



Universidad
Politécnica
de Cartagena

MEMBER OF



EUROPEAN
UNIVERSITY OF
TECHNOLOGY

AI co-scientist agents, and other tools for assisting research tasks: A look to the future

Pablo Pavón Mariño

Vicerrector de Estrategia Digital e Inteligencia Artificial

Universidad Politécnica de Cartagena

pablo.pavon@upct.es

12 February 2026

Agenda

1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages



Agenda

1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages



Sounds like a great idea!

Let's check if it fits, if it has been already investigated, find
potencial industry cases... analyze the state-of-the-art

Specialized tools



...

or deep research functions in...



...

linked to...



...

OpenAI

February 2, 2025 Release

Introducing deep research

An agent that uses reasoning to synthesize large amounts of online information and complete multi-step research tasks for you. Available to Pro users today, Plus and Team next.

Try on ChatGPT ↗

All, global perspective

I'll compile a detailed research report on how the retail industry has transformed over the past three years on a global scale. This will cover:

- The rise of e-commerce and shifts in brick-and-mortar retail
- Changes in consumer behavior and expectations
- Supply chain disruptions and adaptations
- The role of AI, automation, and emerging technologies
- The impact of inflation and economic fluctuations
- Sustainability and ethical retailing trends

I'll structure the report with bullet points and tables where necessary to enhance clarity. I'll let you



nature

Explore content ▾

About the journal ▾

Publish with us ▾

Subscribe

[nature](#) > [news](#) > [article](#)

NEWS | 06 February 2025

OpenAI's 'deep research' tool: is it useful for scientists?

The model produces cited, pages-long reports that might be helpful for generating literature reviews.

By [Nicola Jones](#)



Technology giant OpenAI has unveiled a pay-for-access tool called 'deep research', which synthesizes information from dozens or hundreds of websites into a cited report several pages long. The tool follows a similar one from Google, called 'Deep Research', released last December and acts as a personal assistant, doing the equivalent of hours of work in tens of minutes.



Great, this sounds promising! Now, implement this software, and make all the tests in realistic scenarios, compile the results, so we can validate and refine this idea!!

addressing the “I don’t know how to code” barrier...
the **vibe coding** concept



Andrej Karpathy ✓

@karpathy

There's a new kind of coding I call "vibe coding", where you fully give in to the vibes, embrace exponentials, and forget that the code even exists. It's possible because the LLMs (e.g. Cursor Composer w Sonnet) are getting too good. Also I just talk to Composer with SuperWhisper so I barely even touch the keyboard. I ask for the dumbest things like "decrease the padding on the sidebar by half" because I'm too lazy to find it. I "Accept All" always, I don't read the diffs anymore. When I get error messages I just copy paste them in with no comment, usually that fixes it. The code grows beyond my usual comprehension, I'd have to really read through it for a while. Sometimes the LLMs can't fix a bug so I just work around it or ask for random changes until it goes away. It's not too bad for throwaway weekend projects, but still quite amusing. I'm building a project or webapp, but it's not really coding - I just see stuff, say stuff, run stuff, and copy paste stuff, and it mostly works.

[Traducir post](#)

12:17 a. m. · 3 feb. 2025 · 5M Visualizaciones

<https://www.youtube.com/shorts/b2Ki4-lc5dk>



GitHub Copilot



Antigravity



From vibe coding to SDD (Spec Driven Development)

[Example: Codex](#)

[Example: Google Colab](#)



That was fun to investigate

Unfortunately the party is over: we have to write the paper

Overleaf



Prism

Gemini



ChatGPT

Jenni



Claude

Example: <https://openai.com/prism/>



Now things get worse...

We need to make the midterm report of the project...

compile all the info, write summaries, justify alignment with project objectives...

[Example \(Notebook LM\)](#)



Agenda

1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages

AlphaFold, AlphaGenome

Accelerating scientific Discovery with AI (12 Marzo 2025), Demis Hassabis (CEO Google Deepmind),
<https://www.youtube.com/watch?v=hHooQmmzG4k>

Proteins are the building blocks of life



AlphaFold, AlphaGenome

- **Before 2018.** Understanding the shape of a protein or the effect of a change in DNA required years of laboratory experiments.
- **AlphaFold (2018).** AI learns to predict the 3D shape of a protein.
- **AlphaFold2 (2020 , Nobel Prize in Chemistry 2024).** Solves a historic problem: near-perfect predictions, comparable to real experiments.
- **AlphaFold3 (2024).** AI models how proteins, DNA, and drugs interact, accelerating drug design. <https://alphafoldserver.com>
- **AlphaGenome (2025).** AI predicts what happens when DNA changes, helping to understand genetic diseases. <https://deepmind.google.com/science/alphagenome/>

“Open science” concept

AlphaFold2 available, even in the cloud without local installation!!

<https://colab.research.google.com/github/sokrypton/ColabFold/blob/main/AlphaFold2.ipynb>

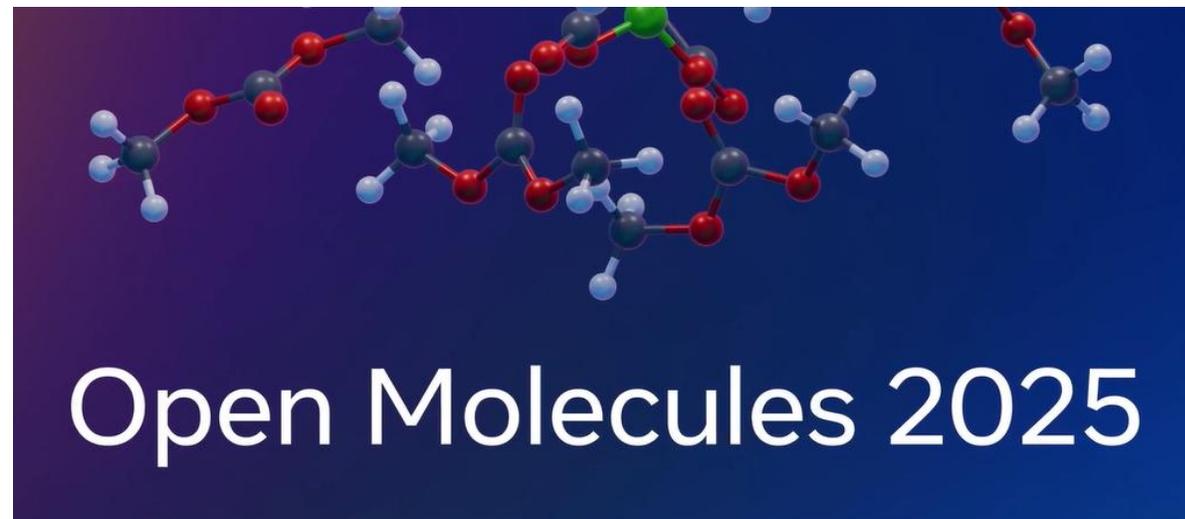
The Open Molecules 2025 (OMol25) Dataset, Evaluations, and Models

Daniel S. Levine^{1,*}, Muhammed Shuaibi^{1,*}, Evan Walter Clark Spotte-Smith², Michael G. Taylor³, Muhammad R. Hasyim⁴, Kyle Michel¹, Ilyes Batatia⁵, Gábor Csányi⁵, Misko Dzamba¹, Peter Eastman⁶, Nathan C. Frey⁷, Xiang Fu¹, Vahe Gharakhanyan¹, Aditi S. Krishnapriyan^{8,9,10}, Joshua A. Rackers⁷, Sanjeev Raja⁸, Ammar Rizvi¹, Andrew S. Rosen¹¹, Zachary Ulissi¹, Santiago Vargas^{12,13}, C. Lawrence Zitnick^{1,†}, Samuel M. Blau^{13,†}, Brandon M. Wood^{1,†}

¹FAIR at Meta, ²Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA, ³Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM, USA, ⁴Simons Center for Computational Physical Chemistry, New York University, New York, NY, USA, ⁵Engineering Laboratory, University of Cambridge, Cambridge, UK, ⁶Department of Chemistry, Stanford University, Stanford, CA, USA, ⁷Prescient Design, Genentech, New York, NY, USA, ⁸Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, USA, ⁹Department of Chemical and Biomolecular Engineering, University of California, Berkeley, CA, USA, ¹⁰Applied Mathematics and Computational Research, Lawrence Berkeley National Laboratory, Berkeley, CA, USA, ¹¹Department of Chemical and Biological Engineering, Princeton University, Princeton, NJ, USA, ¹²Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA, ¹³Energy Technologies Area, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
*Co-first Author, †Co-corresponding Author

Machine learning (ML) models hold the promise of transforming atomic simulations by delivering quantum chemical accuracy at a fraction of the computational cost. Realization of this potential would enable high-throughput, high-accuracy molecular screening campaigns to explore vast regions of chemical space and facilitate *ab initio*-level simulations at sizes and time scales that were previously inaccessible. However, a fundamental challenge to creating ML models that perform well across molecular chemistry is the lack of comprehensive data for training. Despite substantial efforts in data generation, no large-scale molecular dataset exists that combines broad chemical diversity with a high level of accuracy. To address this gap, Meta FAIR introduces Open Molecules 2025 (OMol25), a large-scale dataset composed of more than 100 million density functional theory (DFT) calculations at the ω B97M-V/def2-TZVPD level of theory, representing billions of CPU core-hours of compute.

<https://arxiv.org/pdf/2505.08762>



- Predictive model of molecular properties, with high accuracy for larger sizes and scales.
- Open (model, dataset, and documentation).

<https://huggingface.co/facebook/OMol25>

UMA: A Family of Universal Models for Atoms

Brandon M. Wood^{1,†*}, Misko Dzamba^{1,*}, Xiang Fu^{1,*}, Meng Gao^{1,*}, Muhammed Shuaibi^{1,*}, Luis Barroso-Luque¹, Kareem Abdelmaqsood², Vahe Gharakhanyan¹, John R. Kitchin², Daniel S. Levine¹, Kyle Michel¹, Anuroop Sriram¹, Taco Cohen¹, Abhishek Das¹, Ammar Rizvi¹, Sushree Jagriti Sahoo¹, Zachary W. Ulissi¹, C. Lawrence Zitnick^{1,†}

¹FAIR at Meta, ²Department of Chemical Engineering, Carnegie Mellon University

*Co-first Author, †Co-corresponding Author

The ability to quickly and accurately compute properties from atomic simulations is critical for advancing a large number of applications in chemistry and materials science including drug discovery, energy storage, and semiconductor manufacturing. To address this need, Meta FAIR presents a family of Universal Models for Atoms (UMA), designed to push the frontier of speed, accuracy, and generalization. UMA models are trained on half a billion unique 3D atomic structures (the largest training runs to date) by compiling data across multiple chemical domains, e.g. molecules, materials, and catalysts. We develop empirical scaling laws to help understand how to increase model capacity alongside dataset size to achieve the best accuracy. The UMA small and medium models utilize a novel architectural design we refer to as mixture of linear experts that enables increasing model capacity without sacrificing speed. For example, UMA-medium has 1.4B parameters but only ~50M active parameters per atomic structure. We evaluate UMA models on a diverse set of tasks across multiple domains and find that, remarkably, a single model without any fine-tuning can perform similarly or better than specialized models. We are releasing the UMA code, weights, and associated data to accelerate computational workflows and enable the community to continue to build increasingly capable AI models.

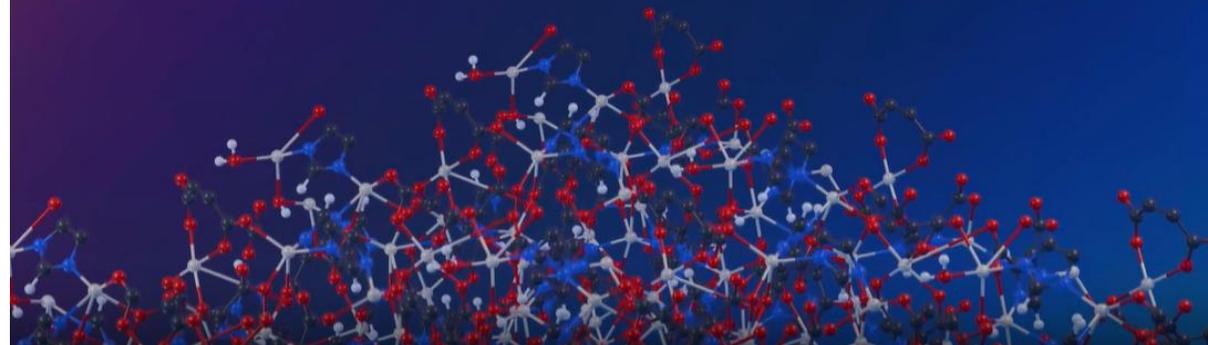
Models: <https://huggingface.co/facebook/UMA>

Code: <https://github.com/facebookresearch/fairchem>

Correspondence: B.M.W. (bmwood@meta.com), C.L.Z. (zitnick@meta.com)



Universal Model for Atoms (UMA)



- Predictor of properties of 3D atomic configurations.
- Open (model and documentation).

<https://huggingface.co/facebook/UMA>

<https://github.com/facebookresearch/fairchem>

<https://ai.meta.com/research/publications/uma-a-family-of-universal-models-for-atoms/>

(May 14th 2025)

The image is a screenshot of a web browser displaying a news article from MIT. The browser's address bar shows the URL: news.mit.edu/2025/new-tool-makes-generative-ai-models-likely-create-breakthrough-materials-0922. The MIT logo and name are visible in the top left, along with navigation links for Education, Research, Innovation, Admissions + Aid, Campus Life, News, Alumni, and About MIT. The article title is "New tool makes generative AI models more likely to create breakthrough materials". Below the title is a sub-headline: "With SCIGEN, researchers can steer AI models to create materials with exotic properties for applications like quantum computing." The author is listed as "Zach Winn | MIT News" and the date as "September 22, 2025". A red circle highlights the date. The main image shows a 3D visualization of a kagome lattice structure, consisting of orange, blue, and white spheres arranged in a complex geometric pattern. To the right of the image, there is a "PRESS INQUIRIES" link.

MIT Massachusetts
Institute of
Technology

Education Research Innovation Admissions + Aid Campus Life **News** Alumni About MIT

MIT News

ON CAMPUS AND AROUND THE WORLD

[SUBSCRIBE](#) [BROWSE](#) [SEARCH NEWS](#)

New tool makes generative AI models more likely to create breakthrough materials

With SCIGEN, researchers can steer AI models to create materials with exotic properties for applications like quantum computing.

Zach Winn | MIT News
September 22, 2025

[PRESS INQUIRIES](#)

The researchers applied their technique to generate millions of candidate materials consisting of geometric lattice structures associated with quantum properties. The kagome lattice, represented here, can support the creation of materials that could be useful for quantum computing.

Image: Jose-Luis Olivares, MIT; iStock

<https://news.mit.edu/2025/new-tool-makes-generative-ai-models-likely-create-breakthrough-materials-0922>
<https://www.nature.com/articles/s41563-025-02355-y> (Sept 2025)

Google DeepMind

Models

Research

Science

About

Try Google AI Studio

Try Gemini



SCIENCE

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

14 MAY 2025

By AlphaEvolve team

Share



<https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/>

May 2025



AlphaEvolve: A coding agent for scientific and algorithmic discovery

Alexander Novikov*, Ngán Vũ*, Marvin Eisenberger*, Emilien Dupont*, Po-Sen Huang*, Adam Zsolt Wagner*, Sergey Shirobokov*, Borislav Kozlovskii*, Francisco J. R. Ruiz, Abbas Mehrabian, M. Pawan Kumar, Abigail See, Swarat Chaudhuri, George Holland, Alex Davies, Sebastian Nowozin, Pushmeet Kohli and Matej Balog*
Google DeepMind¹

In this white paper, we present *AlphaEvolve*, an evolutionary coding agent that substantially enhances capabilities of state-of-the-art LLMs on highly challenging tasks such as tackling open scientific problems or optimizing critical pieces of computational infrastructure. *AlphaEvolve* orchestrates an autonomous pipeline of LLMs, whose task is to improve an algorithm by making direct changes to the code. Using an evolutionary approach, continuously receiving feedback from one or more evaluators, *AlphaEvolve* iteratively improves the algorithm, potentially leading to new scientific and practical discoveries. We demonstrate the broad applicability of this approach by applying it to a number of important computational problems. When applied to optimizing critical components of large-scale computational stacks at Google, *AlphaEvolve* developed a more efficient scheduling algorithm for data centers, found a functionally equivalent simplification in the circuit design of hardware accelerators, and accelerated the training of the LLM underpinning *AlphaEvolve* itself. Furthermore, *AlphaEvolve* discovered novel, provably correct algorithms that surpass state-of-the-art solutions on a spectrum of problems in mathematics and computer science, significantly expanding the scope of prior automated discovery methods (Romera-Paredes et al., 2023). Notably, *AlphaEvolve* developed a search algorithm that found a procedure to multiply two 4×4 complex-valued matrices using 48 scalar multiplications; offering the first improvement, after 56 years, over Strassen’s algorithm in this setting. We believe *AlphaEvolve* and coding agents like it can have a significant impact in improving solutions of problems across many areas of science and computation.

<https://arxiv.org/abs/2506.13131>

https://colab.research.google.com/github/google-deepmind/alphaevolve_results/blob/master/mathematical_results.ipynb

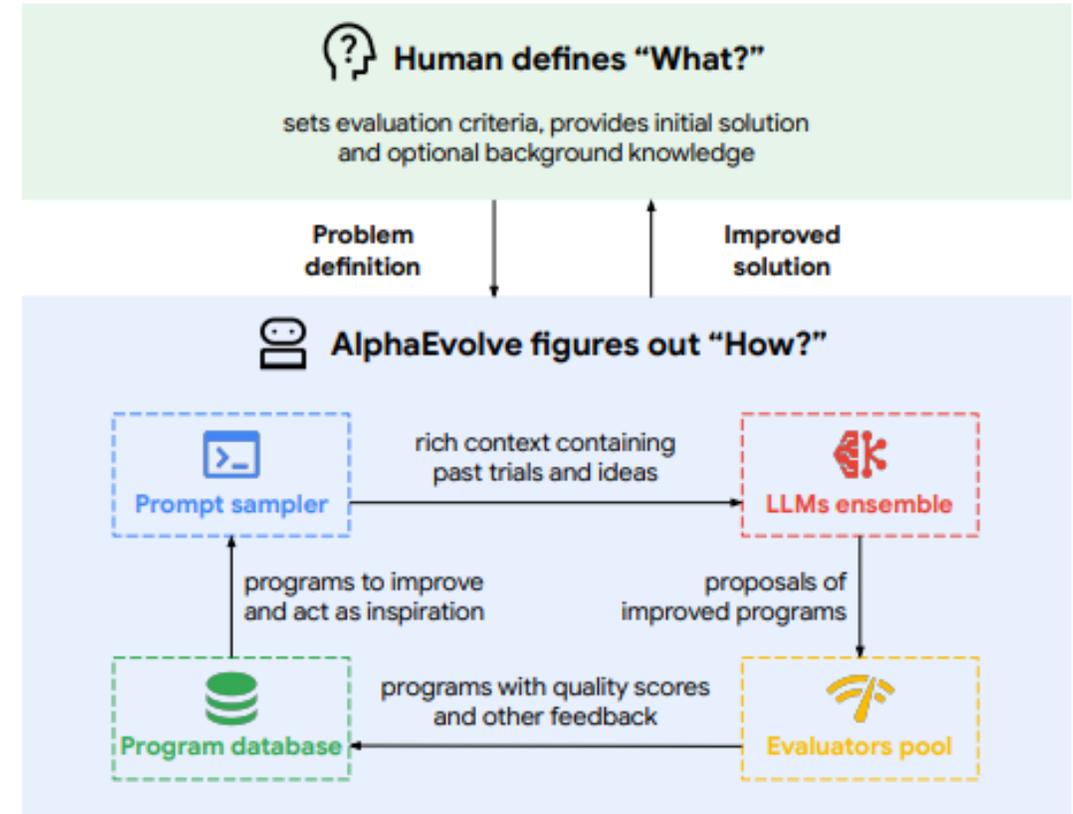


Figure 1 | AlphaEvolve high-level overview.

Formal Mathematical Reasoning: A New Frontier in AI

Kaiyu Yang¹, Gabriel Poesia², Jingxuan He³,
Wenda Li⁴, Kristin Lauter¹, Swarat Chaudhuri⁵, Dawn Song³

¹Meta FAIR, ²Stanford University, ³UC Berkeley, ⁴University of Edinburgh, ⁵UT Austin

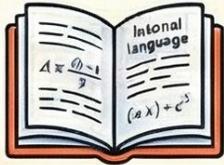
Abstract

AI for Mathematics (AI4Math) is not only intriguing intellectually but also crucial for AI-driven discovery in science, engineering, and beyond. Extensive efforts on AI4Math have mirrored techniques in NLP, in particular, training large language models on carefully curated math datasets in text form. As a complementary yet less explored avenue, *formal mathematical reasoning* is grounded in formal systems such as proof assistants, which can verify the correctness of reasoning and provide automatic feedback. In this position paper, we advocate for formal mathematical reasoning and argue that it is indispensable for advancing AI4Math to the next level. In recent years, we have seen steady progress in using AI to perform formal reasoning, including core tasks such as theorem proving and autoformalization, as well as emerging applications such as verifiable generation of code and hardware designs. However, significant challenges remain to be solved for AI to truly master mathematics and achieve broader impact. We summarize existing progress, discuss open challenges, and envision critical milestones to measure future success. At this inflection point for formal mathematical reasoning, we call on the research community to come together to drive transformative advancements in this field.

<https://arxiv.org/html/2412.16075v1#bib.bib49>

<https://lean-lang.org/>

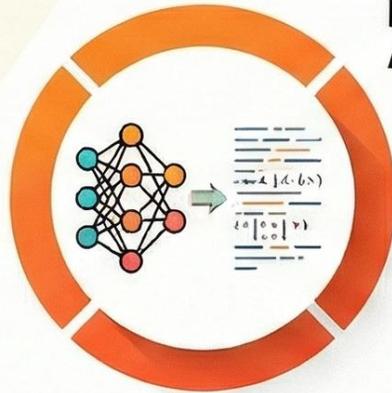
AI for Mathematics: Two Paths to Formal Proof



Informal Math

Modern AI4Math is moving beyond "informal" guessing toward "formal" reasoning grounded in systems like Lean. This transition involves bridging the gap between natural language mathematics and machine-verifiable proofs through two distinct strategic workflows.

Path 1: Sequential Autoformalization and Proving



Autoformalization (Model)

Step 1: Informal to Formal Statement

An AI model translates natural language mathematical text into a structured formal statement.

```

(0 ≤ m < n) → 0 ≤ m
∀ k (0 ≤ k < m) → 0 ≤ k
∃ x (x = (a + i) * θ + a, 1 ≤ i ≤ n)
∃ x ∈ V
x ∈ E ∈ H
    
```

Formal Statement

Neural Theorem Proving (Model)

Step 2: Neural Theorem Proving

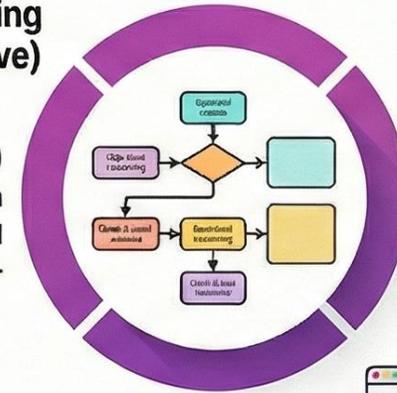
A specialized model searches for a formal proof based on the translated statement.



Outcome: Machine-Verifiable Formal Proof

The final output is a complete, sound proof confirmed by a proof assistant.

Path 2: Verified Reasoning (Draft, Sketch, and Prove)



General LLM Reasoning (Proof Sketch)

Step 1: Informal Math to Proof Sketch

A general LLM generates a high-level reasoning "sketch" in natural language.

Translate Steps to Formal (Lean)

```

def math_eq (a b : ℕ) : Prop := a = b
def math_add (a b : ℕ) : ℕ := a + b
def math_sub (a b : ℕ) : ℕ := a - b
def math_mul (a b : ℕ) : ℕ := a * b
    
```



Formal Verification

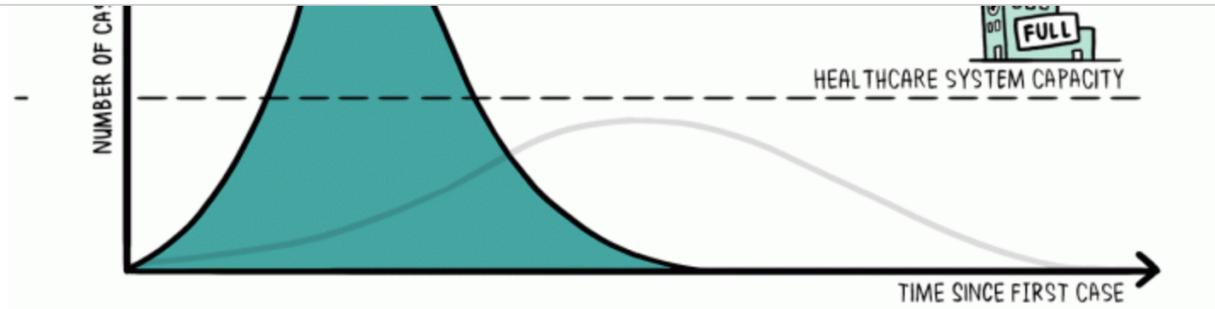
Step 2: Formal Translation & Verification

Informal reasoning steps are translated into Lean code for rigorous symbolic verification.

Outcome: Verified Natural Language Reasoning

Combines the flexibility of human-like reasoning with the rigor of formal systems.





Using
Aristotle AI model +
AlphaEvolve

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 AMS open math notes
- Anonymous on AMS open math notes
- Terence Tao on The Lebesgue differentiation t...
- Terence Tao on 275A, Notes 0: Foundations of...
- Anonymous on The Lebesgue differentiation t...
- Anonymous on 275A, Notes 0: Foundations of...
- Anonymous on Heuristic limitations of the c...

The story of Erdős problem #1026

8 December, 2025 in expository, math.CA, math.CO | Tags: AI, AlphaEvolve, Erdos | by Terence Tao

Problem 1026 on the Erdős problem web site recently got solved through an interesting combination of existing literature, online collaboration, and AI tools. The purpose of this blog post is to try to tell the story of this collaboration, and also to supply a complete proof.

The original problem of Erdős, posed in 1975, is rather ambiguous. Erdős starts by recalling his famous theorem with Szekeres that says that given a sequence of $k^2 + 1$ distinct real numbers, one can find a subsequence of length $k + 1$ which is either increasing or decreasing; and that one cannot improve the $k^2 + 1$ to k^2 , by considering for instance a sequence of k blocks of length k , with the numbers in each block decreasing, but the blocks themselves increasing. He also noted a result of Hanani that every sequence of length $k(k + 3)/2$ can be decomposed into the union of k monotone sequences. He then wrote "As far as I know the following question is not yet settled. Let

<https://terrytao.wordpress.com/2025/12/08/the-story-of-erdos-problem-126/>

<https://aristotle.harmonic.fun/>



pablo.pavon@upct.es



 Ask Aristotle

 API Keys

 Request History

 Docs

What would you like to prove?

Describe what you want to prove...





Agenda

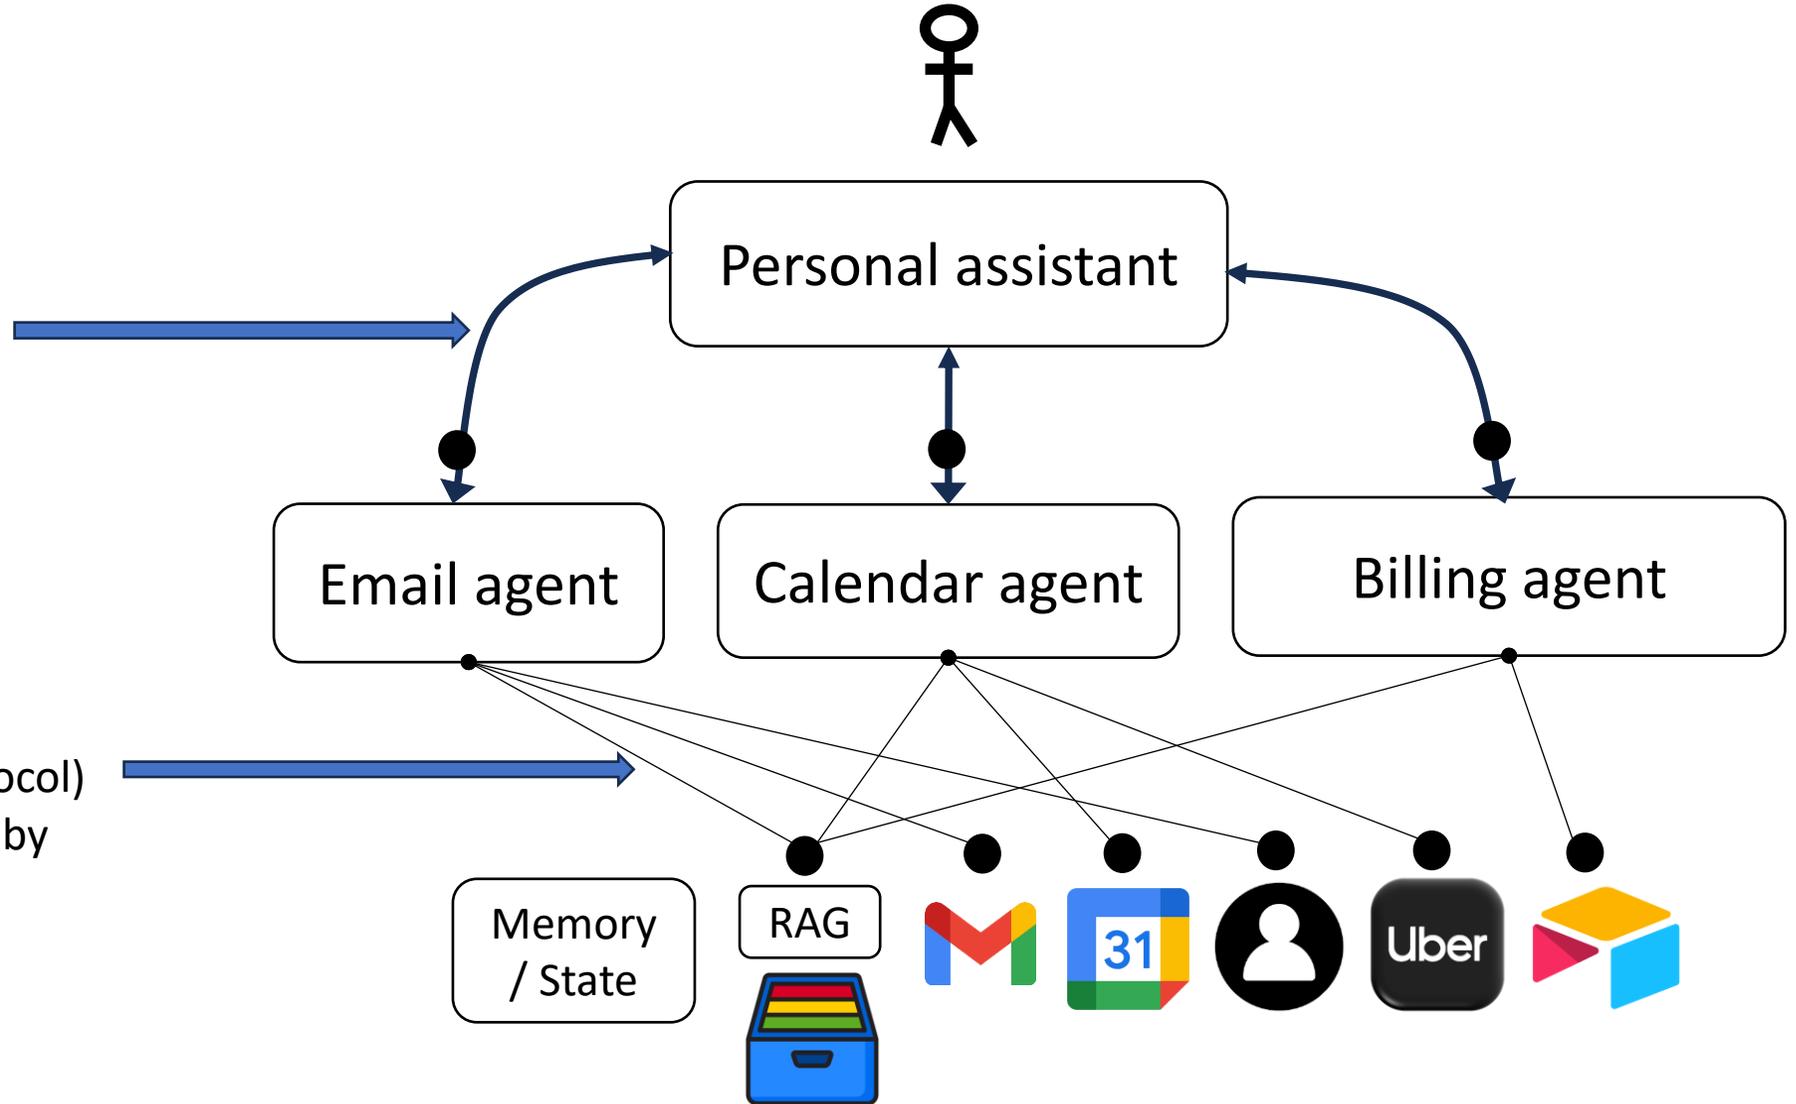
1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages



An AI Agent is an AI system that, in an autonomous form, makes tasks for you

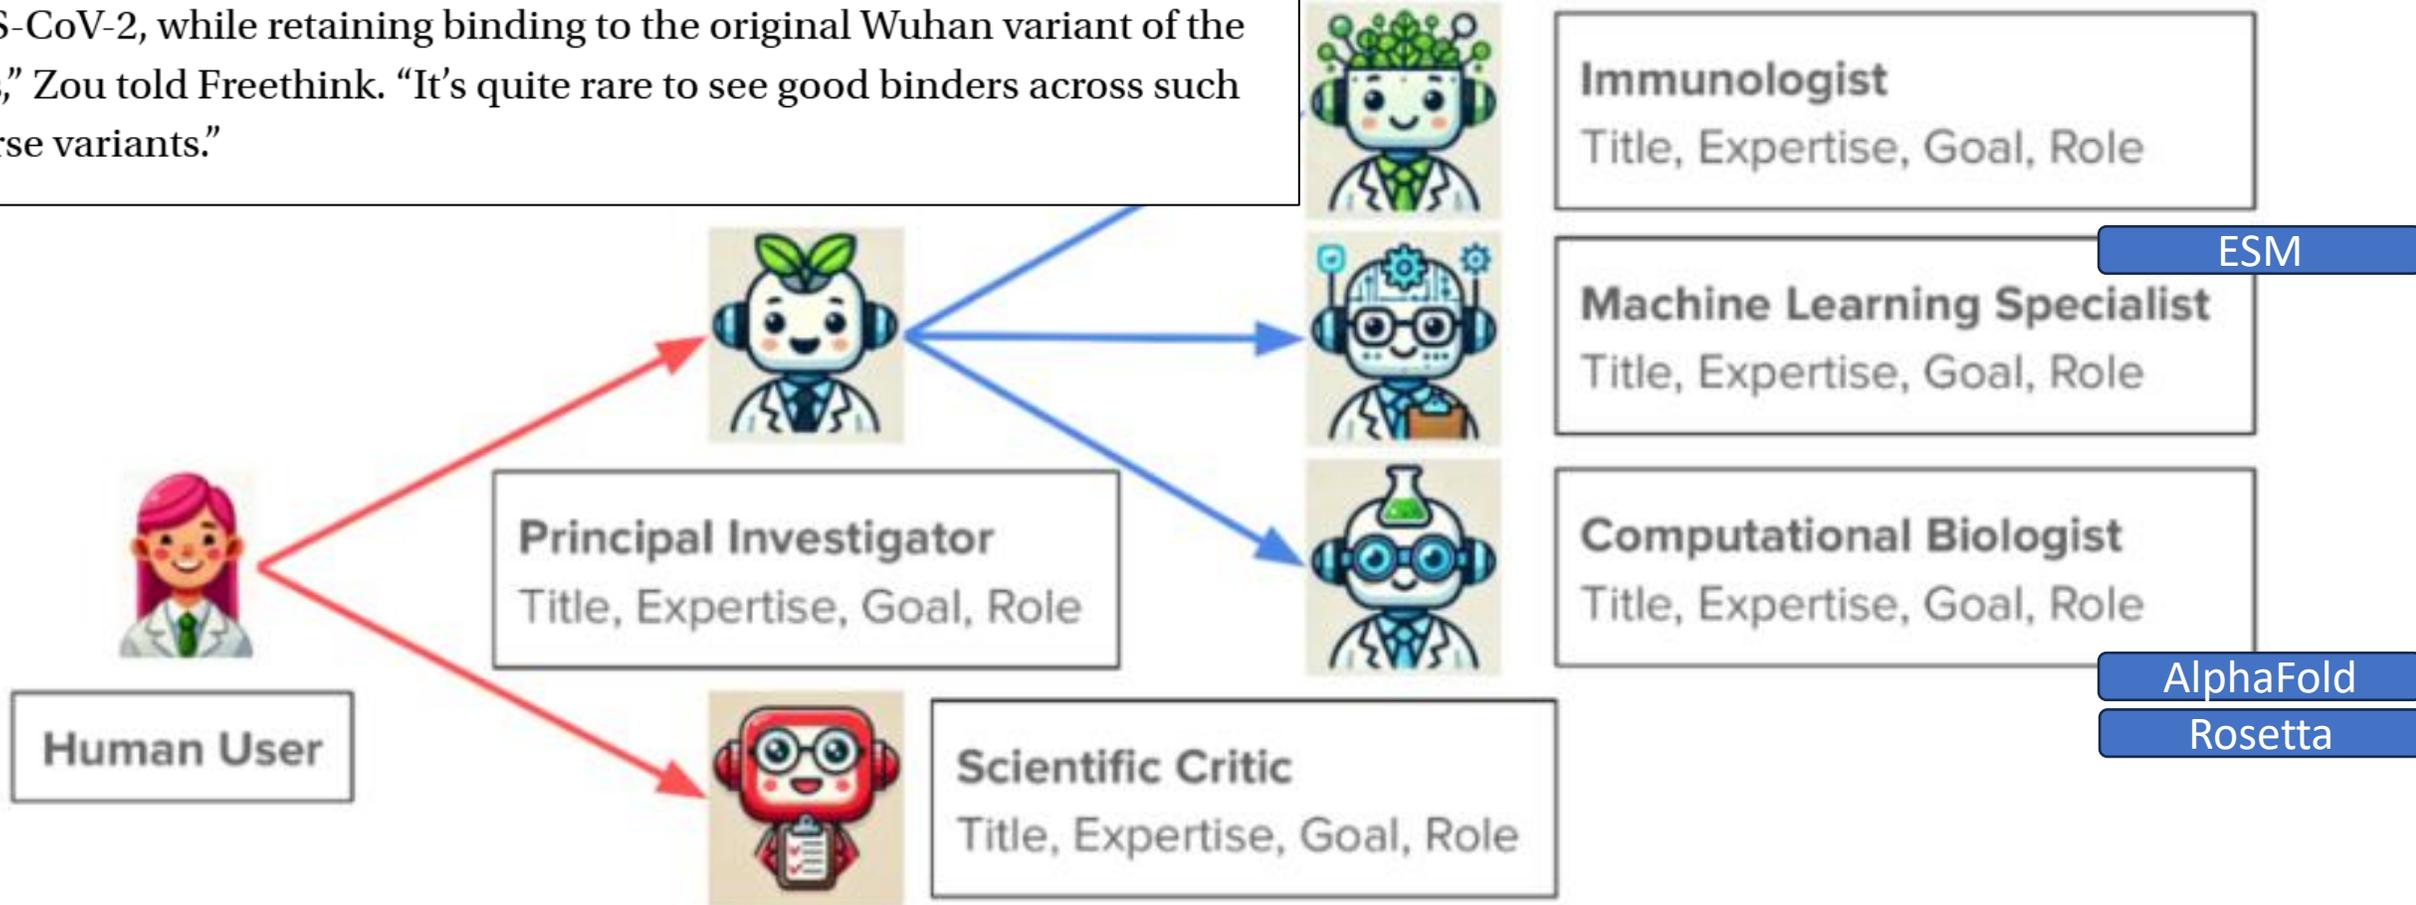
A2A (Agent2Agent - orchestrate agents) – An agent can interact with agents made by others

MCP (Model Context Protocol) Agent can use tools made by others



Virtual Lab example

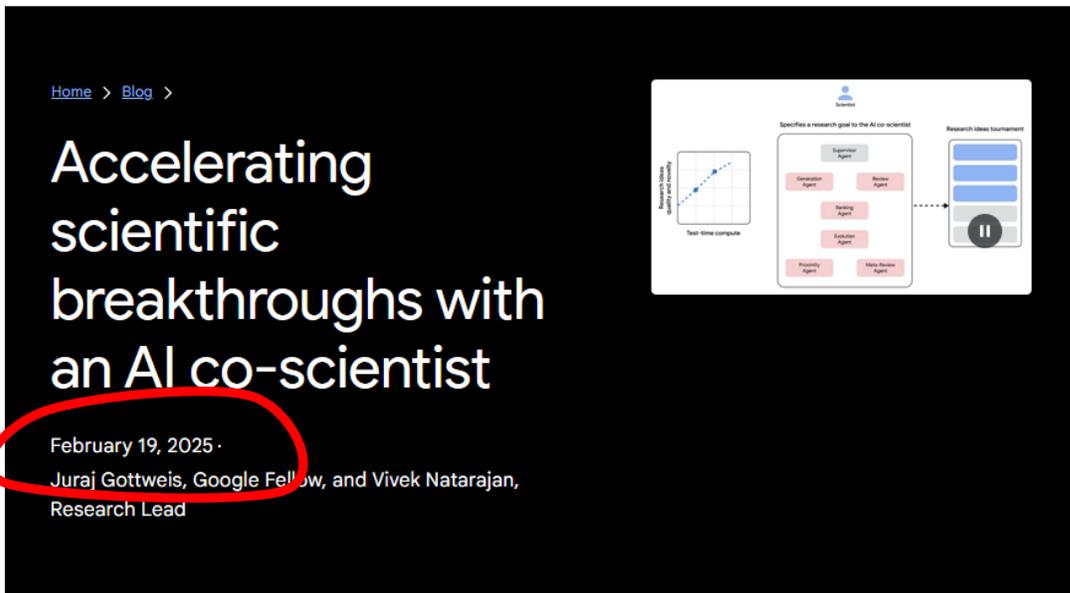
“We are particularly excited that two of the new nanobodies designed by the Virtual Lab show promising binding to the recent JN.1 variant of SARS-CoV-2, while retaining binding to the original Wuhan variant of the virus,” Zou told Freethink. “It’s quite rare to see good binders across such diverse variants.”



<https://www.nature.com/articles/s41586-025-09442-9> (July 2025)

<https://www.freethink.com/artificial-intelligence/virtual-lab-interdisciplinary-research>

Swanson et al. (2024)



- 🔗 AI co-scientist paper
- 🔗 Gene transfer discovery paper
- 🔗 Transfer re-discovery paper
- 🔗 Share

In the pursuit of scientific advances, researchers combine ingenuity and creativity with insight and expertise grounded in literature to generate novel and viable research directions and to guide the exploration that follows. In many fields, this presents a breadth and depth conundrum, since it is challenging to navigate the rapid growth in the rate of scientific publications while integrating insights from unfamiliar domains. Yet overcoming such challenges is critical, as evidenced by the many modern breakthroughs that have emerged from transdisciplinary endeavors. For example, Emmanuelle Charpentier and Jennifer Doudna won the [2020 Nobel Prize in Chemistry](#) for their work on [CRISPR](#), which combined expertise ranging from microbiology to genetics to molecular biology.

Motivated by unmet needs in the modern scientific discovery process and building on [recent AI advances](#), including the ability to synthesize across complex subjects and to perform [long-term planning and reasoning](#), we developed an [AI co-scientist system](#). The AI co-scientist is a multi-agent AI system that is intended to function as a collaborative tool for scientists. Built on [Gemini 2.0](#), [AI co-scientist is designed to mirror the reasoning process underpinning the scientific method. Beyond standard literature review, summarization and “deep research” tools, the AI co-scientist system is intended to uncover new, original knowledge and to formulate demonstrably novel research hypotheses and proposals, building upon prior evidence and tailored to specific research objectives.](#)

We introduce AI co-scientist, a multi-agent AI system built with Gemini 2.0 as a virtual scientific collaborator to help scientists generate novel hypotheses and research proposals, and to accelerate the clock speed of scientific and biomedical discoveries.

QUICK LINKS

- 🔗 AI co-scientist paper

Source: <https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/>

2025-02-18

Towards an AI co-scientist

Juraj Gottweis^{*1}, Wei-Hung Weng^{*2}, Alexander Daryin^{*1}, Tao Tu^{*3},
Anil Palepu², Petar Sirkovic¹, Artiom Myaskovsky¹, Felix Weissenberger¹,
Keran Rong³, Ryutaro Tanno³, Khaled Saab³, Dan Popovici², Jacob Blum⁷, Fan Zhang²,
Katherine Chou², Avinatan Hassidim², Burak Gokturk¹,
Amin Vahdat¹, Pushmeet Kohli³, Yossi Matias²,
Andrew Carroll², Kavita Kulkarni², Nenad Tomasev³,
Vikram Dhillon⁴, Eeshit Dhaval Vaishnav⁵, Byron Lee⁵,
Tiago R D Costa⁶, José R Penadés⁶, Gary Peltz⁷,
Yunhan Xu³, Annalisa Pawlosky¹, Alan Karthikesalingam² and Vivek Natarajan²

¹Google Cloud AI Research, ²Google Research, ³Google DeepMind,
⁴Houston Methodist, ⁵Sequome,
⁶Fleming Initiative and Imperial College London, ⁷Stanford University

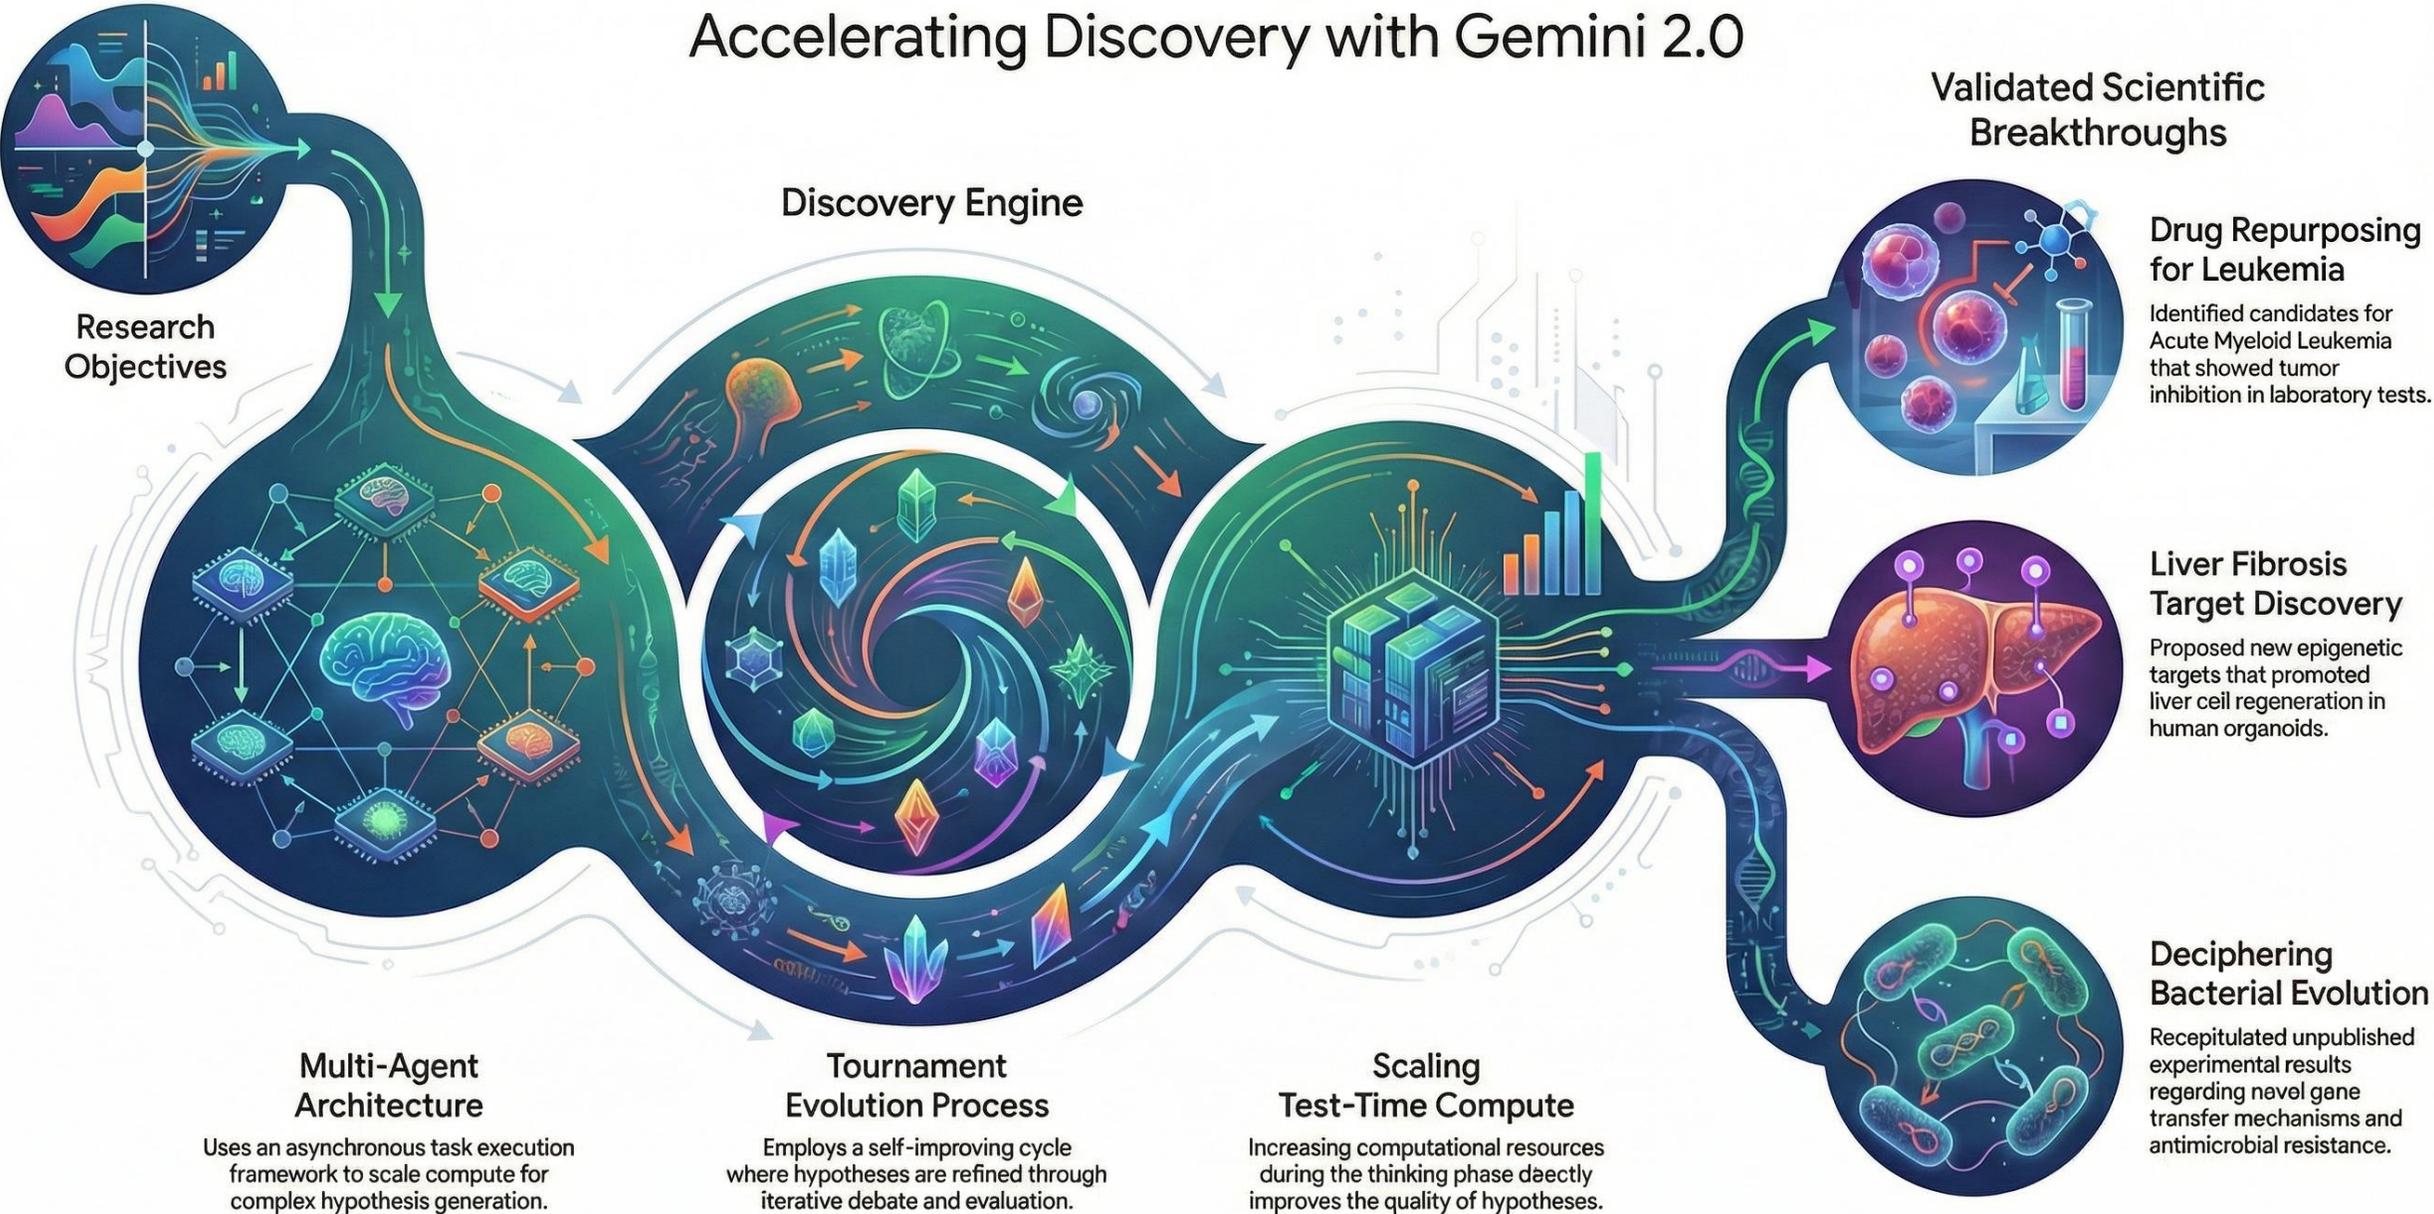
Scientific discovery relies on scientists generating novel hypotheses that undergo rigorous experimental validation. To augment this process, we introduce an AI co-scientist, a multi-agent system built on Gemini 2.0. The AI co-scientist is intended to help uncover new, original knowledge and to formulate demonstrably novel research hypotheses and proposals, building upon prior evidence and aligned to scientist-provided research objectives and guidance. The system's design incorporates a generate, debate, and evolve approach to hypothesis generation, inspired by the scientific method and accelerated by scaling test-time compute. Key contributions include: (1) a multi-agent architecture with an asynchronous task execution framework for flexible compute scaling; (2) a tournament evolution process for self-improving hypotheses generation. Automated evaluations show continued benefits of test-time compute, improving hypothesis quality. While general purpose, we focus development and validation in three biomedical areas: drug repurposing, novel target discovery, and explaining mechanisms of bacterial evolution and anti-microbial resistance. For drug repurposing, the system proposes candidates with promising validation findings, including candidates for acute myeloid leukemia that show tumor inhibition *in vitro* at clinically applicable concentrations. For novel target discovery, the AI co-scientist proposed new epigenetic targets for liver fibrosis, validated by anti-fibrotic activity and liver cell regeneration in human hepatic organoids. Finally, the AI co-scientist recapitulated unpublished experimental results via a parallel *in silico* discovery of a novel gene transfer mechanism in bacterial evolution. These results, detailed in separate, co-timed reports, demonstrate the potential to augment biomedical and scientific discovery and usher an era of AI empowered scientists.

1 Introduction

Human ingenuity and creativity propel the advancement of fundamental research in science and medicine.

<https://arxiv.org/abs/2502.18864>

Towards an AI Co-Scientist: Accelerating Discovery with Gemini 2.0



Research Objectives

Discovery Engine

Validated Scientific Breakthroughs

Drug Repurposing for Leukemia

Identified candidates for Acute Myeloid Leukemia that showed tumor inhibition in laboratory tests.

Liver Fibrosis Target Discovery

Proposed new epigenetic targets that promoted liver cell regeneration in human organoids.

Deciphering Bacterial Evolution

Receptitulated unpublished experimental results regarding novel gene transfer mechanisms and antimicrobial resistance.

Multi-Agent Architecture

Uses an asynchronous task execution framework to scale compute for complex hypothesis generation.

Tournament Evolution Process

Employs a self-improving cycle where hypotheses are refined through iterative debate and evaluation.

Scaling Test-Time Compute

Increasing computational resources during the thinking phase directly improves the quality of hypotheses.

<https://www.futurehouse.org/>

futurehouse.org/about



[About](#)

[Team](#)

[Research](#)

[Contact](#)

[Fellowship](#)

[Platform](#)

[Join Us](#)

MISSION

Our mission is to build an AI scientist, or AI systems, that can automate scientific research and accelerate the pace of discovery, so humanity can find cures for disease, solutions for climate change, and other species-accelerating technologies.

We're building an AI scientist.



01. HUMAN

The Quest

(e.g. figure out how the human brain works; gain the ability to deliver any gene to any cell in the body; and so on)

02. AI SCIENTIST

World Model

Hypothesis
Generation

Experimentation

03. AI SCIENCE
ASSISTANT

Agents for Specific Biological Workflows

(Literature Search, Functional protein annotation,
Designing new proteins, Single cell seq, etc.)

04. AI TOOLS

Predictive
Models

(e.g. AlphaFold)

APIs

Laboratory
Experiments

Healthcare & Life Sciences

Four agentic workflows you can build for life sciences for R&D

November 22, 2025

Pranav Mehrotra

Head of GTM & Partnerships - New Frontiers, Google Cloud

Joe Ledsam

Google Health JAPAC

Try Gemini 3

Our most intelligent model is now available on Vertex AI and Gemini Enterprise

AI agents, powered by generative AI, are rapidly transforming industries by acting as intelligent, collaborative partners that can interpret goals, plan multi-step actions, and work independently across systems, marking a significant shift in how businesses can find, understand, and act on their data. Our [recent blog](#) outlines how AI agents are transforming several industries.

Below we describe how to create a modular, end-to-end platform that accelerates the discovery and





Agenda

1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages



**Our goal in this section:
Use “regular” ChatGPT as an AI-coscientist in an
engineering problem to illustrate potential benefits**

Knowing that: what works (or doesn't) in one example doesn't imply that it will work (or not) in others... it's up to each of us to evaluate "if" and "how" this can be useful in our research tasks



Role

[Link to chat](#)

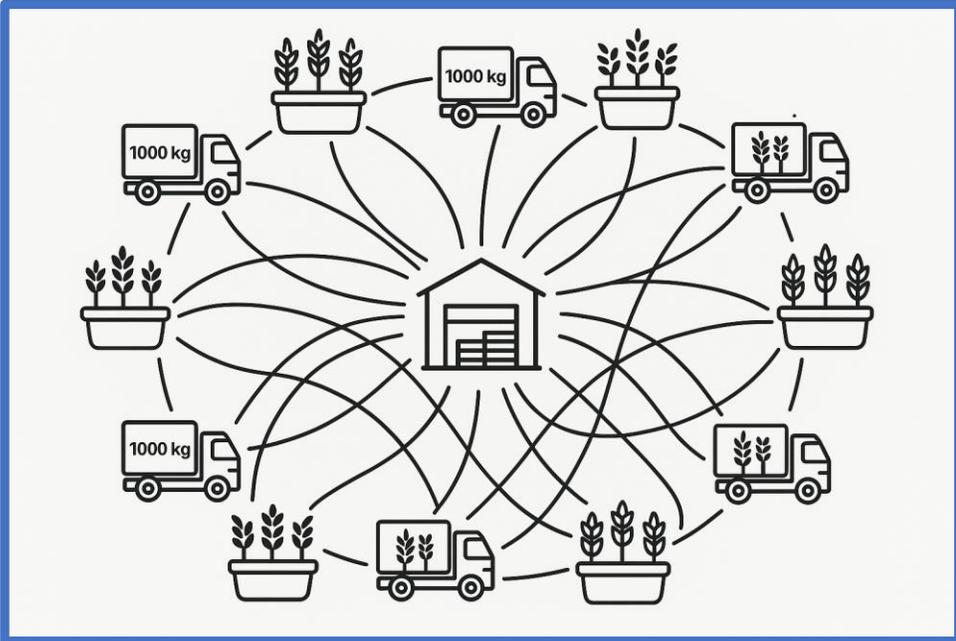
I am a researcher in **Agricultural Economics**, with basic concepts of programming, and basic understanding of heuristic algorithms for optimization problems

I am investigating the structure of transportation costs in the economies of Spanish agricultural distribution companies: impact of truck costs, fuel costs, driver costs, working hours...

I need an algorithm that, given a particular location of farms of a particular company, helps me to estimate their transportation costs. I will use AI for creating such an algorithm, details will be stated later, wait for instructions

I will use it in 100 company examples I have. This will help me to address questions like the average impact of a change in fuel costs in Spain, working hours regulations, etc.

Mechanisms for efficient transport of agricultural products



Algorithm description

- I have trucks with a maximum capacity of 1000 kg (input param). With these, I must collect agricultural products from 30 farms (input param), starting and ending at a central warehouse.
- There is a fixed cost of €300 for using a truck per day (input param). The sections between different locations have associated costs for tolls and fuel, and an estimated travel time (input params). No daily route can exceed 8 hours (input param).
- We should determine how many trucks to use each day, and what routes, so that all the product is collected, minimizing the total cost.
- I should spend no more than 5 minutes in each algorithm run for a particular setup

Information on possible strategies for creating algorithms for this problem

Regarding the problem, indicate what type of algorithms might be useful for planning routes. Consider the 5 minute limit of running time in a conventional computer, and typical number of farms in the order of 30-50. The algorithm and all needed scripts and explanations will be implemented as a Jupyter Notebook for Google Colab

First prototype implementation: understanding the general operating scheme

I like the local search approach: starting from the solution of a fast heuristic, then apply best improving change until no improvement or time limit hit. If no improvement is possible use the extra time starting again from a different solution.

The algorithm, in each iteration, must print relevant information to track its operation, e.g. , the evolution of the current cost (in euros, and with auxiliary information such as cost per route, cost per truck, and cost per overtime hour) and the cost of the best solution. The algorithm must terminate based on time, with execution time being an input parameter.

The code should call the algorithm with a generated example scenario like the following:

- Randomly generate the positions of the 30 farms on a 100 x 100 km grid. The cost of the path between farms should be generated based on the geographical distance, plus a random component of 10%, counting 80 cents per km traveled.
- The time between farms must be proportional to the distance, assuming an average speed of 40 km/h, with a random variation of 10%.
- The quantities of product to be collected must be between 100 kg and 300 kg.
- Print information about the instance: positions of the farms, total amount to be collected, and any other information you consider relevant.

Provide me with the Python code to include in the Google colab

Generate script for test suite

Develop a Python code to perform a battery of tests generating variations in the amount to be collected on each farm.

With each variation, it calls the algorithm for 1 minute.

Compile the information from the battery of tests into a table, including the input parameters used, with statistics such as the total amount to be collected, and information on the results obtained, such as costs and the breakdown into different types, number of trucks used, etc.

In-depth research on a topic

[Link to chat](#)

I would like your help in compiling a state-of-the-art overview of the different types of trucks available for transporting refrigerated vegetables. I am interested in trucks with a maximum payload ranging from 500 kg to 2000 kg. I need information on the most common models and their characteristics, such as price, dimensions, and maximum load capacity. I would also appreciate information that will help me estimate the fuel costs per kilometer, assuming a full load, for each model.

I'm interested in the situation in Spain. I'm interested in new trucks, any type of engine; I'm only interested in fuel consumption, and for electric trucks, the energy cost. For electric trucks, please include the battery range in kilometers and the charging time. There are no restrictions on the type of cooling system.

Agenda

1. General purpose AI for lowering friction
2. Domain-Specific models: a game changer
3. AI co-scientist agents, the next step
4. A self-made AI co-scientist use case example
5. Final (and I hope that useful) messages



- AI is “superpowers” for researchers. Use it to discover more, understand deeper, heal faster, design smarter...
- A great time to be a researcher: your value to society is x100... “research” is the long-term *killer application* for AI :-)

nature medicine

Explore content ▾ About the journal ▾ Publish with us ▾

[nature](#) > [nature medicine](#) > [articles](#) > article

Article | [Open access](#) | Published: 03 June 2025

A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial

[Zuojun Xu](#) , [Feng Ren](#), [Ping Wang](#), [Jie Cao](#), [Chunting Tan](#), [Dedong Ma](#), [Li Zhao](#), [Jinghong Dai](#), [Yipeng Ding](#), [Haohui Fang](#), [Huiping Li](#), [Hong Liu](#), [Fengming Luo](#), [Ying Meng](#), [Pinhua Pan](#), [Pingchao Xiang](#), [Zuke Xiao](#), [Sujata Rao](#), [Carol Satler](#), [Sang Liu](#), [Yuan Lv](#), [Heng Zhao](#), [Shan Chen](#), [Hui Cui](#), ... [Alex Zhavoronkov](#) 

+ Show authors

Nature Medicine **31**, 2602–2610 (2025) | [Cite this article](#)

46k Accesses | 8 Citations | 328 Altmetric | [Metrics](#)

Abstract

Despite substantial progress in artificial intelligence (AI) for generative chemistry, few novel AI-discovered or AI-designed drugs have reached human clinical trials. Here we present the results of the first phase 2a multicenter, double-blind, randomized, placebo-controlled trial testing the safety and efficacy of **rentosertib** (formerly ISM001-055), a first-in-class AI-generated small-molecule inhibitor of TNIK, a first-in-class target in idiopathic pulmonary fibrosis (IPF) discovered using generative AI. IPF is an age-related progressive lung condition with no current therapies available that reverse the degenerative course of disease. Patients

<https://www.nature.com/articles/s41591-025-03743-2>

AI-discovered medicines

MIT News

ON CAMPUS AND AROUND THE WORLD

 [SUBSCRIBE](#)

Using generative AI, researchers design compounds that can kill drug-resistant bacteria

The team used two different AI approaches to design novel antibiotics, including one that showed promise against MRSA.

Anne Trifon | MIT News
August 14, 2025



<https://news.mit.edu/2025/using-generative-ai-researchers-design-compounds-kill-drug-resistant-bacteria-0814>



Ekin Dogus Cubuk, left, and Liam Fedus are the co-founders of the start-up Periodic Labs in San Francisco. Jason Henry for The New York Times

Founded by a co-creator of ChatGPT, Periodic Labs aims to build artificial intelligence that can accelerate discoveries in physics, chemistry and other fields.

<https://www.nytimes.com/2025/09/30/technology/ai-meta-google-openai-periodic.html>

<https://periodic.com/>

- AI is “superpowers” for researchers. Use it to discover more, understand deeper, heal faster, design smarter...
- A great time to be a researcher: your value to society is x100... “research” is the long-term *killer application* for AI :-)
- Academic rigour, which implies ethical use of AI is the **ONLY** way
- Basic programming skills open A LOT of doors now... invest time in that if you are not already there, it is worthy
- Growing importance of **deeply learning the fundamentals**, how to write, think, reason and understand your discipline core concepts. Or else, you will not be able to exploit AI



Universidad
Politécnica
de Cartagena

MEMBER OF



EUROPEAN
UNIVERSITY OF
TECHNOLOGY

Thanks!

Pablo Pavón Mariño
Vicerrector de Estrategia Digital e Inteligencia Artificial
Universidad Politécnica de Cartagena
pablo.pavon@upct.es