since 0
Multiplying by 28 and simplifying gives $0 \leq 28a^2 -$
follows.

Formal Proof Sketch:

```
theorem algebra_binomnegdiscrineq_10alt28asc
  fixes a :: real
  shows "10 * a ≤ 28 * a^2 + 1"
proof - (* it suffices to show 0 <= 28a^2 -
  have c0: "0 ≤ 28a^2 - 10a + 1"
  proof - (* observe that (a - (5/28))^2 = a
    have c1: "(a - (5/28))^2 = a^2 - 10/28a
    (* we get 0 <= a^2 - (10/28)a + (5/28)^2
    have c2: "0 ≤ a^2 - 10/28a + (5/28)^2"
    (* Multiplying by 28 and simplifying giv
    have c3: "0 ≤ 28a^2 - 10a + 28((5/28)^2
    have c4: "0 ≤ 28a^2 - 10a + 28((5/28)*(
    have c5: "0 ≤ 28a^2 - 10a + (25/28)" us
    (* Since 25/28 < 1, the result follows.
    show ?thesis using c5 < ··· >
  qed
  show ?thesis < ··· >
qed
```

Produits Solutions Recherche Blog Clients Entreprise Contact commercial Essayer Studio

Leanstral: Open-Source foundation for trustworthy vibe-coding

First open-source code agent for Lean 4.

Research

Mar 16, 2026 Mistral AI

AI agents have proven to be highly capable tools at code generation. Yet, as we push these models to high-stakes domains, ranging from frontier research mathematics to mission-critical software, we encounter a scaling bottleneck: the human review. The time and specialized expertise required to manually verify become the primary impedance of engineering velocity.

We envision a more helpful generation of coding agents to both carry out their tasks and formally prove their implementations against strict specifications. Instead of debugging machine-generated logic, humans dictate what they want. Today, we are taking the first major step toward that vision.

Introducing Leanstral

We release Leanstral, the first open-source code agent designed for Lean 4. Lean4 is a proof assistant capable of expressing complex mathematical objects such as perfectoid spaces and software specifications like properties of Rust fragments. Unlike existing proving systems that act as wrappers around large generalist models or focus on single math problems, Leanstral is designed to be highly efficient (with 6B active parameters) and trained for operating in realistic formal repositories.

Partager

Plus

Actualités

Modèles

Services IA

LEAN

ROCQ

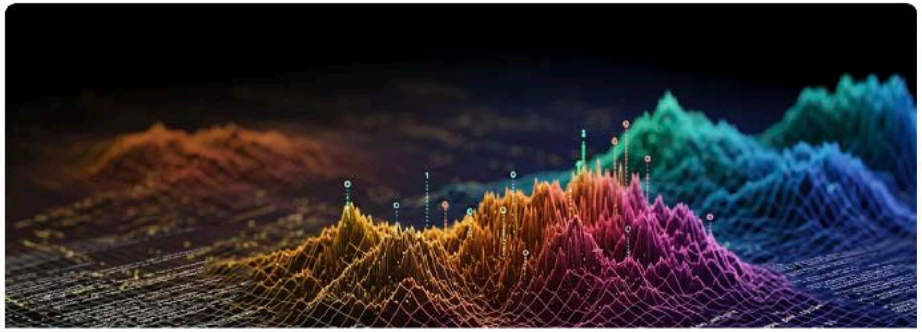
Isabelle

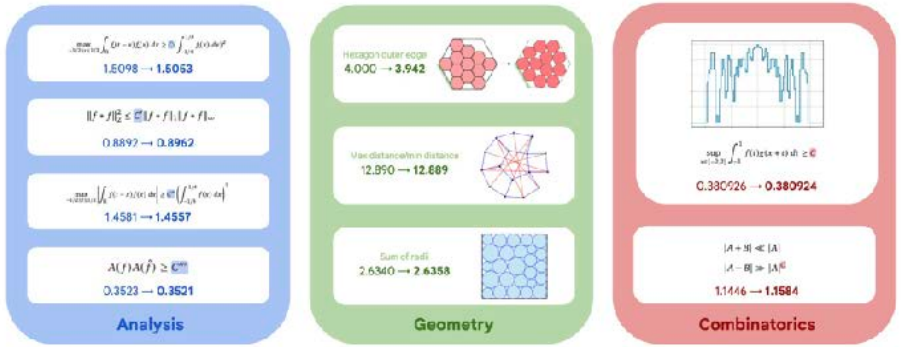
May 14, 2025 Science

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

AlphaEvolve team

Share





What's new

Updates on my research and expository papers, discussion of open problems, and other maths-related topics. By Terence Tao

Subscribe to feed | Home | About | Career advice | On writing | Books | Mastodon+ | Applets

RECENT COMMENTS

Terence Tao on Mathematical exploration and d...
Anonymous on Mathematical exploration and d...
Anonymous on Mathematical exploration and d...
QNFT on Mathematical exploration and d...
Terence Tao on Mathematical exploration and d...
Alan Chang on 247B, Notes 4: almost everywhe...
Anonymous on Mathematical exploration and d...
Vance Faber on Mathematical exploration and d...
Terence Tao on Mathematical exploration and d...
Jas, the Physicist on Mathematical exploration and d...
QNFT on Mathematical exploration and d...
inspiringcd947a018e on Mathematical exploration and d...
Terence Tao on Sumset and inverse sumset theo...
zawagner22 on Mathematical exploration and d...
Terence Tao on Mathematical exploration and d...

Mathematical exploration and discovery at scale

5 November, 2025 in math.CA, math.CO, math.MG, paper | Tags: Adam Zsolt Wagner, AlphaEvolve, Artificial Intelligence, Bogdan Georgiev, Javier Gomez-Serrano, optimization | by Terence Tao | 25 comments

Bogdan Georgiev, Javier Gómez-Serrano, Adam Zsolt Wagner, and I have uploaded to the arXiv our paper "Mathematical exploration and discovery at scale". This is a longer report on the experiments we did in collaboration with Google Deepmind with their AlphaEvolve tool, which is in the process of being made available for broader use. Some of our experiments were already reported on in a previous white paper, but the current paper provides more details, as well as a link to a repository with various relevant data such as the prompts us and the evolution of the tool outputs.

AlphaEvolve is a variant of more traditional optimization tools that are designe to extremize some given score function over a high-dimensional space of possible inputs. A traditional optimization algorithm might evolve one or more trial inputs over time by various methods, such as stochastic gradient descent, that are intended to locate increasingly good solutions while trying to avoid getting stuck at local extrema. By contrast, AlphaEvolve does not evolve the score function inputs directly, but uses an LLM to evolve computer code (often written in a standard language such as Python) which will in turn be run to generate the inputs that one tests the score function on. This reflects the belie that in many cases, the extremizing inputs will not simply be an arbitrary looki string of numbers, but will often have some structure that can be efficiently described, or at least approximated, by a relatively short piece of code. The to then works with a population of relatively successful such pieces of code, with the code from one generation of the population being modified and combined t the LLM based on their performance to produce the next generation. The stochastic nature of the LLM can actually work in one's favor in such an evolutionary environment: many "hallucinations" will simply end up being prun out of the pool of solutions being evolved due to poor performance, but a smal number of such mutations can add enough diversity to the pool that one can break out of local extrema and discover new classes of viable solutions. The LL can also accept user-supplied "hints" as part of the context of the prompt; in some cases, even just uploading PDFs of relevant literature has led to improve performance by the tool. Since the initial release of AlphaEvolve, similar tools have been developed by others, including OpenEvolve, ShinkaEvolve and DeepEvolve.

New result distribution

Visualization of results across 67 problems.



arXiv:2511.02864v1 [cs.NE] 3 Nov 2025



MATHEMATICAL EXPLORATION AND DISCOVERY AT SCALE

BOGDAN GEORGIEV, JAVIER GÓMEZ-SERRANO, TERENCE TAO, AND ADAM ZSOLT WAGNER

ABSTRACT. AlphaEvolve [223] is a generic evolutionary coding agent that combines the generative capabilities of LLMs with automated evaluation in an iterative evolutionary framework that proposes, tests, and refines algorithmic solutions to challenging scientific and practical problems. In this paper we showcase AlphaEvolve as a tool for autonomously discovering novel mathematical constructions and advancing our understanding of long standing open problems.

To demonstrate its breadth, we considered a list of 67 problems spanning mathematical analysis, combinatorics, geometry, and number theory. The system rediscovered the best known solutions in most of the cases and discovered improved solutions in several. In some instances, AlphaEvolve is also able to *generalize* results for a finite number of input values into a formula valid for all input values. Furthermore, we are able to combine this methodology with Deep Think [148] and AlphaProof [147] in a broader framework where the additional proof-assistants and reasoning systems provide automated proof generation and further mathematical insights.

These results demonstrate that large language model guided evolutionary search can autonomously discover mathematical constructions that complement human intuition, at times matching or even improving the best known results, highlighting the potential for significant new ways of interaction between mathematicians and AI systems. We present AlphaEvolve as a powerful new tool for mathematical discovery, capable of exploring vast search spaces to solve complex optimization problems at scale, often with significantly reduced requirements on preparation and computation time.

1. Introduction

The landscape of mathematical discovery has been fundamentally transformed by the emergence of computational tools that can autonomously explore mathematical spaces and generate novel constructions [56, 119, 241, 287]. AlphaEvolve represents a step in this evolution, demonstrating that large language models, when combined with evolutionary computation and rigorous automated evaluation, can discover explicit constructions that either match or improve upon the best-known bounds to long-standing mathematical problems, at large scales.

AlphaEvolve is not a general-purpose solver for all types of mathematical problems; it is primarily designed to attack problems in which a key objective is to construct a complex mathematical object that satisfies good quantitative properties, such as obeying a certain inequality with a good numerical constant. In this paper, we report on our experiments testing the performance of AlphaEvolve on a wide variety of such problems, primarily in the areas of analysis, combinatorics, and geometry. In many cases, the constructions provided by AlphaEvolve were not merely numerical in nature, but can be interpreted and generalized by human mathematicians, by other tools such as Deep Think, and even by AlphaEvolve itself. AlphaEvolve was not able to match or exceed previous results in all cases, and some of the individual improvements it was able to achieve could likely also have been matched by more traditional computational or theoretical methods performed by human experts. However, in contrast to such methods, we have found that AlphaEvolve can be readily scaled up to study large classes of problems at a time, without requiring extensive expert supervision for each new problem. This demonstrates that evolutionary computational approaches can systematically explore the space of mathematical objects in ways that complement traditional techniques, thus helping answer questions about the relationship between computational search and mathematical existence proofs.

We have also seen that in many cases, besides the scaling, in order to get AlphaEvolve to output comparable results to the literature and in contrast to traditional ways of doing mathematics, very little overhead is needed: on average the usual preparation time for the setup of a problem using AlphaEvolve took only up to a few hours.

1 [cs.AI] 5 Feb 2026

First Proof

Mohammed Abouzaid* Andrew J. Blumberg Martin Hairer
Stanford University *Columbia University* *EPFL and Imperial*
Joe Kileel Tamara G. Kolda Paul D. Nelson
University of Texas at Austin *MathSci.ai* *Aarhus University*
Daniel Spielman Nikhil Srivastava†
Yale University *University of California, Berkeley*
Rachel Ward† Shmuel Weinberger
University of Texas at Austin *University of Chicago*
Lauren Williams§
Harvard University

February 6, 2026

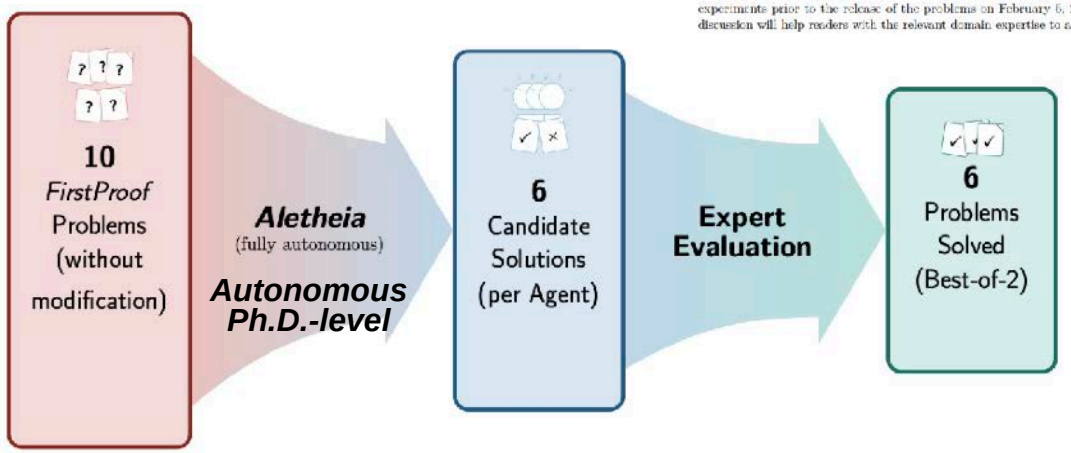
First Proof solutions and comments

Mohammed Abouzaid* Andrew J. Blumberg Martin Hairer
Stanford University *Columbia University* *EPFL and Imperial*
Joe Kileel Tamara G. Kolda Paul D. Nelson
University of Texas at Austin *MathSci.ai* *Aarhus University*
Daniel Spielman Nikhil Srivastava†
Yale University *University of California, Berkeley*
Rachel Ward† Shmuel Weinberger
University of Texas at Austin *University of Chicago*
Lauren Williams§
Harvard University

February 14, 2026

Abstract

Here we provide our solutions to the First Proof questions. We also discuss the best responses from publicly available AI systems that we were able to obtain in our experiments prior to the release of the problems on February 5, 2026. We hope this discussion will help readers with the relevant domain expertise to assess such responses.



FirstProof # 8

A polyhedral Lagrangian surface K in $\mathbb{R}^4$ is a finite polyhedral complex all of whose faces are Lagrangians, and which is a topological submanifold of $\mathbb{R}^4$. A Lagrangian smoothing of K is a Hamiltonian isotopy K_t of smooth Lagrangian submanifolds, parameterised by $(0, 1]$, extending to a topological isotopy, parametrised by $[0, 1]$, with endpoint $K_0 = K$.
Let K be a polyhedral Lagrangian surface with the property that exactly 4 faces meet at every vertex. Does K necessarily have a Lagrangian smoothing?

Model Response (Agent A)

Let $K \subset \mathbb{R}^4$ be a polyhedral Lagrangian surface. By definition, K is a finite polyhedral complex all of whose faces are Lagrangian, and which forms a topological submanifold of $\mathbb{R}^4$. A Lagrangian smoothing of K is a Hamiltonian isotopy K_t of smooth Lagrangian submanifolds, parameterised by $t \in (0, 1]$, that extends to a continuous topological isotopy for $t \in [0, 1]$ such that $K_0 = K$.
The local combinatorial structure of K strongly restricts its symplectic geometry. In this paper, we investigate the specific setting where exactly 4 faces meet at every vertex.

Theorem 1. *Every 4-valent polyhedral Lagrangian surface $K \subset \mathbb{R}^4$ admits a Lagrangian smoothing.*

1 Local Symplectic Factorization at the Vertices

Let $v \in K$ be an arbitrary vertex. Without loss of generality, we translate v to the origin of $\mathbb{R}^4$. The local tangent space of K at v consists of exactly 4 flat planar sectors bounded by 4 cyclically ordered outgoing rays e_1, e_2, e_3, e_4. The 4 incident faces are the strictly positive planar cones $F_i = \mathrm{span}_{\geq 0}(e_i, e_{i+1})$ for $i = 1, \ldots, 4$ (with indices taken modulo 4). Let $P_i = \mathrm{span}(e_i, e_{i+1})$ denote the full linear span of each face.
Because K is a Lagrangian surface, the standard symplectic form ω inherently evaluates to zero on each face, imposing strict cyclic orthogonality on the generating rays:

$$\omega(e_1, e_2) = \omega(e_2, e_3) = \omega(e_3, e_4) = \omega(e_4, e_1) = 0. \tag{1}$$

Before analyzing the general configuration, we must account for the possibility that the linear spans of adjacent faces coincide. If, for instance, $P_1 = P_2$, the faces F_1 and F_2 are coplanar. In this setting, their intersection is a 2-dimensional plane rather than a 1-dimensional line, and the faces can form adjacent, non-overlapping sectors. If this occurs, the local geometry of K at v is either entirely flat or reducible to a

1 [cs.AI] 5 Feb 2026

First Proof

Mohammed Abouzaid* *Stanford University*
Andrew J. Blumberg *Columbia University*
Martin Hairer *EPFL and Imperial*
Joe Kileel *University of Texas at Austin*
Tamara G. Kolda *MathSci.ai*
Paul D. Nelson *Aarhus University*
Daniel Spielman *Yale University*
Nikhil Srivastava† *University of California, Berkeley*
Rachel Ward† *University of Texas at Austin*
Shmuel Weinberger *University of Chicago*
Lauren Williams§ *Harvard University*

February 6, 2026

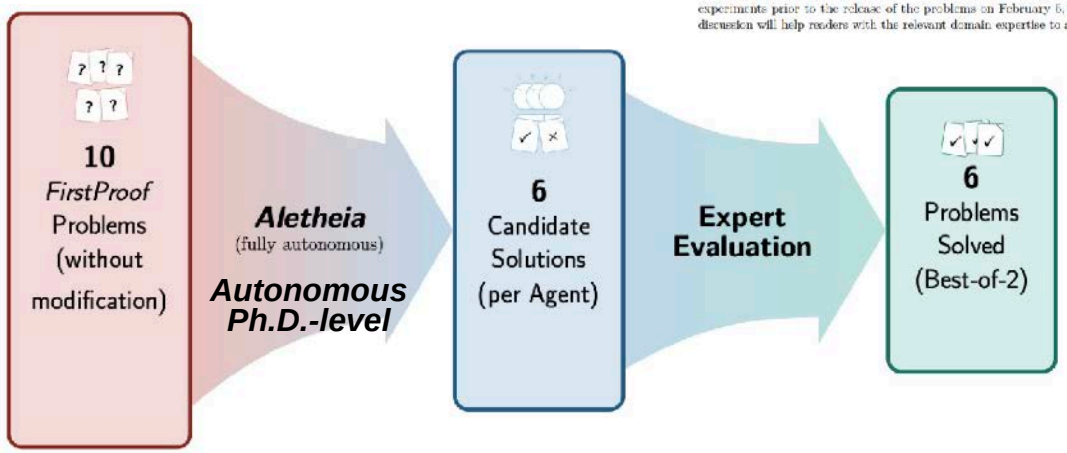
First Proof solutions and comments

Mohammed Abouzaid* *Stanford University*
Andrew J. Blumberg *Columbia University*
Martin Hairer *EPFL and Imperial*
Joe Kileel *University of Texas at Austin*
Tamara G. Kolda *MathSci.ai*
Paul D. Nelson *Aarhus University*
Daniel Spielman *Yale University*
Nikhil Srivastava† *University of California, Berkeley*
Rachel Ward† *University of Texas at Austin*
Shmuel Weinberger *University of Chicago*
Lauren Williams§ *Harvard University*

February 14, 2026

Abstract

Here we provide our solutions to the First Proof questions. We also discuss the best responses from publicly available AI systems that we were able to obtain in our experiments prior to the release of the problems on February 5, 2026. We hope this discussion will help readers with the relevant domain expertise to assess such responses.



First Proof Third Batch

October 14, 2026

Mohammed Abouzaid[1] *Stanford University*
Nikhil Srivastava[2] *University of California, Berkeley*
Rachel Ward[3] *University of Texas at Austin*
Lauren Williams[4] *Harvard University*

Problem Contributors

1. Leo Herr, *Virginia Tech*[5]; Jonathan Wise, *University of Colorado Boulder*[6]
2. Markus Reiß, *Humboldt-Universität zu Berlin*[7]; Sven Wang, *EPFL*[8]
3. Adam S. Sikora, *University at Buffalo*[9]
4. Bernd Siebert, *The University of Texas at Austin*[10]
5. Tuong Le, *Princeton University*[11]; Chayim Lowen, *Princeton University*[12]
6. Cameron Gordon, *The University of Texas at Austin*[13]
7. Tanya Firsova, *Kansas State University*[14]; Mikhail Mazin, *Kansas State University*[15]
8. Simion Filip, *University of Chicago*[16]
9. Benjamin Dozier, *Cornell University*[17]; Sidhanth Raman, *University of Chicago*[18]; Carlos Serván, *University of Michigan*[19]
10. Charles Doran, *University of Alberta*[20]; James M. Jones, *Philipps-Universität Marburg*[21]; Eric Pichon-Pharabod, *University of Oxford*[22]; Alan Thompson, *Loughborough University*[23].

hexagon:2610.00244v1

- *5 AI systems on 10 research-level problems in various fields.*
- *The ProofCouncil harness and GPT-6 Astra resp. reaching 9/10 and 7/10 compared to Batch 2.*
- *Still citation and exposition issues.*

FirstProof # 8

A polyhedral Lagrangian surface K in $\mathbb{R}^4$ is a finite polyhedral complex all of whose faces are Lagrangians, and which is a topological submanifold of $\mathbb{R}^4$. A Lagrangian smoothing of K is a Hamiltonian isotopy K_t of smooth Lagrangian submanifolds, parameterised by $(0, 1]$, extending to a topological isotopy, parametrised by $[0, 1]$, with endpoint $K_0 = K$.
Let K be a polyhedral Lagrangian surface with the property that exactly 4 faces meet at every vertex. Does K necessarily have a Lagrangian smoothing?

Model Response (Agent A)

Let $K \subset \mathbb{R}^4$ be a polyhedral Lagrangian surface. By definition, K is a finite polyhedral complex all of whose faces are Lagrangian, and which forms a topological submanifold of $\mathbb{R}^4$. A Lagrangian smoothing of K is a Hamiltonian isotopy K_t of smooth Lagrangian submanifolds, parameterised by $t \in (0, 1]$, that extends to a continuous topological isotopy for $t \in [0, 1]$ such that $K_0 = K$.
The local combinatorial structure of K strongly restricts its symplectic geometry. In this paper, we investigate the specific setting where exactly 4 faces meet at every vertex.

Theorem 1. *Every 4-valent polyhedral Lagrangian surface $K \subset \mathbb{R}^4$ admits a Lagrangian smoothing.*

1 Local Symplectic Factorization at the Vertices

Let $v \in K$ be an arbitrary vertex. Without loss of generality, we translate v to the origin of $\mathbb{R}^4$. The local tangent space of K at v consists of exactly 4 flat planar sectors bounded by 4 cyclically ordered outgoing rays e_1, e_2, e_3, e_4. The 4 incident faces are the strictly positive planar cones $F_i = \mathrm{span}_{\geq 0}(e_i, e_{i+1})$ for $i = 1, \ldots, 4$ (with indices taken modulo 4). Let $P_i = \mathrm{span}(e_i, e_{i+1})$ denote the full linear span of each face.
Because K is a Lagrangian surface, the standard symplectic form ω inherently evaluates to zero on each face, imposing strict cyclic orthogonality on the generating rays:

$$\omega(e_1, e_2) = \omega(e_2, e_3) = \omega(e_3, e_4) = \omega(e_4, e_1) = 0. \quad (1)$$

Before analyzing the general configuration, we must account for the possibility that the linear spans of adjacent faces coincide. If, for instance, $P_1 = P_2$, the faces F_1 and F_2 are coplanar. In this setting, their intersection is a 2-dimensional plane rather than a 1-dimensional line, and the faces can form adjacent, non-overlapping sectors. If this occurs, the local geometry of K at v is either entirely flat or reducible to a

Yoshua Bengio
3 h ·

If this is not fake news, frontier AI is beginning to prove long-standing conjectures, from no less than Donald Knuth, a god for computer scientists.

Claude's Cycles

Don Knuth, Stanford Computer Science Department
(28 February 2026; revised 02 March 2026)

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 $0 \le i, j, k < m$, 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 $m > 2$.

I had solved the problem for $m = 3$, and asked for a generalization as part of the answer to an exercise in [3]. My friend Filip Stappers rose to the challenge, and empirically discovered solutions for $4 \le m \le 16$; therefore it became highly likely that the desired decompositions do exist, except when $m \le 2$.

Indeed, it was Filip who had the gumption to pose this question to Claude, using exactly the wording above. He also gave guidance/coaching, instructing Claude to summarize its ongoing progress:

> ** After EVERY `exploreXX.py` run, IMMEDIATELY update this file [`plan.md`] before doing anything else. ** No exceptions. Do not start the next exploration until the previous one is documented here.

February 13, 2026 Research Publication

GPT-5.2 derives a new result in theoretical physics

In a new preprint, GPT-5.2 proposed a formula for a gluon amplitude later proved by an internal OpenAI model and verified by the authors.

Read the preprint ↗

Listen to article 3:54 Share

We've published a new preprint showing that a type of particle interaction many physicists expected would not occur can in fact arise under specific conditions. The work focuses on gluons, the particles that carry the strong nuclear force. The preprint is available on arXiv and is being submitted for publication. In the meantime, we welcome feedback from the community.

October 6, 2026 Research Publication

Sharing AI progress in mathematics

View on GitHub ↗

The public openai/math repository currently contains 722 manuscripts organized into 372 related result families. It includes papers, supporting proof artifacts, Lean formalizations for many results, and 10 abridged summaries of the model's reasoning.

vitrupo @vitrupo · Jan 5

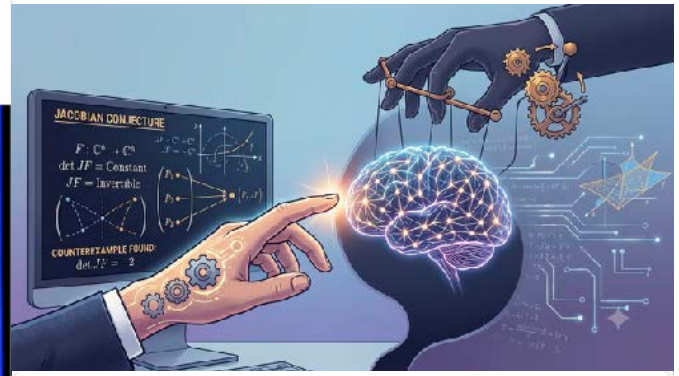
Geoffrey Hinton says mathematics is a closed system, so AIs can play it like a game.

They can pose problems to themselves, test proofs, and learn from what works, without relying on human examples....



2:29 PM · Jan 6, 2026 · 139.2K Views

> I agree with Demis Hassabis, the leader of DeepMind, who for many years has said AI is going to be very important for making scientific progress...
>
> There's one area in which that's particularly easy, which is mathematics, because mathematics is a closed system...
>
> I think AI will get much better at mathematics than people, maybe in the next 10 years or so. And within mathematics, it's much like things like Go or Chess that are closed systems with rules...



levent
@__alpoge__ · Follow

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 + 3 x (1+xy)^2 z + 3 x y^2 (4+3xy), 2 x - 3 x^2 y - x^3 z): \C^3\to \C^3, Show more

4:19 AM · Jul 20, 2026

September 8, 2026 Research Publication

On the Navier–Stokes Millennium Prize Problem

Read the paper Link to Lean formalized proof ›

May 20, 2026 Research Milestone

An OpenAI model has disproved a central conjecture in discrete geometry

Read the proof ↗ Read the companion remarks ↗

REMARKS ON THE DISPROOF OF THE UNIT DISTANCE CONJECTURE

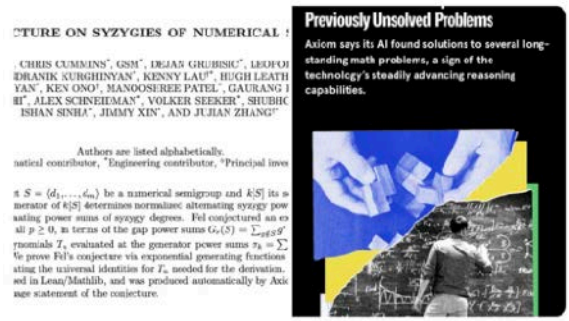
NOGA ALON, THOMAS F. BLOOM, W. T. GOWERS, DANIEL LITT, WILL SAWIN, ARUL SHANKAR, JACOB TSIMERMAN, VICTOR WANG, AND MELANIE MATCHETT WOOD

ABSTRACT. We present a short, digested, human-verified version of the recent OpenAI-generated counterexample to the Erdős unit distance conjecture, and a sequence of reflections on it. The argument relies crucially on ideas that may, at least in retrospect, be attributed to Ellenberg-Venkatesh, Golod-Shafarevich, and Hajir-Maire-Ramakrishna.

Axiom
@axiommathai
Suivre

1/ AxiomProver has solved Fel's open conjecture on syzygies of numerical semigroups, autonomously generating a formal proof in Lean with zero human guidance.

This is the first time an AI system has settled an unsolved research problem in theory-building math and self verifies.



17:34 · 05/02/2026 · 212k vues

20 203 819 316

Les réponses les plus pertinentes

Justin Curry
@currying

Aug 23, 2026

Kudos to Levent and @AnthropicAI for building a god-like machine that appears to have settled the Hopf problem: the six-dimensional sphere does admit a complex structure. Open since 1947, this problem has spurned many faulty advances, including Atiyah's embarrassing 2016 supposed disproof of any such structure existing. To have a computer explicitly construct one 10 years later and after his death, is perhaps an illustration of Planck's Principle for the AI age.

levent @__alpoge__ 9:02 AM · Aug 19, 2026 · 99.7K Views

a week ago i chatted with @AchimWar about favourite problems and the caratheodory conjecture came up. i literally wanted to just read about it but i did see the complicated literature situation and so in hopes i could help asked him if it was ok to look at with a model. me, John-Paul Smith, and of course claude checked and (i hope im being a better citizen, hademard led to a lot of flaming) i sent over a pdf writeup, and for the past week it's circulated round. i did wanna have it recorded somewhere public though so i'll try to again give a form certifiable by models

Yoshua Bengio
3 h ·

If this is not fake news, frontier AI is beginning to prove long-standing conjectures, from no less than Donald Knuth, a god for computer scientists.

Claude's Cycles

Don Knuth, Stanford Computer Science Department
(28 February 2026; revised 02 March 2026)

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 $0 \le i, j, k < m$, 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 $m > 2$.

I had solved the problem for $m = 3$, and asked for a generalization as part of the answer to an exercise in [3]. My friend Filip Stappers rose to the challenge, and empirically discovered solutions for $4 \le m \le 16$; therefore it became highly likely that the desired decompositions do exist, except when $m \le 2$.

Indeed, it was Filip who had the gumption to pose this question to Claude, using exactly the wording above. He also gave guidance/coaching, instructing Claude to summarize its ongoing progress:

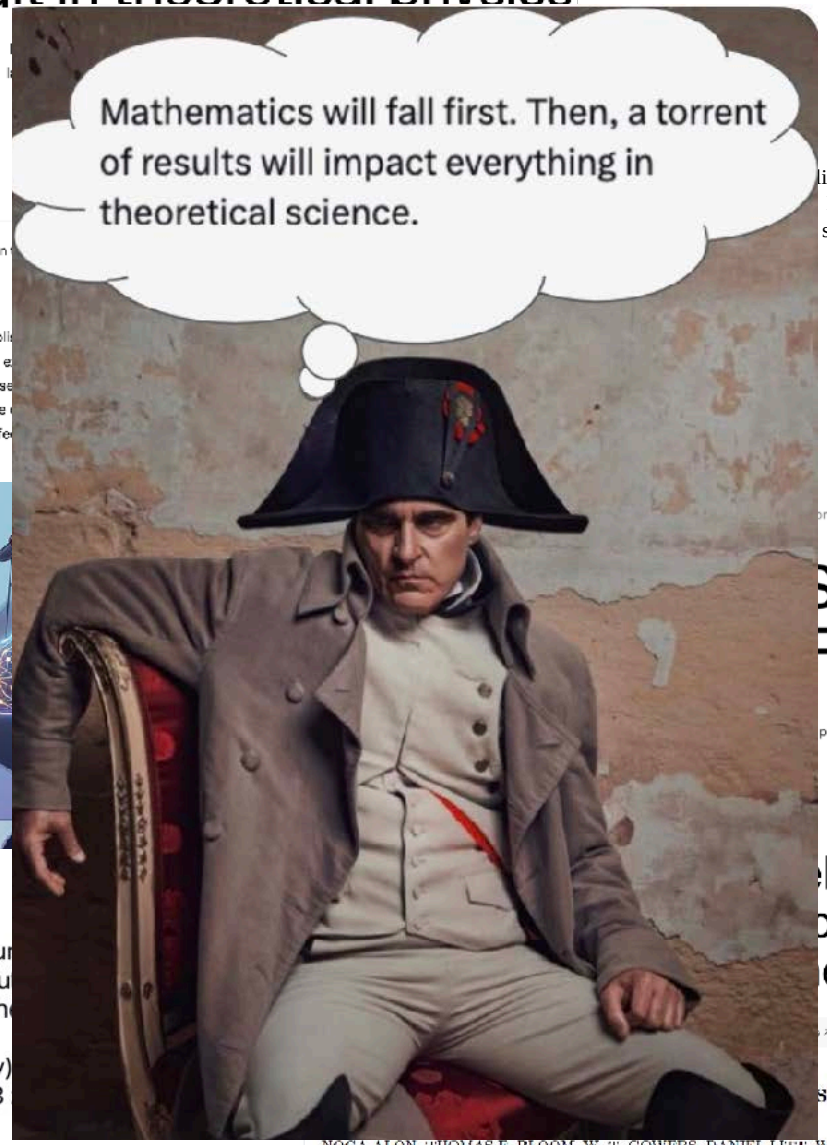
> ** After EVERY `exploreXX.py` run, IMMEDIATELY update this file [`plan.md`] before doing anything else. ** No exceptions. Do not start the next exploration until the previous one is documented here.

February 13, 2026 Research Publication

GPT-5.2 derives a new result in theoretical physics

Listen

We've publi
physicists e
work focuse
is available
welcome fe



October 6, 2026 Research Publication

Sharing AI progress in mathematics

View on GitHub ↗

lic openai/math repository currently contains 722 manuscripts organized into 372 related result
It includes papers, supporting proof artifacts, Lean formalizations for many results, and 10
summaries of the model's reasoning.

vitrupo @vitrupo · Jan 5

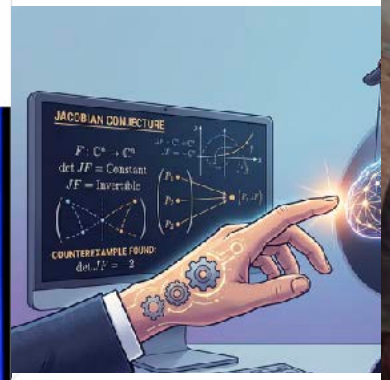
Geoffrey Hinton says mathematics is a closed system, so AIs can play it like a game.

They can pose problems to themselves, test proofs, and learn from what works, without relying on human examples....



2:29 PM · Jan 6, 2026 · 139.2K Views

> I agree with Demis Hassabis, the leader of DeepMind, who for many years has said AI is going to be very important for making scientific progress...
>
> There's one area in which that's particularly easy, which is mathematics, because mathematics is a closed system...
>
> I think AI will get much better at mathematics than people, maybe in the next 10 years or so. And within mathematics, it's much like things like Go or Chess that are closed systems with rules...



levent
@__alpoge__ · Follow

hello there the jacobian conjectur
close friend akhil for asking abou
friend fable for working during th

((1+xy)^3 z + y^2 (1+xy) (4+3xy)
y^2 (4+3xy), 2 x - 3 x^2 y - x^3
more

4:19 AM · Jul 20, 2026

Stokes
Problem
proof ›

l has
onjecture
etry

STANCE CONJECTURE

NOGA ALON, THOMAS F. BLOOM, W. T. GOWERS, DANIEL LITT, WILL SAWIN, ARUL SHANKAR, JACOB TSIMERMAN, VICTOR WANG, AND MELANIE MATCHETT WOOD

Abstract. We present a short, digested, human-verified version of the recent OpenAI-generated counterexample to the Erdős unit distance conjecture, and a sequence of reflections on it. The argument relies crucially on ideas that may, at least in retrospect, be attributed to Ellenberg-Venkatesh, Golod-Shafarevich, and Hajir-Maire-Ramakrishna.

Axiom
@axiommathai
Suivre

1/ AxiomProver has solved Fel's open conjecture on syzygies of numerical semigroups, autonomously generating a formal proof in Lean with zero human guidance.

This is the first time an AI system has settled an unsolved research problem in theory-building math and self verifies.

CTURE ON SYZYGIES OF NUMERICAL

CHRIS CUMMINS*, GSM*, DEJAN GRUBISIC*, LEOFOI
IDRANIK KURGHINYAN*, KENNY LAU*, HUGH LEATH
YAN*, KEN ONO*, MANOOSEREE PATEL*, GAURANG
HI*, ALEX SCHNEIDMAN*, VOLKER SEEKER*, SHUBHC
ISHAN SINHA*, JIMMY XIN*, AND JUJIAN ZHANG*

Authors are listed alphabetically.
natical contributor, *Engineering contributor, *Principal inves

Previously Unsolved Problems
Axiom says its AI found solutions to several long-standing math problems, a sign of the technology's steadily advancing reasoning capabilities.

17:34 · 05/02/2026 · 212k vues

20 203 819 316

Les réponses les plus pertinentes

Justin Curry
@currying
Aug 23, 2026

Kudos to Levent and @AnthropicAI for building a god-like machine that appears to have settled the Hopf problem: the six-dimensional sphere does admit a complex structure. Open since 1947, this problem has spurned many faulty advances, including Atiyah's embarrassing 2016 supposed disproof of any such structure existing. To have a computer explicitly construct one 10 years later and after his death, is perhaps an illustration of Planck's Principle for the AI age.

levent @__alpoge__ 9:02 AM · Aug 19, 2026 · 99.7K Views

a week ago i chatted with @AchimWar about favourite problems and the caratheodory conjecture came up. i literally wanted to just read about it but i did see the complicated literature situation and so in hopes i could help asked him if it was ok to look at with a model. me, John-Paul Smith, and of course claude checked and (i hope im being a better citizen, hademard led to a lot of flaming) i sent over a pdf writeup, and for the past week it's circulated round. i did wanna have it recorded somewhere public though so i'll try to again give a form certifiable by models

Jeremy Avigad

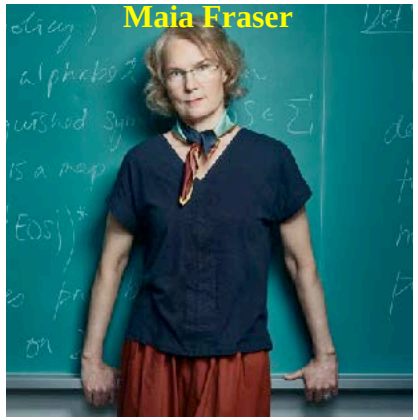


“Understanding is just a wonderful feeling when you’ve been thinking long and hard about something complex, difficult, and then—all of a sudden—it just comes together.”

We need to remember our strengths: **mathematicians are problem solvers and theory builders *extraordinaire***. Rather than fight the use of AI in mathematics, **we should *own* it**. It is not enough to keep up with current events and design benchmarks for AI researchers; **we need to play an active role in deploying the technology and molding it to our purposes**.

arXiv:2603.03684v3 [math.HO] 6 Apr 2026

Maia Fraser



“Mathematics (and science) is more than finding answers. The struggle to understand a problem is one of the greatest rewards”

Mathematicians in the Age of AI*

Jeremy Avigad

March 3, 2026

Abstract

Recent developments show that AI can prove research-level theorems in mathematics, both formally and informally. This essay urges mathematicians to stay up-to-date with the technology, to consider the ways it will disrupt mathematical practice, and to respond appropriately to the challenges and opportunities we now face.

1 Background

This past year was not an easy one for the U.S. National Science Foundation, but last fall it managed to launch the Institute for Computer-Aided Reasoning in Mathematics (ICARM), even while facing serious budgetary constraints. The institute’s mission is to support the adoption of new reasoning technologies in mathematics, many of them involving neural and symbolic AI. As director, I have given several variants of a talk titled “Mathematics in the Age of AI,” in which I survey the new technologies, including formalization and proof assistants, symbolic AI and automated reasoning, and machine learning and neural networks. These technologies are distinct, but they overlap and interact in interesting ways.

To date, early adopters of formal methods have contributed to a formal library for mathematics, Mathlib, and several collaborative projects have used it to verify contemporary mathematical results, even before or in the absence of the usual peer review process. Automated reasoning tools such as SAT solvers have solved open problems in combinatorics, algebra, number theory, and discrete geometry. Machine learning techniques have been used to detect patterns in and relationships between computational data, leading to the discovery of new results verified by conventional proofs. They have also been used to find combinatorial objects of interest, including counterexamples to conjectures. Neural networks are also used to compute solutions to PDEs, identify phenomena of interest, and search for parameter settings that give rise to such phenomena. Although the technologies are still niche and have seen only a handful of notable successes, the advances make it clear that they are bound to have a significant impact on the subject.

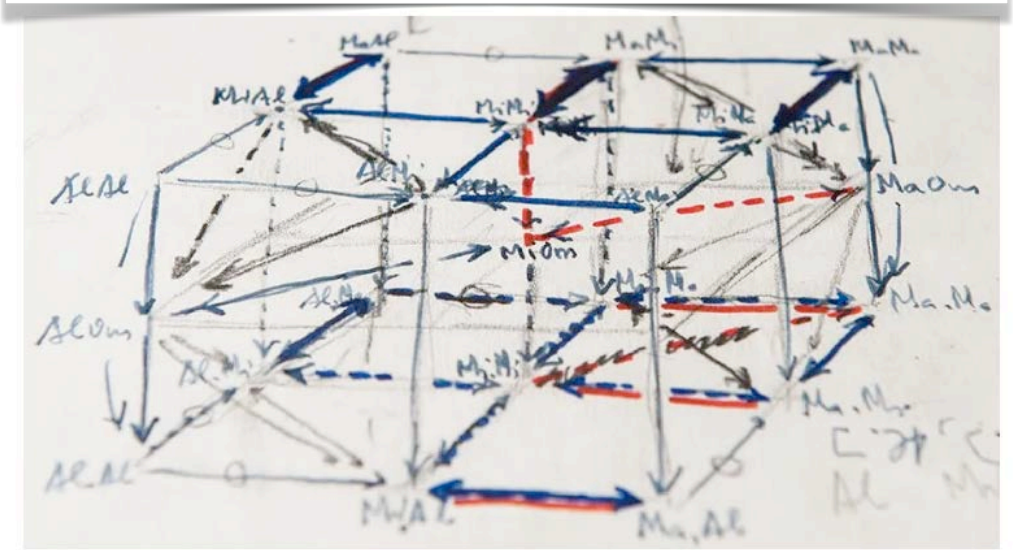
The fall of the theorem economy

How AI could destroy mathematics and barely touch it

DAVID BESSIS
APR 21, 2026

653 101 129 Share

“The product of mathematics is clarity and understanding. Not theorems, by themselves.”
—Bill Thurston



Handwritten diagram by Alexander Grothendieck

“Just want to know” camp

Yang-Hui He



AI taking on tasks that are currently the preserve of humans. Human mathematicians could become “priests to oracles”: curating and canonizing AI proofs.

	AI as an oracle
Role of AI	Autonomous researcher
What matters most?	Answers

“Just want to know” camp

Yang-Hui He



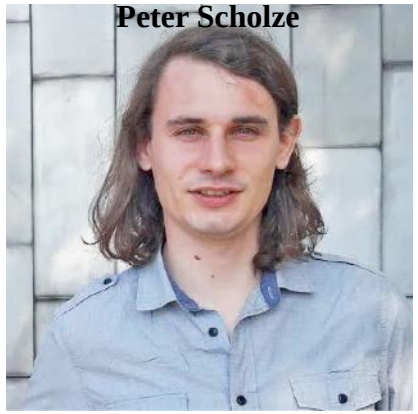
AI taking on tasks that are currently the preserve of humans. Human mathematicians could become “priests to oracles”: curating and canonizing AI proofs.

Human-centric camp

Akshay Venkatesh



Peter Scholze



Mathematics as a human endeavor independent of AI corporates. Prioritize human understanding. Treat AI as a tool, much like a calculator, or not used at all.

	AI as an oracle	AI as a tool
Role of AI	Autonomous researcher	Assistant or not used at all
What matters most?	Answers	Human understanding

“Just want to know” camp

Yang-Hui He



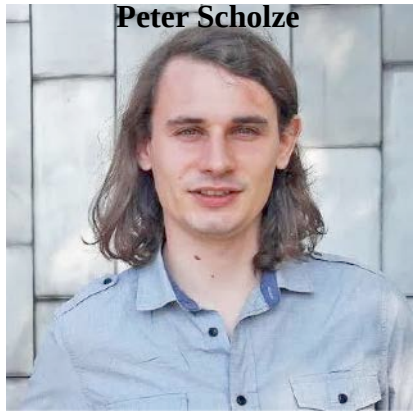
AI taking on tasks that are currently the preserve of humans. Human mathematicians could become “priests to oracles”: curating and canonizing AI proofs.

Human-centric camp

Akshay Venkatesh



Peter Scholze



Mathematics as a human endeavor independent of AI corporates. Prioritize human understanding. Treat AI as a tool, much like a calculator, or not used at at all.

Collaborative teamwork camp

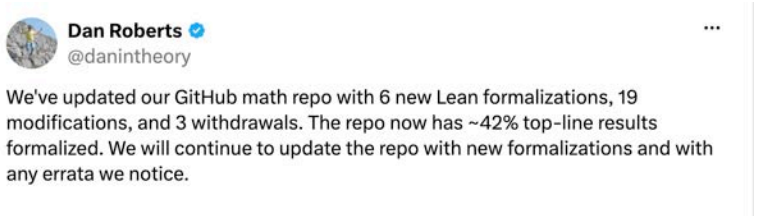
Terry Tao

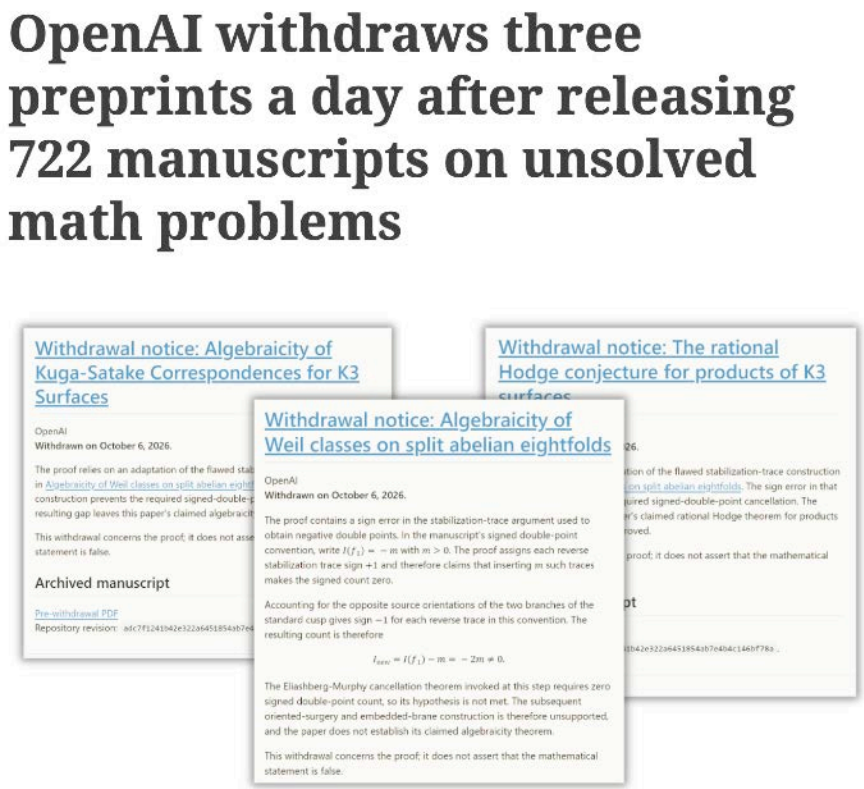


An era of “Big Mathematics (and science),” where humans and machines work together on complex problems that neither could solve alone.

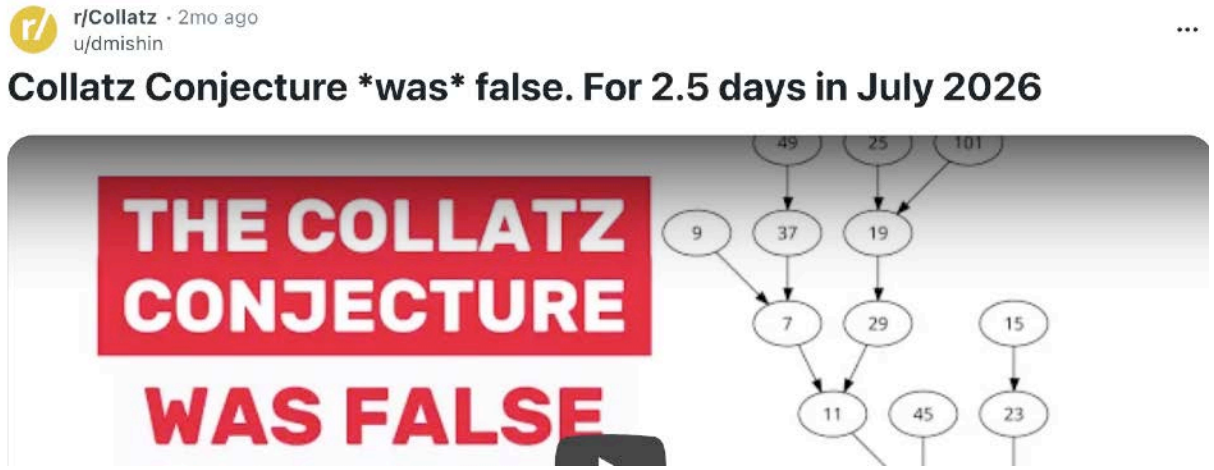
	AI as an oracle	AI as a tool	AI as a partner
Role of AI	Autonomous researcher	Assistant or not used at all	Collaborator
What matters most?	Answers	Human understanding	Verifiable Shared discovery

The risks of AI in mathematics and science





Only a fraction of the manuscripts were lean verified. On Oct. 7, OpenAI announced the withdrawal of three manuscripts because of a sign error invalidating several results.



- July 25, 2026: a machine-checked Lean proof asserting that the Collatz conjecture was false.
- **A bug in Lean**: the proof did not solve or disprove the Collatz conjecture. It bypassed the normal type-checking rules and allowed Lean to prove False with no axioms.
- The fake disproof was accepted as valid for ~2.5 days before Lean developers identified and patched the issue.

The risks of AI in mathematics and science

The risks of AI in mathematics and science

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.

The risks of AI in mathematics and science

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?



The risks of AI in mathematics and science

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?
- **Motivation**: jump straight to answers, skip the struggle, miss an opportunity to build the foundations and the **next generation** may suffer from intellectual atrophy.

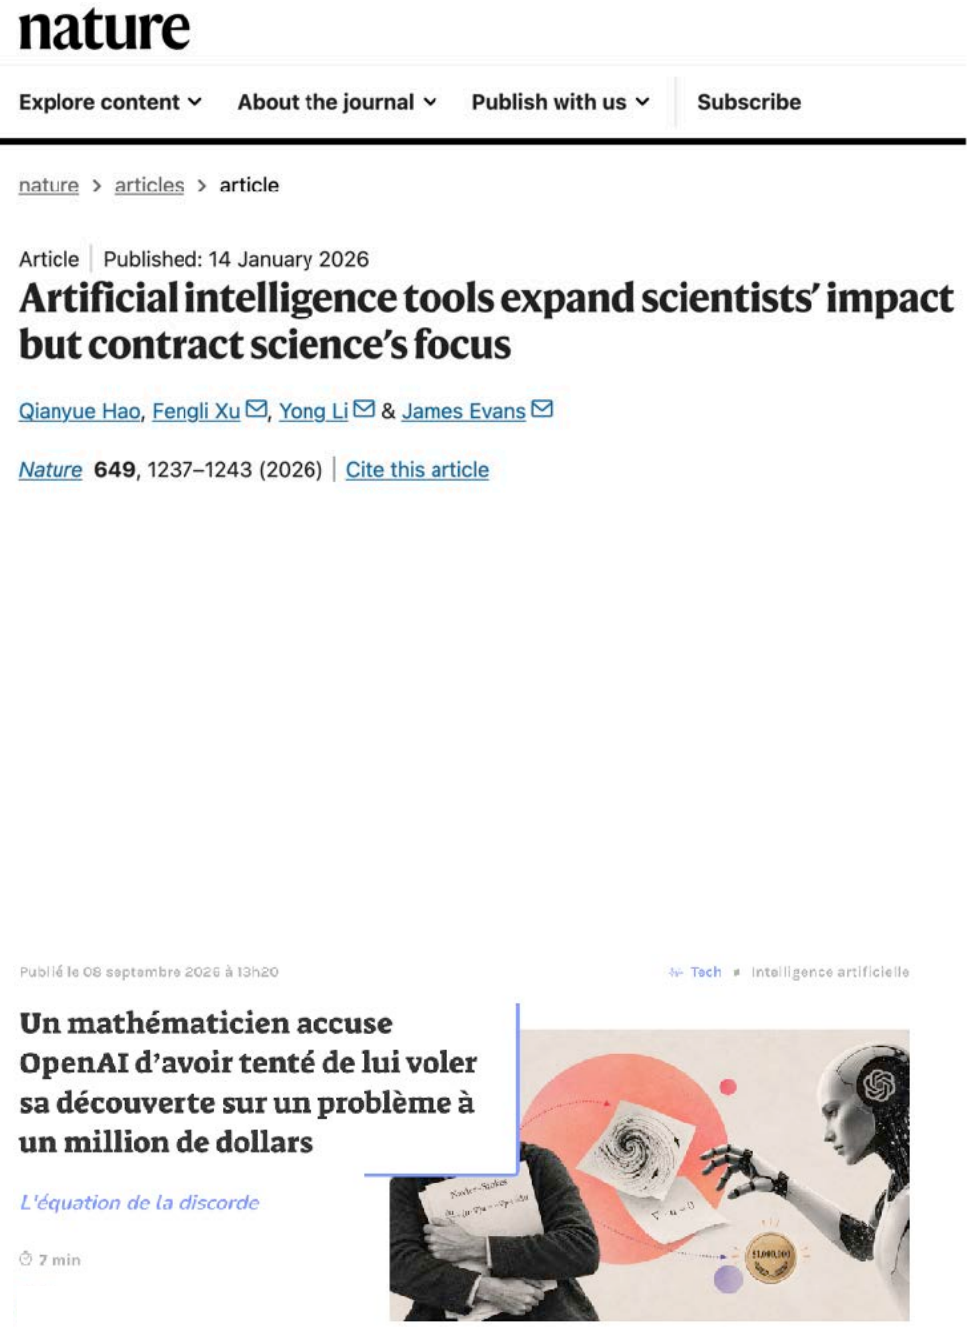


The risks of AI in mathematics and science



- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?
- **Motivation**: jump straight to answers, skip the struggle, miss an opportunity to build the foundations and the **next generation** may suffer from intellectual atrophy.
- **Values alignment**: the goals of major AI corporations may not always align with the aims of scientific research.

The risks of AI in mathematics and science



nature

Explore content · About the journal · Publish with us · Subscribe

nature > articles > article

Article | Published: 14 January 2026

Artificial intelligence tools expand scientists' impact but contract science's focus

Qianyue Hao, Fengli Xu, Yong Li & James Evans

Nature **649**, 1237–1243 (2026) | Cite this article

Publié le 08 septembre 2026 à 13h20

Tech · Intelligence artificielle

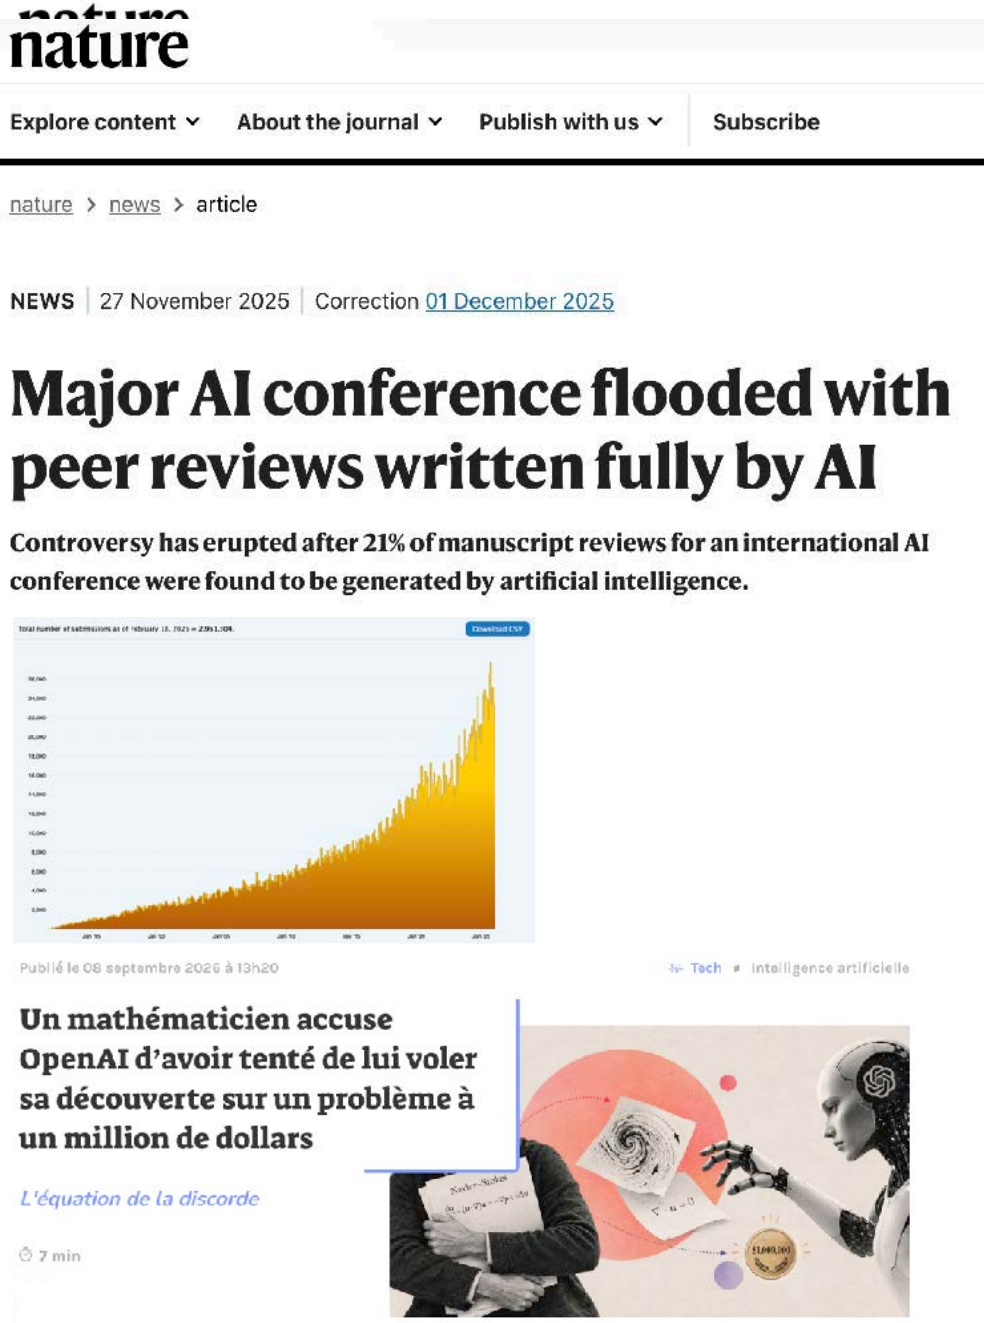
Un mathématicien accuse OpenAI d'avoir tenté de lui voler sa découverte sur un problème à un million de dollars

L'équation de la discorde

7 min

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?
- **Motivation**: jump straight to answers, skip the struggle, miss an opportunity to build the foundations and the **next generation** may suffer from intellectual atrophy.
- **Values alignment**: the goals of major AI corporations may not always align with the aims of scientific research.
- **Future role of researchers and students:** science understanding vs solving, verification burdens on the broader community, amplify unhealthy incentives around AI-generated results.

The risks of AI in mathematics and science


nature

Explore content ▾ About the journal ▾ Publish with us ▾ Subscribe

nature › news › article

NEWS | 27 November 2025 | Correction 01 December 2025

Major AI conference flooded with peer reviews written fully by AI

Controversy has erupted after 21% of manuscript reviews for an international AI conference were found to be generated by artificial intelligence.

Publié le 08 septembre 2026 à 13h20 Tech • Intelligence artificielle

Un mathématicien accuse OpenAI d'avoir tenté de lui voler sa découverte sur un problème à un million de dollars

L'équation de la discorde

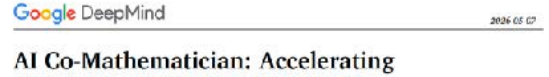
7 min


sakana.ai

2025.4.14

THE AI SCIENTIST-V2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search

Yutaro Yamada[1,*], Robert Tjarko Lange[1,*], Cong Lu[1,2,3,*], Shengran Hu[1,2,3], Chris Lu[1], Jakob Foerster[1], Jeff Clune[2,3,5,†] and David Ha[1,†]
[*]Equal Contribution, [1]Sakana AI, [2]University of British Columbia, [3]Vector Institute, [4]FLAIR, University of Oxford, [5]Canada CIFAR AI Chair, [†]Equal Advising


Google DeepMind

2026 05 07

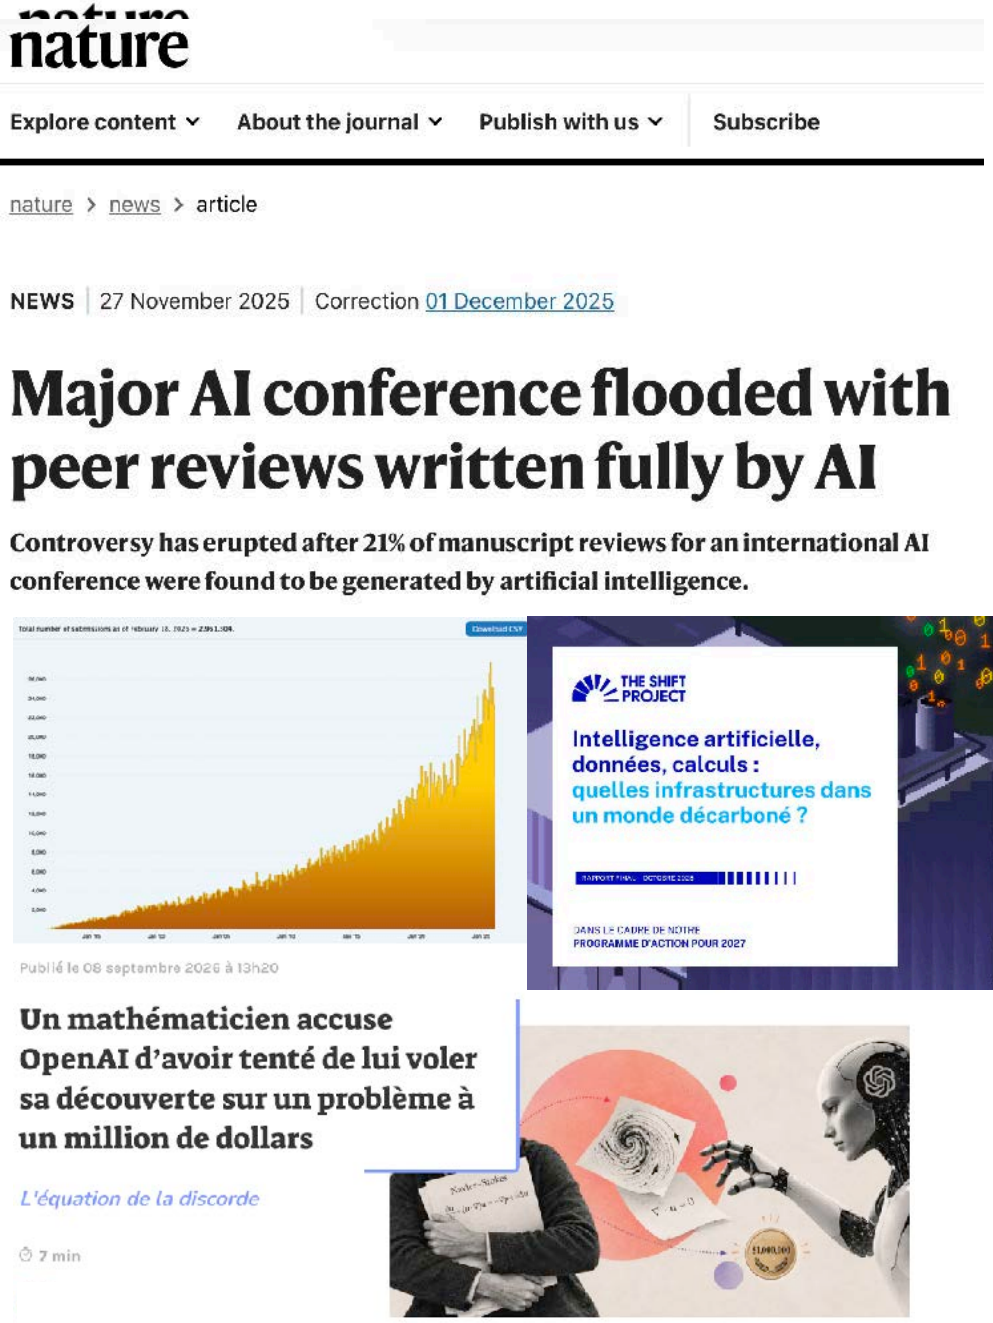
AI Co-Mathematician: Accelerating Mathematicians with Agentic AI

Daniel Zheng[*], Ingrid von Glehn[*], Yori Zwols[*], Iuliya Beloshapka[*], Lars Buesing[*], Daniel M. Roy[1,3], Martin Wattenberg[1], Bogdan Georgiev[1], Tatiana Schmidt[1], Andrew Cowie[1], Fernanda Viegas[1], Dimitri Kanevsky[1], Vineet Kahlon[2], Hartmut Maennel[1], Sophia Ali[2], George Holland[1], Alex Davies[1], and Pushmeet Kohli[1]
[1]Google DeepMind, [2]Google

We introduce the AI co-mathematician, a workbench for mathematicians to interactively leverage AI agents to pursue open-ended research. The AI co-mathematician is optimized to provide holistic support for the exploratory and iterative reality of mathematical workflows, including ideation, literature search, computational exploration, theorem proving and theory building. By providing an asynchronous, stateful workspace that manages uncertainty, refines user intent, tracks failed hypotheses, and outputs native mathematical artifacts, the system mirrors human collaborative workflows. In early tests, the AI co-mathematician helped researchers solve open problems, identify new research directions, and uncover overlooked literature references. Besides demonstrating a highly interactive paradigm for AI-assisted mathematical discovery, the AI co-mathematician also achieves state of the art results on hard problem-solving benchmarks, including scoring 48% on FrontierMath Tier 4, a new high score among all AI systems evaluated.

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?
- **Motivation**: jump straight to answers, skip the struggle, miss an opportunity to build the foundations and the **next generation** may suffer from intellectual atrophy.
- **Values alignment**: the goals of major AI corporations may not always align with the aims of scientific research.
- **Future role of researchers and students:** science understanding vs solving, verification burdens on the broader community, amplify unhealthy incentives around AI-generated results.
- Deluge of **publications** and **peer review:** mass production could destroy fertile ground instead of breathing life into new ideas.

The risks of AI in mathematics and science


nature

Explore content ▾ About the journal ▾ Publish with us ▾ Subscribe

nature › news › article

NEWS | 27 November 2025 | Correction 01 December 2025

Major AI conference flooded with peer reviews written fully by AI

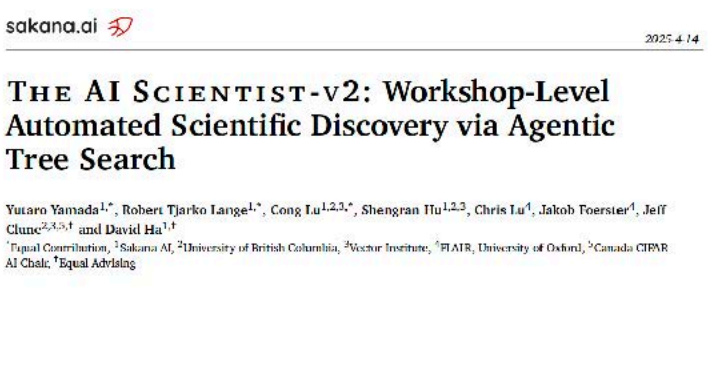
Controversy has erupted after 21% of manuscript reviews for an international AI conference were found to be generated by artificial intelligence.

Publié le 08 septembre 2026 à 13h20

Un mathématicien accuse OpenAI d'avoir tenté de lui voler sa découverte sur un problème à un million de dollars

L'équation de la discorde

7 min

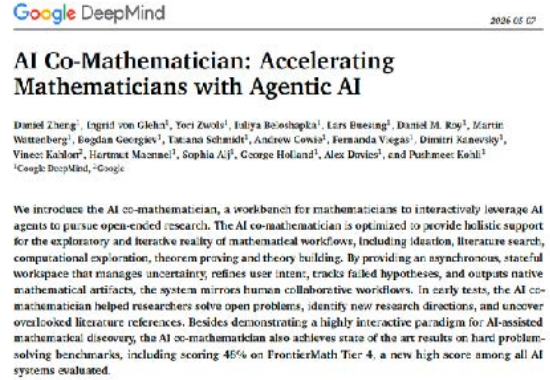

sakana.ai

2025.4.14

THE AI SCIENTIST-V2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search

Yutaro Yamada[1,*], Robert Tjarko Lange[1,*], Cong Lu[1,2,3,*], Shengran Hu[1,2,3], Chris Lu[1], Jakob Foerster[1], Jeff Clune[2,3,5,†] and David Ha[1,†]

[*]Equal Contribution, [1]Sakana AI, [2]University of British Columbia, [3]Vector Institute, [4]FLAIR, University of Oxford, [5]Canada CIFAR AI Chair, [†]Equal Advising


Google DeepMind

2026.05.07

AI Co-Mathematician: Accelerating Mathematicians with Agentic AI

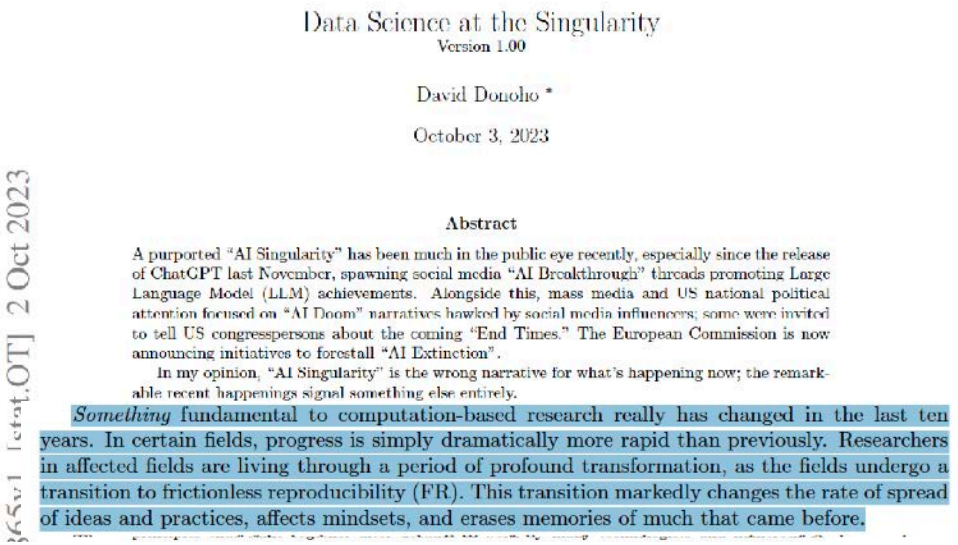
Daniel Zheng[1], Ingrid von Glehn[1], Yori Zwols[1], Iuliya Beloshapka[1], Lars Buesing[1], Daniel M. Roy[1,3], Martin Wattenberg[1], Bogdan Georgiev[1], Tatiana Schmidt[1], Andrew Cowie[1], Fernanda Viegas[1], Dimitri Kanevsky[1], Vineet Kahlon[2], Hartmut Maennel[1], Sophia Ali[1], George Holland[1], Alex Davies[1], and Pushmeet Kohli[1]

[1]Google DeepMind, [2]Google

We introduce the AI co-mathematician, a workbench for mathematicians to interactively leverage AI agents to pursue open-ended research. The AI co-mathematician is optimized to provide holistic support for the exploratory and iterative reality of mathematical workflows, including ideation, literature search, computational exploration, theorem proving and theory building. By providing an asynchronous, stateful workspace that manages uncertainty, refines user intent, tracks failed hypotheses, and outputs native mathematical artifacts, the system mirrors human collaborative workflows. In early tests, the AI co-mathematician helped researchers solve open problems, identify new research directions, and uncover overlooked literature references. Besides demonstrating a highly interactive paradigm for AI-assisted mathematical discovery, the AI co-mathematician also achieves state of the art results on hard problem-solving benchmarks, including scoring 48% on FrontierMath Tier 4, a new high score among all AI systems evaluated.

- **Equal opportunities and access to AI tools**: if the slow, deliberative research process is no longer valued by society and research funders, then research could become an elitist activity, only practiced by organizations that can afford proprietary AI models.
- **Research misconduct**: IP, honest credit, integrity, ethics, what scientific discovery/result, authorship, collaboration, and intelligence now mean ?
- **Motivation**: jump straight to answers, skip the struggle, miss an opportunity to build the foundations and the **next generation** may suffer from intellectual atrophy.
- **Values alignment**: the goals of major AI corporations may not always align with the aims of scientific research.
- **Future role of researchers and students:** science understanding vs solving, verification burdens on the broader community, amplify unhealthy incentives around AI-generated results.
- Deluge of **publications** and **peer review:** mass production could destroy fertile ground instead of breathing life into new ideas.
- **Ecological** impact.

A profound transformation



Data Science at the Singularity
Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

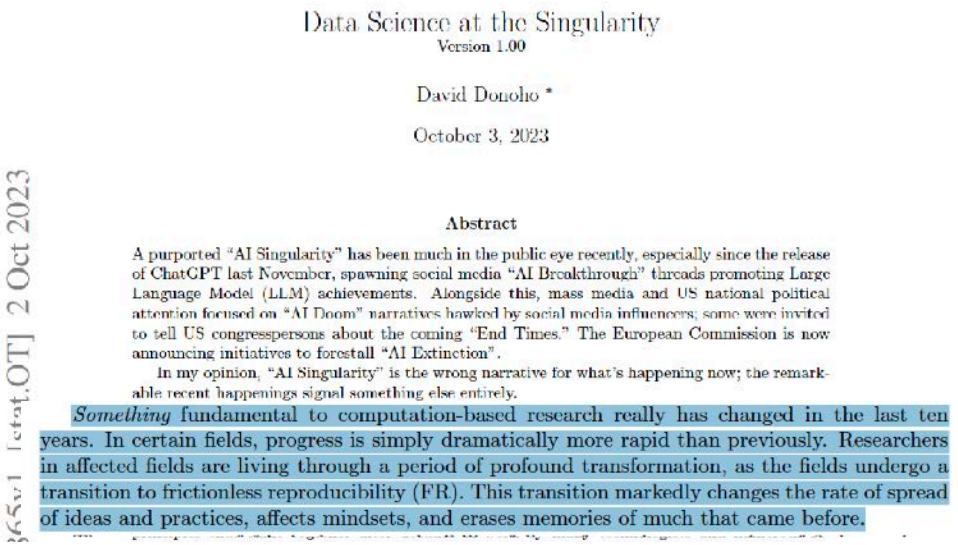
A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.



Data Science at the Singularity
Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

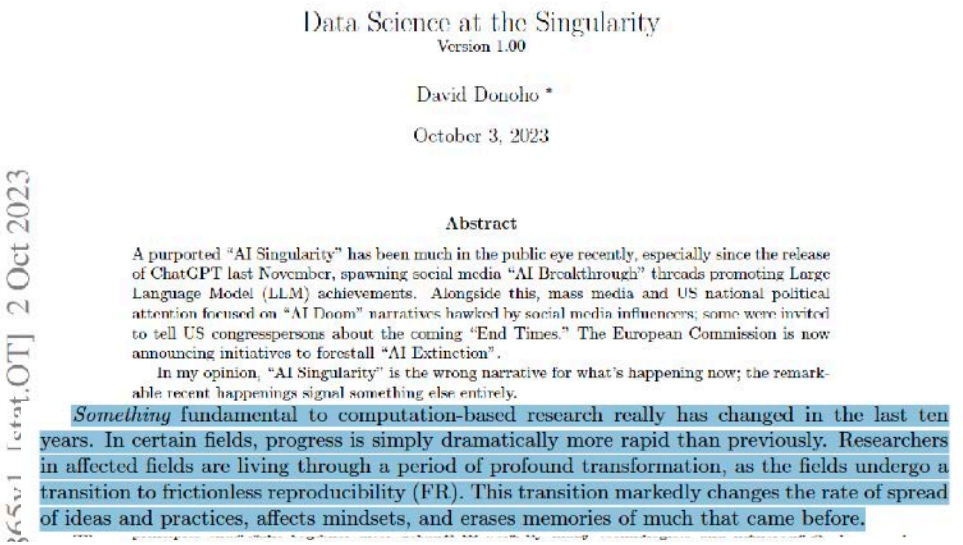
A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.
- What is happening should not be underestimated.



Data Science at the Singularity

Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

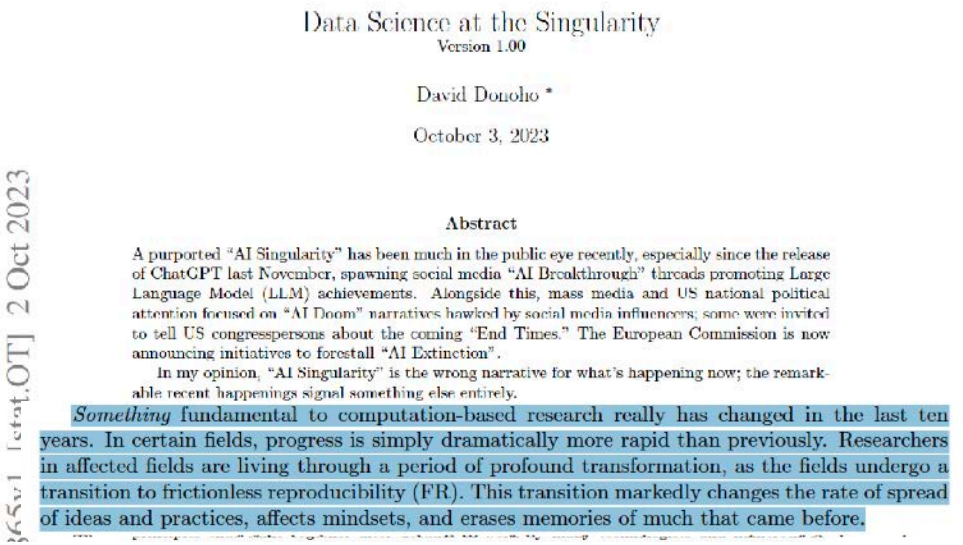
A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.
- What is happening should not be underestimated.
- AI changes the way research work is done.



Data Science at the Singularity

Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

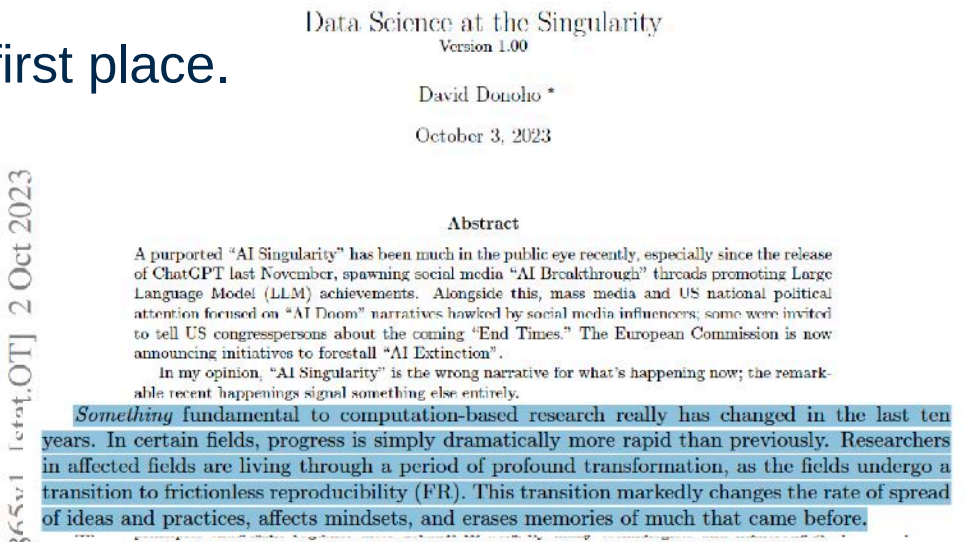
A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.
- What is happening should not be underestimated.
- AI changes the way research work is done.
- Do not lose sight of what research was meant to achieve in the first place.



Data Science at the Singularity

Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

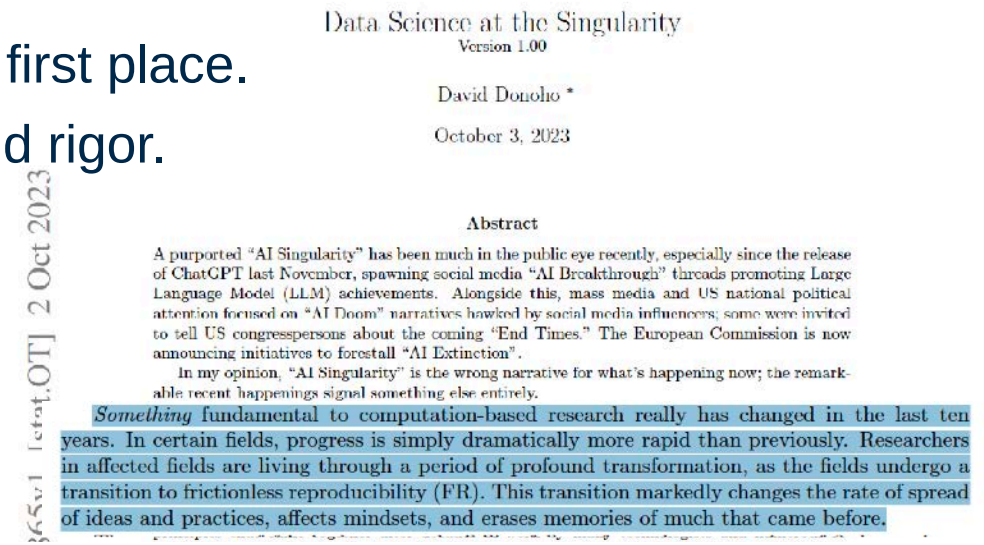
A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.
- What is happening should not be underestimated.
- AI changes the way research work is done.
- Do not lose sight of what research was meant to achieve in the first place.
- The scientific community should engage with it with curiosity and rigor.



Data Science at the Singularity

Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

A profound transformation

- An extraordinary opportunity and transformation.
- What is happening should not be underestimated.
- AI changes the way research work is done.
- Do not lose sight of what research was meant to achieve in the first place.
- The scientific community should engage with it with curiosity and rigor.
- Moving swiftly but not hastily.

https://agmai.org/

Advisory Group on Mathematics and Artificial Intelligence

Math and AI.

https://mathandai.org/
25 fields medalists

A Severe Misalignment of AI in Mathematics

The Association for Human Mathematics (AHM)

https://www.ahmath.org/

The Association for Human Mathematics (AHM) protects mathematics as a human endeavor against the threat of artificial intelligence. We work to maintain the independence of the mathematical community from the encroachment of corporate interest. In light of the grave societal, environmental, and economic impacts of the AI industry and the threat of unfettered AI growth, we strive not to be complicit.

Data Science at the Singularity
Version 1.00

David Donoho *

October 3, 2023

arXiv:2310.00865v1 [stat.OT] 2 Oct 2023

Abstract

A purported "AI Singularity" has been much in the public eye recently, especially since the release of ChatGPT last November, spawning social media "AI Breakthrough" threads promoting Large Language Model (LLM) achievements. Alongside this, mass media and US national political attention focused on "AI Doom" narratives hawked by social media influencers; some were invited to tell US congresspersons about the coming "End Times." The European Commission is now announcing initiatives to forestall "AI Extinction".

In my opinion, "AI Singularity" is the wrong narrative for what's happening now; the remarkable recent happenings signal something else entirely.

Something fundamental to computation-based research really has changed in the last ten years. In certain fields, progress is simply dramatically more rapid than previously. Researchers in affected fields are living through a period of profound transformation, as the fields undergo a transition to frictionless reproducibility (FR). This transition markedly changes the rate of spread of ideas and practices, affects mindsets, and erases memories of much that came before.

Leiden Declaration on Artificial Intelligence and Mathematics

https://leidendeclaration.ai/

Sign Signatories News About Contact Report

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).

Sign this Declaration | Read Declaration | Download Declaration | Endorsements | View signatories

Share the declaration



Training in the era of AI



Training in the era of AI



« À mauvais ouvrier, point de bon outil »

« No tool is good enough for a poor craftsman »

CNRS
AISSAI CENTER

OLIVIER CAPPÉ CLAIRE MARC

TOUT COMPRENDRE
(OU PRESQUE)
SUR L'INTELLIGENCE ARTIFICIELLE

