

10,000 Einsteins: AI and the future of theoretical physics

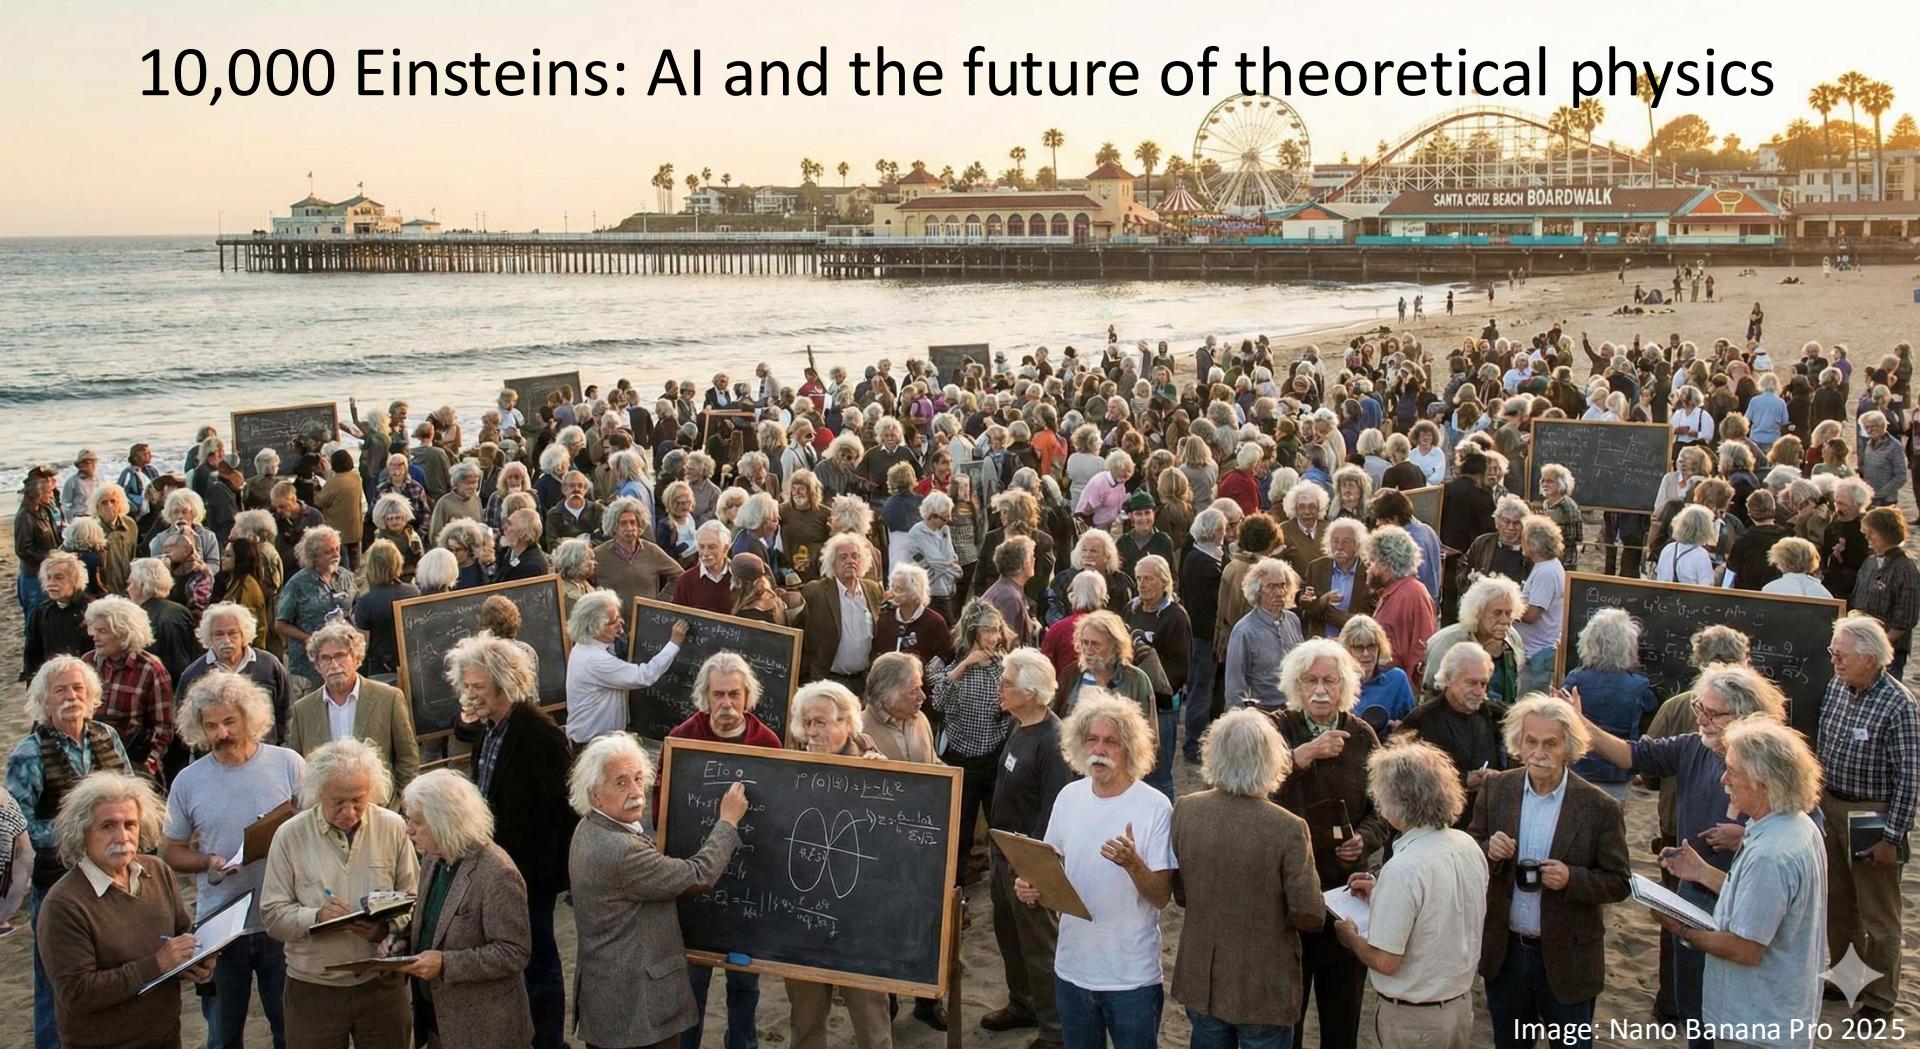


Image: Nano Banana Pro 2025

UCSC
Physics Colloquium
Dec 4, 2025

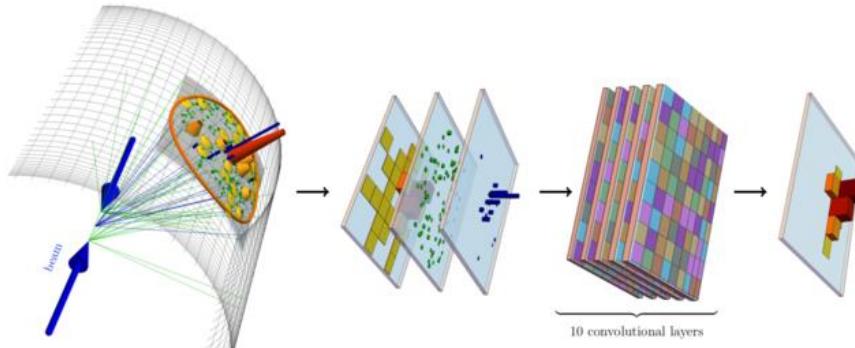
Matthew Schwartz
Harvard University
Institute for Artificial Intelligence
and Fundamental Interactions



Outline

Part 1: The Past

- Machine learning as data science



Part 2: The Present

- Augmented intelligence



Part 3: The Future

- Vibe Physics



1. The Past

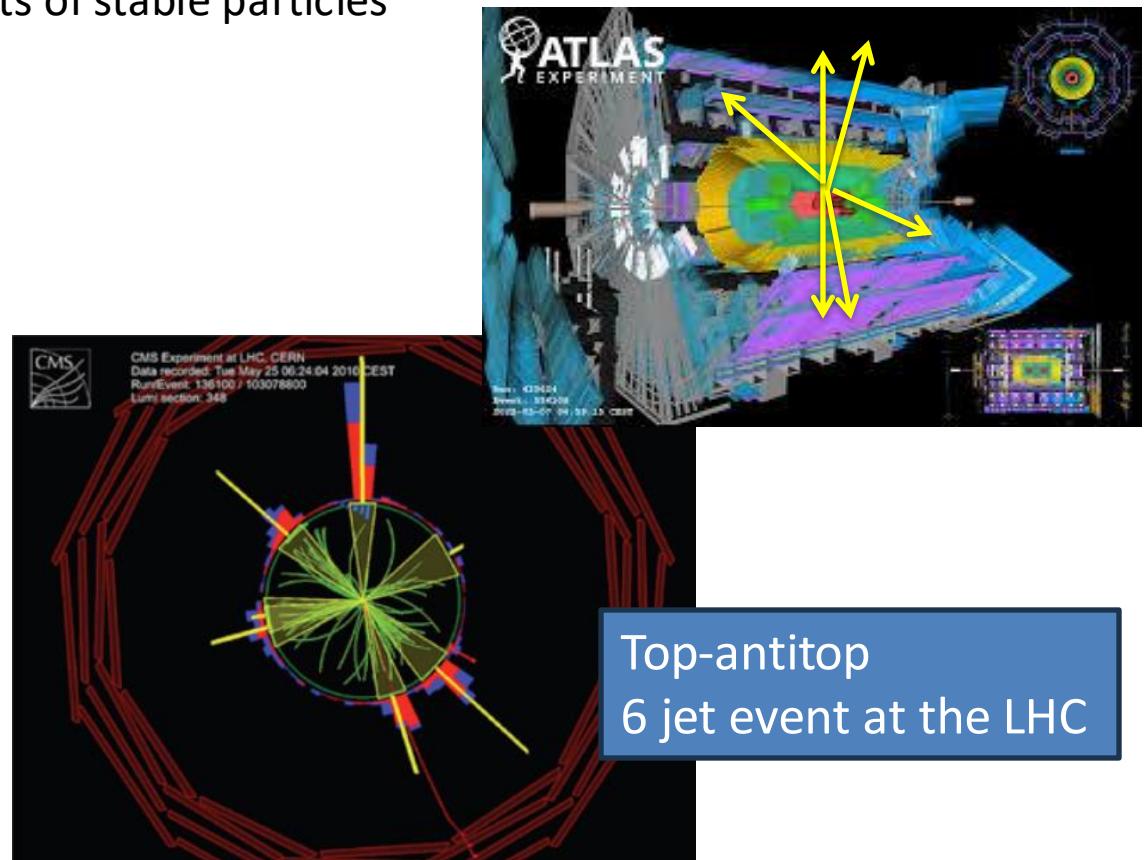
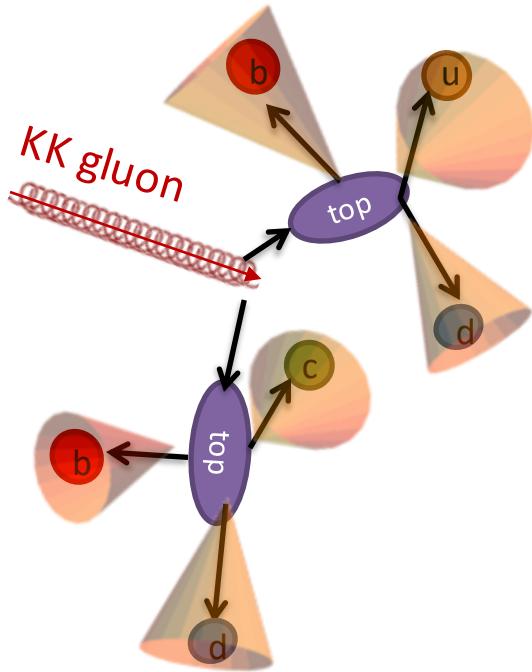


Image: Nano Banana Pro 2025

ML in particle physics

String theory predicts a new very heavy particle called the **KK gluon**

- KK gluons are unstable: they decay to two **top quarks**
 - top quarks decay to other quarks
 - quarks radiate into jets of stable particles



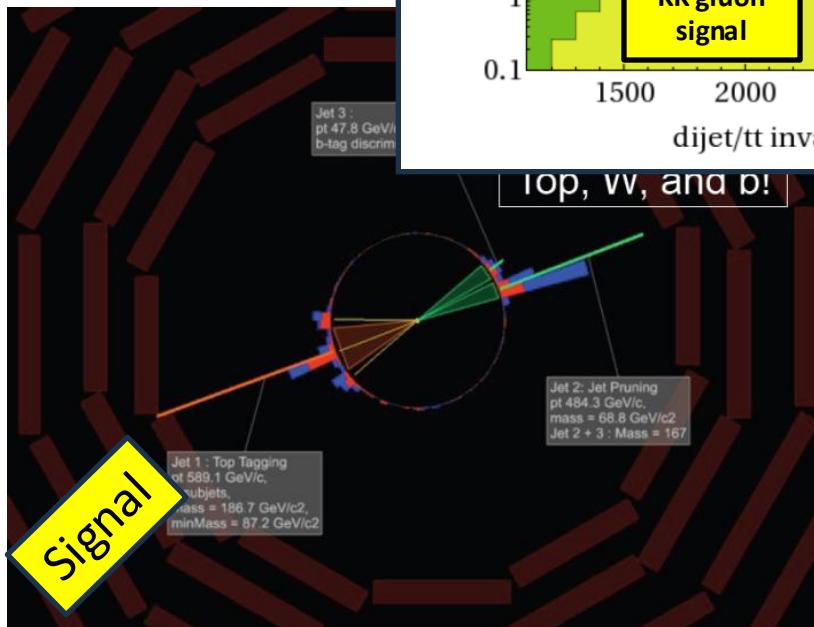
Boosted tops

For TeV mass KK gluons, top quarks are **ultrarelativistic** $m = 1\text{TeV}$

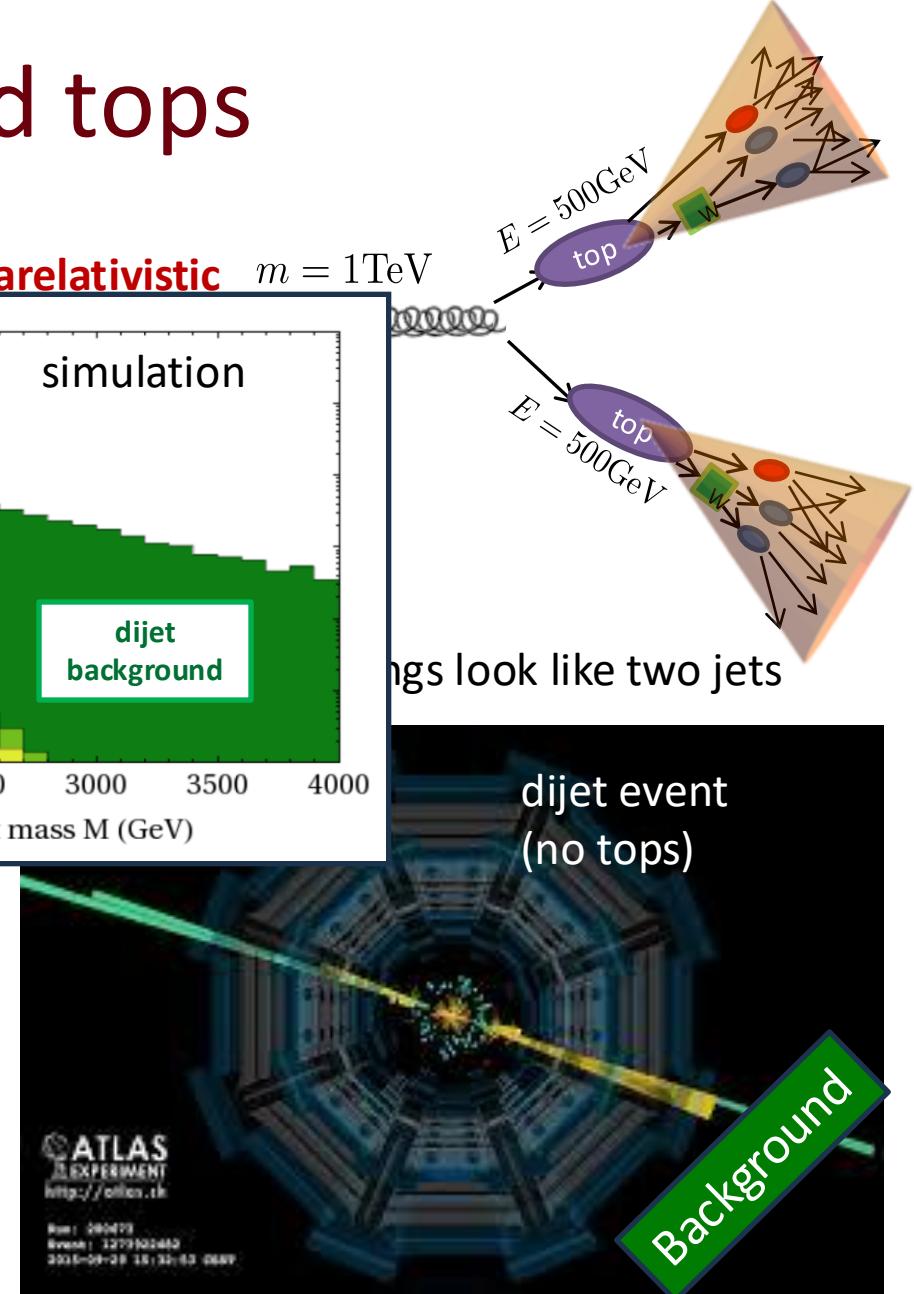
- decay prod



Now looks like 2



predicted: 5 events/year



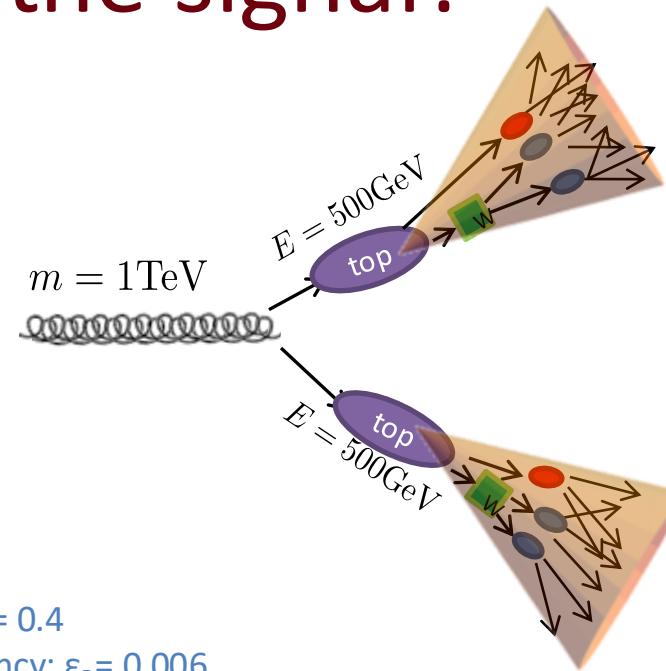
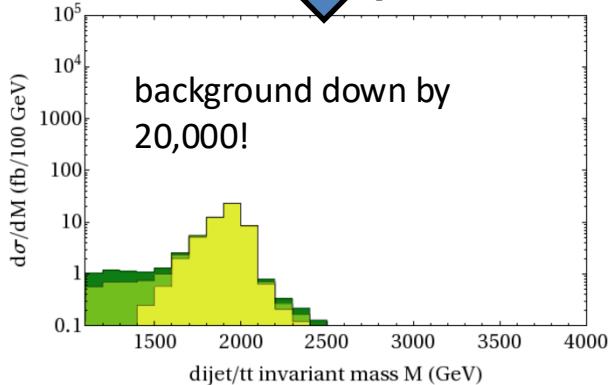
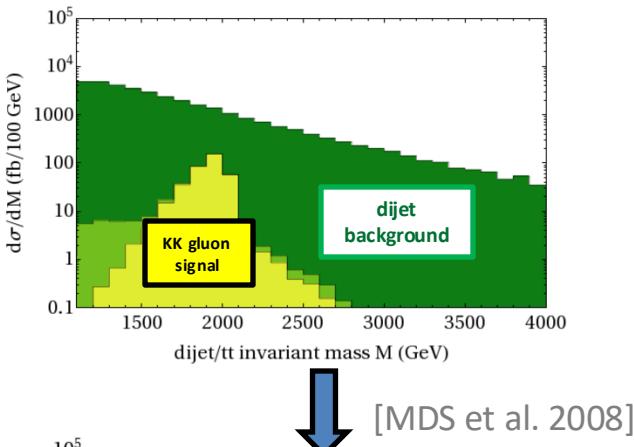
predicted 10,000 events/year

How can we find the signal?

traditional approach (2008-2018)

Think about the physics

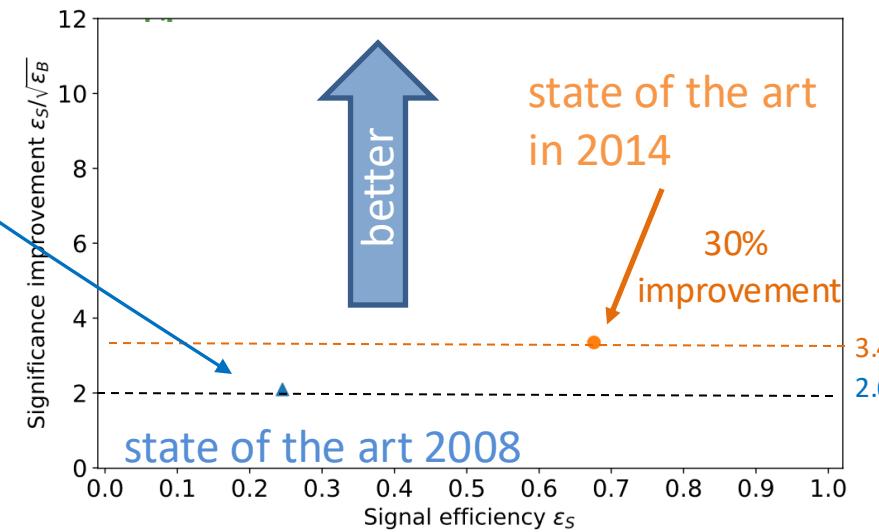
- Top quarks are different from dijets
 - look for quarks inside jets
 - look for mass scale of top quark
 - loop for W boson
 - ...



signal efficiency: $\varepsilon = 0.4$

background efficiency: $\varepsilon_B = 0.006$

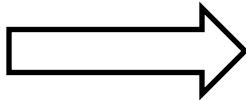
significance improvement: $\varepsilon_S / \sqrt{\varepsilon_B} = 2.6$



Machine learning (ML) approach

Take some tool highly engineered for another purpose and **shoehorn it into physics**

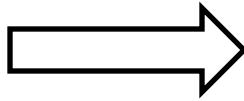
vase designed for flowers



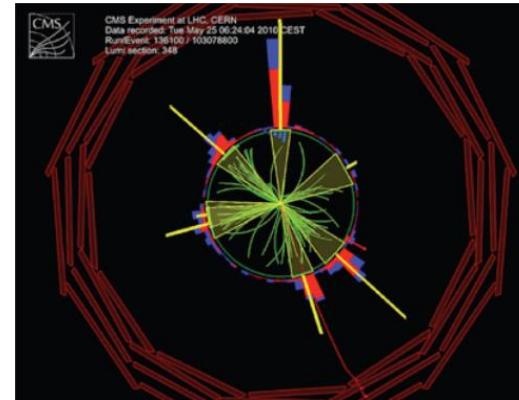
cat uses as a house



Machine learning for facial recognition



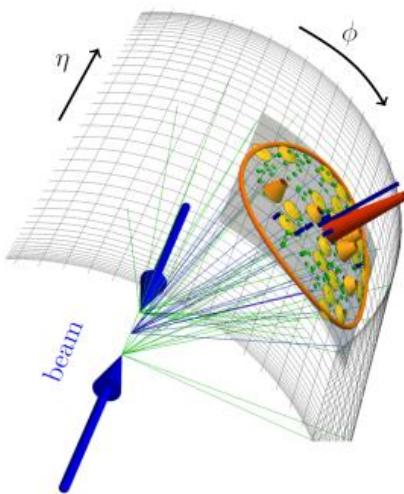
used for particle identification in physics



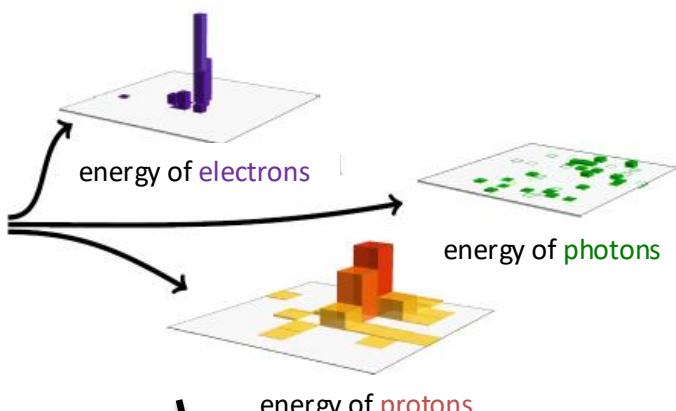
Jet images

Shoehorn LHC data into image format [Cogan et al. 2014]

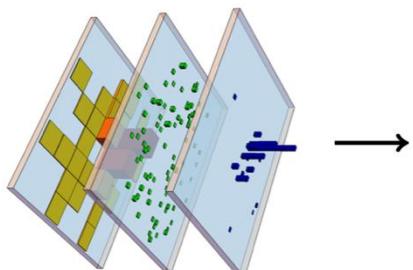
1. experiments measure energy in a detector



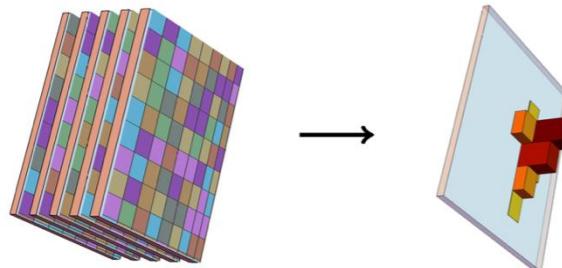
2. Separate energies into classes



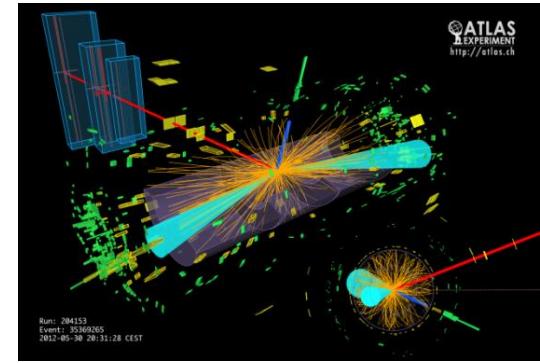
3. map to a color image



4. Input to convolutional neural network designed for facial recognition

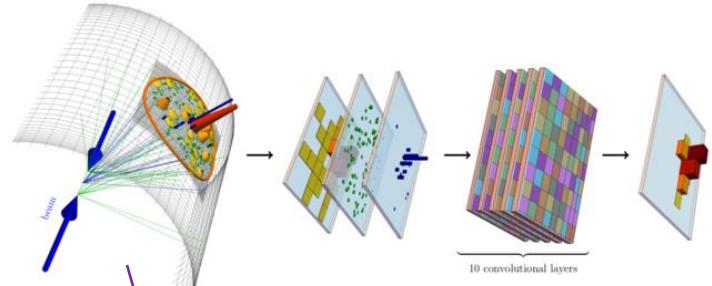


[MDS et al. 2016]

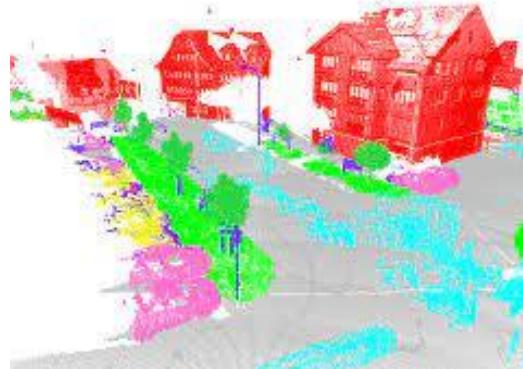


Results for ML

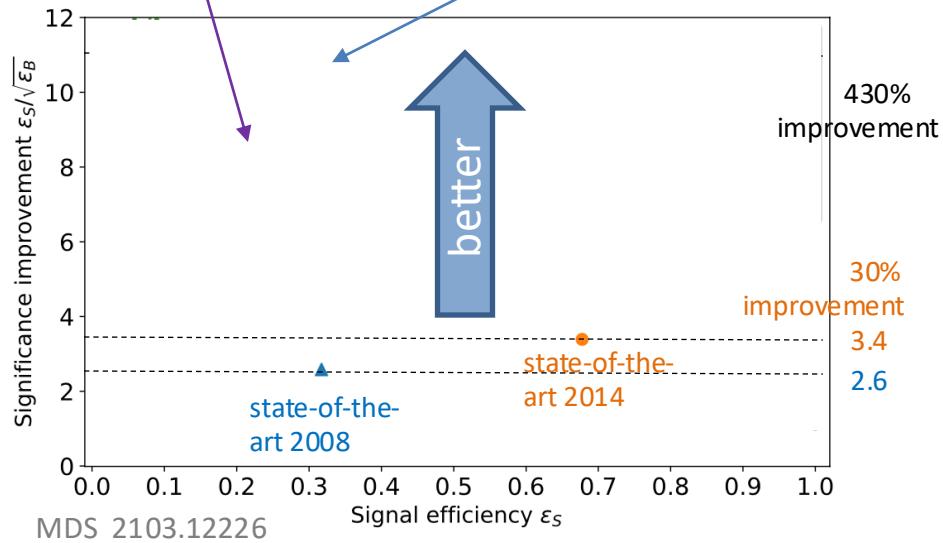
Convolutional networks for facial recognition



Point clouds: developed for self-driving cars



Machine learning methods are much better



Machine learning works amazingly well!

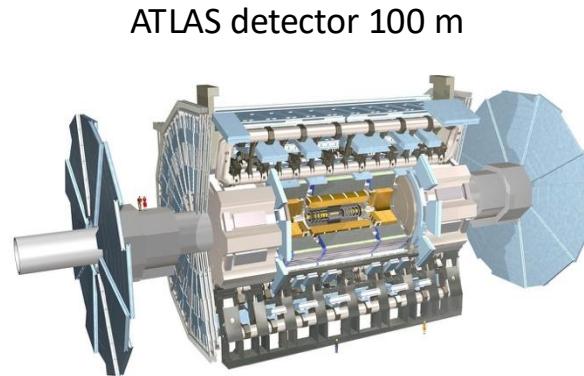
- Blows traditional approaches away!
- ML requires less “thinking”
- Better performance
- Provides less physical insight

Particle physics is perfect for ML

- We understand physics over **24 orders of magnitude!**
- Can model and simulate details of signal and background to amazing accuracy
 - Unparalleled environment to train artificial neural networks



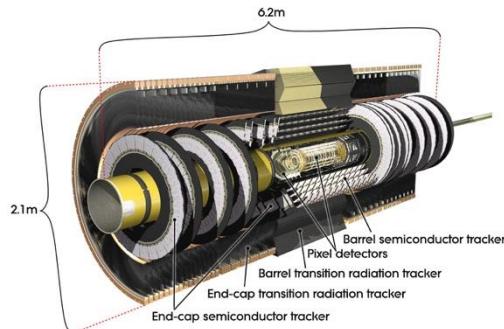
Large Hadron Collider $25\text{km} = 10^4\text{ m}$



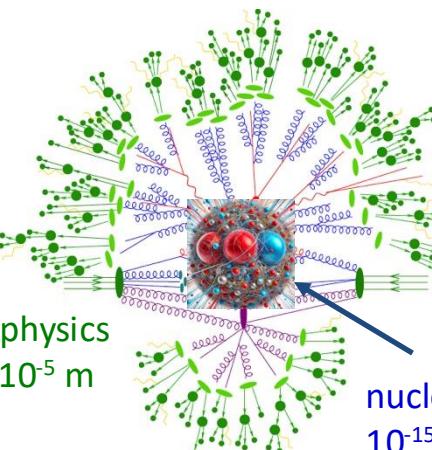
ATLAS detector 100 m



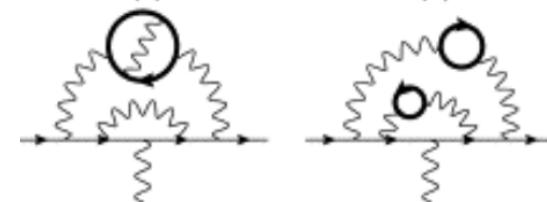
calorimeters $0.1\text{ m} - 10\text{m}$



tracker $10^{-5}\text{ m} - 10^{-1}\text{ m}$



nuclear physics
 $10^{-15}\text{ m} - 10^{-10}\text{ m}$



quantum field theory
 $10^{-20}\text{ m} - 10^{-15}\text{ m}$

↑
size of a top quark

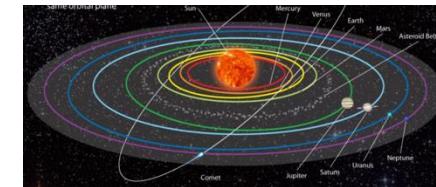
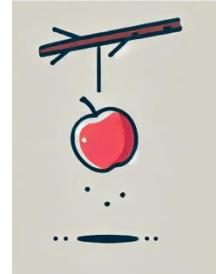
What about theoretical physics?

Many advances in theoretical physics come from simplification/unification



Newton: motion of apples and planets governed by the same simple law

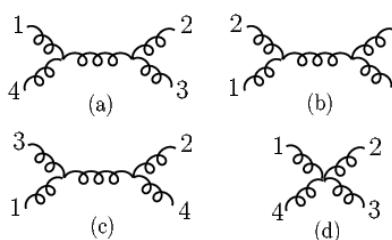
$$F = G \frac{m_1 m_2}{r^2}$$



Can machine learning learn to simplify?

calculate some Feynman diagrams for
4 gluon scattering

get a messy answer



$$\begin{aligned} (a) &= -\frac{\langle 12 \rangle \langle 13 \rangle \langle 24 \rangle [13] [24]}{\langle 23 \rangle \langle 34 \rangle^2 [12] [23]} + \frac{\langle 12 \rangle \langle 13 \rangle \langle 24 \rangle [14]}{\langle 23 \rangle \langle 34 \rangle^2 [12]} - \frac{\langle 12 \rangle \langle 24 \rangle [13] [24]^2}{\langle 34 \rangle^2 [12] [14] [23]} + \frac{\langle 12 \rangle \langle 24 \rangle [24]}{\langle 34 \rangle^2 [12]} + \frac{\langle 12 \rangle [13] [24] [34]}{\langle 34 \rangle [12] [14] [23]} \\ (b) &= -\frac{\langle 12 \rangle [34]}{\langle 34 \rangle [12]} - \frac{\langle 13 \rangle \langle 14 \rangle [13] [34]}{\langle 34 \rangle^2 [12] [23]} - \frac{\langle 13 \rangle \langle 24 \rangle [13] [24] [34]}{\langle 34 \rangle^2 [12] [14] [23]} + \frac{\langle 13 \rangle \langle 24 \rangle [34]}{\langle 34 \rangle^2 [12]} - \frac{\langle 13 \rangle \langle 24 \rangle [13] [24]}{\langle 34 \rangle^2 [12]^2} + \frac{\langle 13 \rangle \langle 24 \rangle [14] [23]}{\langle 34 \rangle^2 [12]^2} \\ (c) &= +\frac{\langle 13 \rangle [13] [34]^2}{\langle 34 \rangle [12] [14] [23]} - \frac{\langle 14 \rangle \langle 23 \rangle [34]}{\langle 34 \rangle^2 [12]} - \frac{\langle 23 \rangle \langle 24 \rangle [13] [24]^2}{\langle 34 \rangle^2 [12]^2 [14]} + \frac{\langle 23 \rangle \langle 24 \rangle [23] [24]}{\langle 34 \rangle^2 [12]^2} + \frac{\langle 23 \rangle [13] [24] [34]}{\langle 34 \rangle [12]^2 [14]} - \frac{\langle 23 \rangle [23] [34]}{\langle 34 \rangle [12]^2} \\ (d) &= \end{aligned}$$

$$\text{simplify!} = -\frac{\langle 12 \rangle [34]^2}{\langle 34 \rangle [14] [23]}$$

- simpler form suggests **deeper structure**
- is there a **better way** to do the calculation?
 - In this case, yes! (BCFW recursion)

Machine learning approach

- Generate training data by “scrambling”

Input:
$$\frac{-(34)^2[12]^3[34][35] + (34)^2[12]^2[13][23][45] - (34)^2[12]^2[14][23][35]}{\langle 23 \rangle \langle 34 \rangle [13]^2[34][35] + \langle 23 \rangle \langle 45 \rangle [13]^2[35][45] - \langle 23 \rangle \langle 45 \rangle [13][14][35]^2}$$

$[15] \rightarrow \frac{[14][35] - [13][45]}{[34]}$ Schouten

$\frac{-(34)^2[12]^3[35] - (34)^2[12]^2[15][23]}{\langle 23 \rangle \langle 34 \rangle [13]^2[35] - \langle 23 \rangle \langle 45 \rangle [13][15][35]}$

$[25] \rightarrow \frac{[12][35] + [15][23]}{[13]}$ Schouten

$\langle 24 \rangle \rightarrow \frac{\langle 45 \rangle [15] - \langle 34 \rangle [13]}{[12]}$ Momentum conservation

$\frac{\langle 34 \rangle^2[12]^2[25]}{\langle 23 \rangle \langle 34 \rangle [13][35] - \langle 23 \rangle \langle 45 \rangle [15][35]}$

$\frac{\langle 34 \rangle^2[12][25]}{\langle 23 \rangle \langle 24 \rangle [35]}$ Desired output

[MDS et al. 2024]

Learning the Simplicity of Scattering Amplitudes

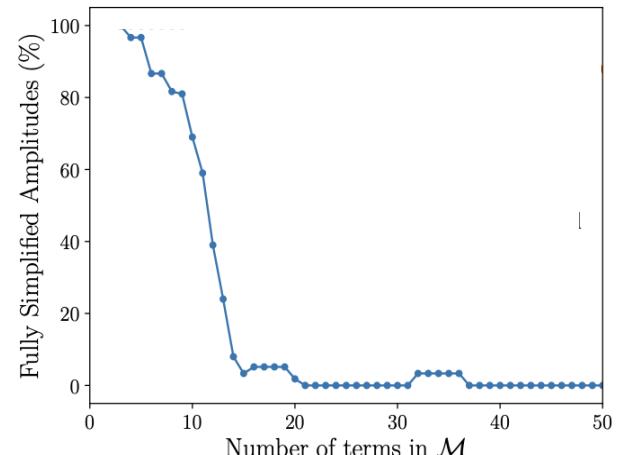
Clifford Cheung (Caltech), Aurélien Dersy (Harvard U. and IAIPI, Cambridge),

Matthew D. Schwartz (Harvard U. and IAIPI, Cambridge) (Aug 8, 2024)

e-Print: 2408.04720 [hep-th]

train
neural network
to unscramble

It learns to simplify short expressions



Machine learning approach

- Generate training data by “scrambling”

Input:
$$\frac{-(34)^2[12]^3[34][35] + (34)^2[12]^2[13][23][45] - (34)^2[12]^2[14][23][35]}{(23)\langle 34 \rangle [13]^2[34][35] + \langle 23 \rangle \langle 45 \rangle [13]^2[35][45] - \langle 23 \rangle \langle 45 \rangle [13][14][35]^2}$$

$[15] \rightarrow \frac{[14][35] - [13][45]}{[34]}$ Schouten

$[25] \rightarrow \frac{[12][35] + [15][23]}{[13]}$ Schouten

$\langle 24 \rangle \rightarrow \frac{\langle 45 \rangle [15] - \langle 34 \rangle [13]}{[12]}$ Momentum conservation

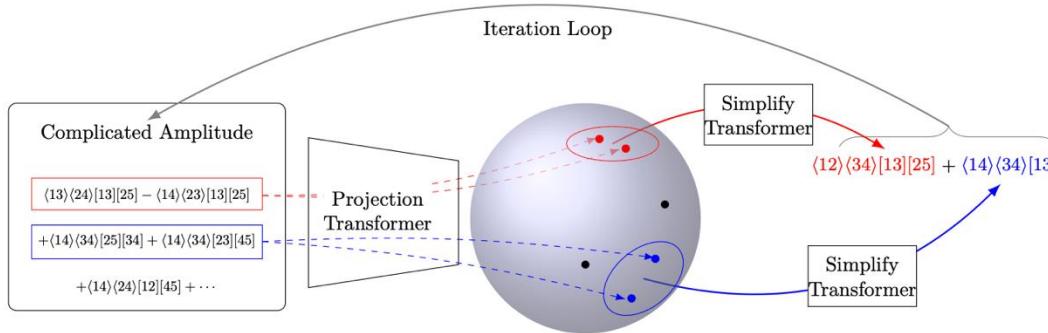
Desired output: $\frac{(34)^2[12]^2[25]}{(23)\langle 34 \rangle [13][35] - \langle 23 \rangle \langle 45 \rangle [15][35]}$

Learning the Simplicity of Scattering Amplitudes

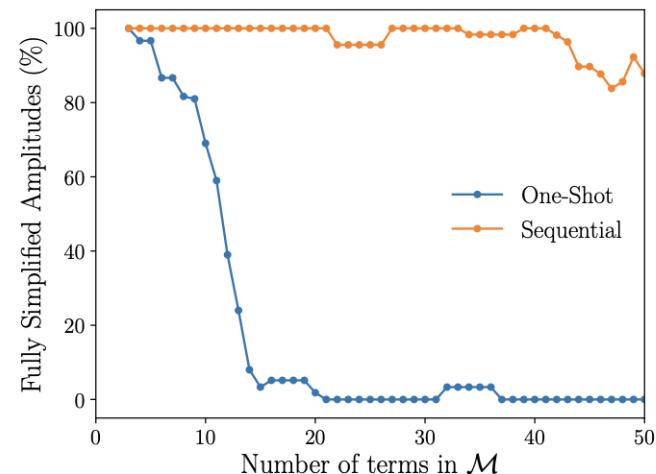
Clifford Cheung (Caltech), Aurélien Dersy (Harvard U. and IAIFl, Cambridge),
Matthew D. Schwartz (Harvard U. and IAIFl, Cambridge) (Aug 8, 2024)
e-Print: 2408.04720 [hep-th]

train
neural network
to unscramble

Supplement with contrastive learning
(adapted from computer vision)



It learns to simplify short expressions
and long expressions



Example application: graviton-scalar scattering

Iteration Loop

Complicated Amplitude

Projection Transformer

Tracer

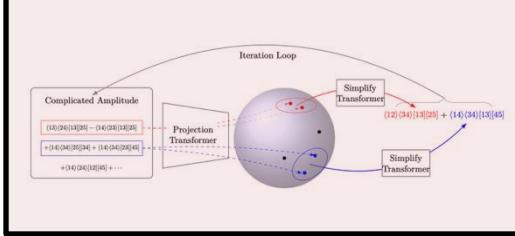
Iteration Loop

Learning the Simplicity of Scattering Amplitudes

Clifford Cheung (Caltech), Aurélien Dersy (Harvard U. and IAIFI, Cambridge),
Matthew D. Schwartz (Harvard U. and IAIFI, Cambridge) (Aug 8, 2024)

e-Print: 2408.04720 [hep-th]

298 terms



simplifies to 2 terms

$$\overline{\mathcal{M}} = \frac{\langle 12 \rangle \langle 13 \rangle \langle 23 \rangle}{\langle 24 \rangle \langle 25 \rangle \langle 45 \rangle} \left(\frac{[14][35]}{\langle 14 \rangle \langle 35 \rangle} - \frac{[15][34]}{\langle 15 \rangle \langle 34 \rangle} \right)$$

- No known better way to do calculation
 - Can we (or ML) find deeper structure?

2. The present: generative AI



Image: Nano Banana Pro 2025
Animation: Runway

Current AI models can do *all* college-level physics

Problem from my Spring 2025 Statistical Mechanics final

Lithium-Ion batteries work by having lithium in both the anode and the cathode with a electrolytic barrier of around $\Delta\epsilon = 0.55\text{eV}$ in between. The lithium can discharge when the circuit is closed and also thermally over the barrier. If a battery lasts for 20 hours at 10°C estimate how long it lasts in the sun at 40°C .

To estimate how long a lithium-ion battery will last at a higher temperature (40°C) compared to a lower temperature (10°C), assuming **thermal discharge over a barrier of $\Delta\epsilon = 0.55\text{ eV}$** , we can use the **Arrhenius equation**:

$$\text{Rate} \propto e^{-\Delta\epsilon/kT}$$

Where:

- $\Delta\epsilon = 0.55\text{ eV}$
- $k = 8.617 \times 10^{-5}\text{ eV/K}$ (Boltzmann constant)
- T is the absolute temperature in kelvin

Let:

- $T_1 = 10^\circ\text{C} = 283.15\text{ K}$
- $T_2 = 40^\circ\text{C} = 313.15\text{ K}$

Let R_1 and R_2 be the rates of discharge at T_1 and T_2 . Then:

$$\frac{R_2}{R_1} = \frac{e^{-\Delta\epsilon/kT_2}}{e^{-\Delta\epsilon/kT_1}} = e^{\Delta\epsilon(\frac{1}{T_1} - \frac{1}{T_2})}$$

Calculate the exponent:

$$\frac{1}{kT_1} - \frac{1}{kT_2} = \frac{1}{8.617 \times 10^{-5} \times 283.15} - \frac{1}{8.617 \times 10^{-5} \times 313.15} \approx 41.07 - 36.91 =$$

Then:

$$\frac{R_2}{R_1} = e^{0.55 \times 4.16} \approx e^{2.29} \approx 9.87$$

So the **rate** of thermal discharge increases by a factor of $e^{2.29}$, meaning the **battery life shortens by the same factor**.

If the battery lasts 20 hours at 10°C :

$$\text{Battery life at } 40^\circ\text{C} = \frac{20}{9.87} \approx 2.03 \text{ hours}$$

 **Final Answer:**

The battery would last approximately **2 hours** in the sun at 40°C .

Correct!

Not only does it get the right answer, but it uses the right equations and shows its work

Nano-banana pro (Nov 2025) is scary:

Physics 222, Modern Physics, Exam 2

NAME _____

You are graded on your work, with partial credit where it is deserved.

Please be clear and well-organized in all your steps.

1. (15) Show that

$$\Psi(x, t) = e^{-ipx/\hbar} e^{-iEt/\hbar}$$

satisfies the time-dependent Schrödinger equation for a free particle

$$i\hbar \frac{\partial \Psi(x, t)}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi(x, t)}{\partial x^2}$$

and obtain the energy E in terms of the momentum p .



Nano-banana

Physics 222, Modern Physics, Exam 2

NAME Alex Smith

You are graded on your work, with partial credit where it is deserved.

Please be clear and well-organized in all your steps.

1. (15) Show that

$$\Psi(x, t) = e^{-ipx/\hbar} e^{-iEt/\hbar} = e^{-i(px+Et)/\hbar}$$

satisfies the time-dependent Schrödinger equation for a free particle

$$i\hbar \frac{\partial \Psi(x, t)}{\partial t} = -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi(x, t)}{\partial x^2}$$

and obtain the energy E in terms of the momentum p .

$$\begin{aligned} \text{LHS: } i\hbar \frac{\partial \Psi}{\partial t} &= i\hbar \frac{\partial}{\partial t} \left[e^{-i(px+Et)/\hbar} \right] \\ &= i\hbar \left(-i \frac{E}{\hbar} \right) e^{-i(px+Et)/\hbar} \\ &= -i^2 E \Psi = E \Psi \quad \checkmark \end{aligned}$$

$$\begin{aligned} \text{RHS: } -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi}{\partial x^2} &= \frac{\partial}{\partial x} \left[e^{-i(px+Et)/\hbar} \right] = -\frac{ip}{\hbar} \Psi \\ \frac{\partial^2 \Psi}{\partial x^2} &= \frac{\partial}{\partial x} \left[\left(-\frac{ip}{\hbar} \right) \Psi \right] = -\frac{ip}{\hbar} \frac{\partial \Psi}{\partial x} \\ &= \left(-\frac{ip}{\hbar} \right) \left(-\frac{ip}{\hbar} \right) \Psi = -\frac{p^2}{\hbar^2} \Psi \end{aligned}$$

$$\begin{aligned} \frac{\partial^2 \Psi}{\partial x^2} &= \frac{\partial}{\partial x} \left[\left(-\frac{ip}{\hbar} \right) \Psi \right] = \left(-\frac{ip}{\hbar} \right) \frac{\partial \Psi}{\partial x} \Psi \\ \text{So, } -\frac{\hbar^2}{2m} \frac{\partial^2 \Psi}{\partial x^2} &= -\frac{\hbar^2}{2m} \left(-\frac{p^2}{\hbar^2} \right) \Psi \\ &= \frac{p^2}{2m} \Psi \quad \checkmark \end{aligned}$$

Set LHS = RHS:

$$E \Psi = \frac{p^2}{2m} \Psi$$

Since $\Psi \neq 0$,

$$\boxed{E = \frac{p^2}{2m}} \quad \text{Q.E.D.}$$

(b) (5) Is this wave traveling in the positive or the negative x direction? Explain.

The wave function is $\Psi(x, t) = e^{-i(px+Et)/\hbar}$.

The phase is $\phi = -\frac{(px+Et)}{\hbar}$.

For a point of constant phase, $\frac{d\phi}{dt} = 0$.

$$\frac{d}{dt} (px + Et) = 0 \Rightarrow p \left(\frac{dx}{dt} \right) + E = 0$$

$$\text{velocity } v = \frac{dx}{dt} = -\frac{E}{p}.$$

Since $E = \frac{p^2}{2m} > 0$, and assuming $p > 0$ (momentum magnitude),

$$v = -\frac{E}{p} \text{ is negative.}$$

Therefore, the wave is traveling in the negative x direction.



Current AI models can help with research

Equation from
arXiv:2402.19356

$$M(s) = \int d^2k \frac{\sqrt{k^2}}{(k+p)^2 + m^2} \left(\frac{k^2}{m^2} \right)^s = (1+x)^{s+1} B(s+2, -1-s) {}_2F_1 \left(s+2, -1-s; 1; \frac{x}{1+x} \right).$$

Do the integral

Mathematica: Fail

Maple: Fail

Claude 2.7: Fail

Gemino 2.5 Pro: Fail

GPT 4o-mini-high: Success!

GPT 5.1: Success!

Genimi 3.0 Success!

can't even input
to Mathematica/Maple
in this form

Do this integral in polar coordinates. Just the angular integral and
simplify into a single radial integral

one often rewrites

$$M(s) = \frac{2\pi}{m^{2s}} \int_0^\infty dk \frac{k^{2s+2}}{\sqrt{((k+p)^2 + m^2)((k-p)^2 + m^2)}},$$

which is a single one-dimensional integral over k .

Then change variables to $y = k^2/m^2$ and $x = p^2/m^2$.

Putting everything together gives the single "radial" integral in y :

$$M(s) = \pi m \int_0^\infty dy \frac{y^{s+\frac{1}{2}}}{\sqrt{(y+x+1)^2 - 4xy}}, \quad x = \frac{p^2}{m^2}.$$

$$M(s) = B(-s-1, s+2) {}_2F_1(-s-1, -s-1; 1; -x),$$

. Equivalently, by the reflection formula one may write

$$M(s) = \frac{\pi}{\sin(\pi s)} {}_2F_1(-s-1, -s-1; 1; -x),$$

simpler form than in the paper!

Followup questions

where the poles in s of this expression

for what range of x and s does the integral converge?

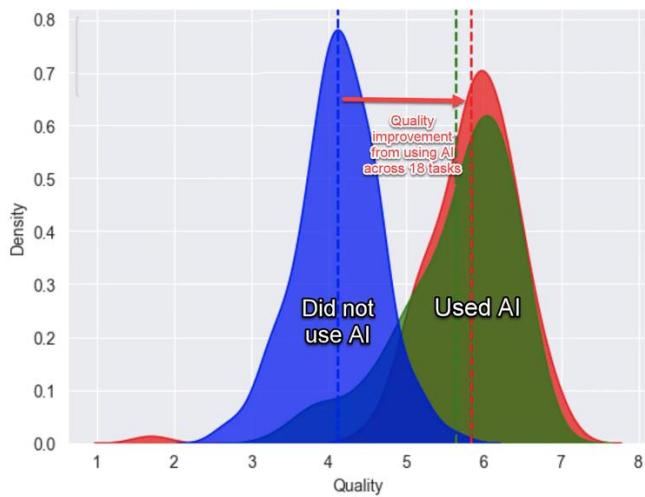
write python code to numerically verify your solution to the integral

- This is completely revolutionary for research
- No going back
- These are the **worst** AI tools we will ever have

Augmented intelligence: LLMs help us thrive

Dell'Acqua et al, "Navigating the Jagged Technological Frontier"
(Harvard Buisness School, 2023)

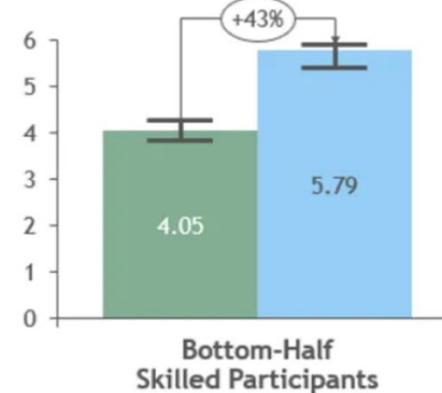
Can LLMs help consultants?



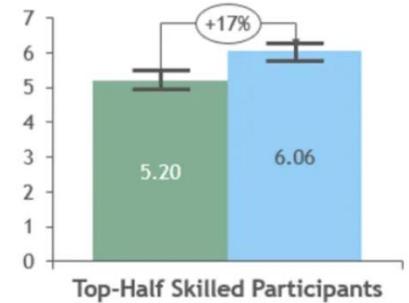
Yes!
By a lot!

AI is a “skill-leveler”

bottom half improved 43%



top half improved 17%



- Mediocre consultants using AI become as good as the best consultants
- The best consultants also improve

10,000 Einsteins

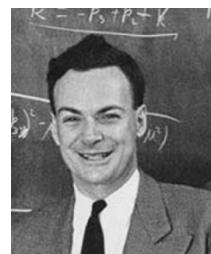
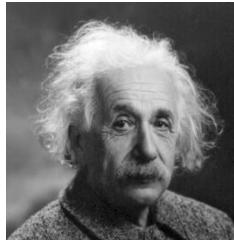
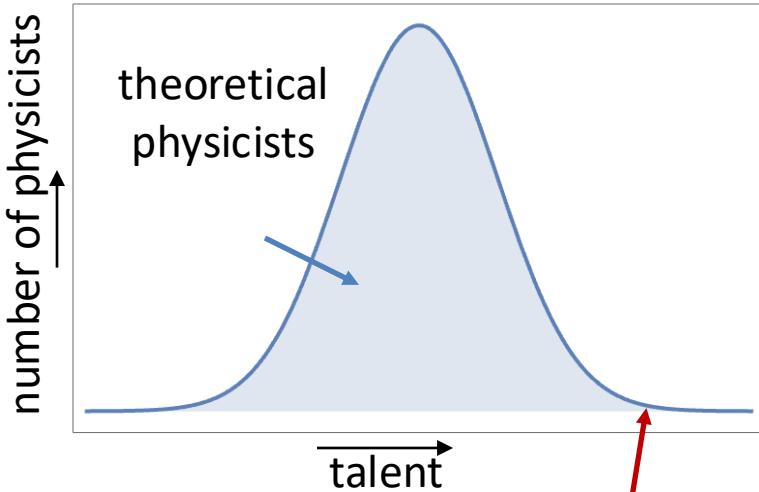
MA

You

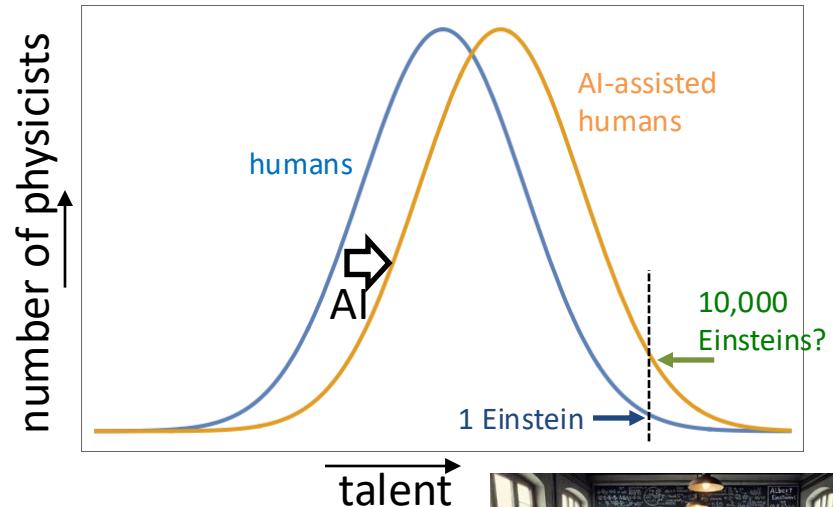
is progress in theoretical physics due to the outsized contribution of a small number of individuals?

ChatGPT

Yes, progress in theoretical physics, like many fields of science, often appears to be driven by the outsized contributions of a relatively small number of individuals.



Can AI be a skill-leveler for theoretical physics?

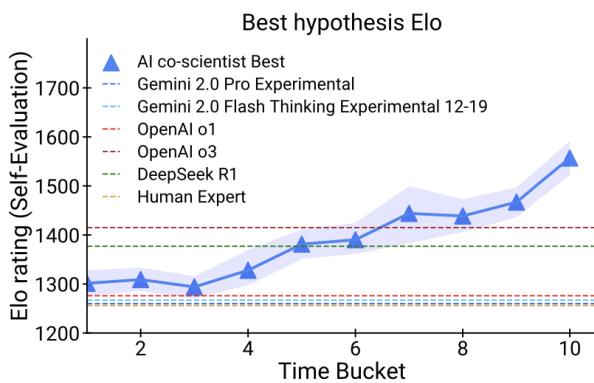


AI Scientists

Google (Feb 2025)

Towards an AI co-scientist

- Generates novel testable hypotheses



(May 2025)

Google DeepMind

2025-5-14

AlphaEvolve: A coding agent for scientific and algorithmic discovery

- Broke a 56-year-old record in matrix multiplication
- Improved on the “kissing problem”

sakana.ai (August 2024)

The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery

August 13, 2024

first principles (Feb 2025)

AI to Explore the Laws of Nature

AgentRxiv: Towards Collaborative Autonomous Research

Samuel Schmidgall, Michael Moor

(March 2025)

The AI Cosmologist I: (April 2025)

An Agentic System for Automated Data Analysis

Adam Moss

Meet Denario — An AI Assistant for Every Step of the Scientific Process

Simons Foundation (May 2025)

periodic labs (Sep 2025)

Today, we introduce Periodic Labs. Our goal is to create an AI scientist.

- Automated materials design

#AI4Math

AlphaProof

Thu 13 November 2025

- Uses LEAN: language to formalize proofs
- Reinforcement learning to find proofs
 - Same approach as AlphaZero for chess



Harmonic

Article | Published: 12 November 2025

Olympiad-level formal mathematical reasoning with reinforcement learning

Erdős problems

- 1535 unsolved problems by Erdős, written 1930-1990
- 580 Solved Pre-2024 (pre #AI4Math)
- Oct 2025: Open AI solves 15 erdos problems
- Nov 30, 2025: Harmonic solves **#124**
 - 6 hours, with LEAN.
- Dec 1, 2025: Harmonic solves #481

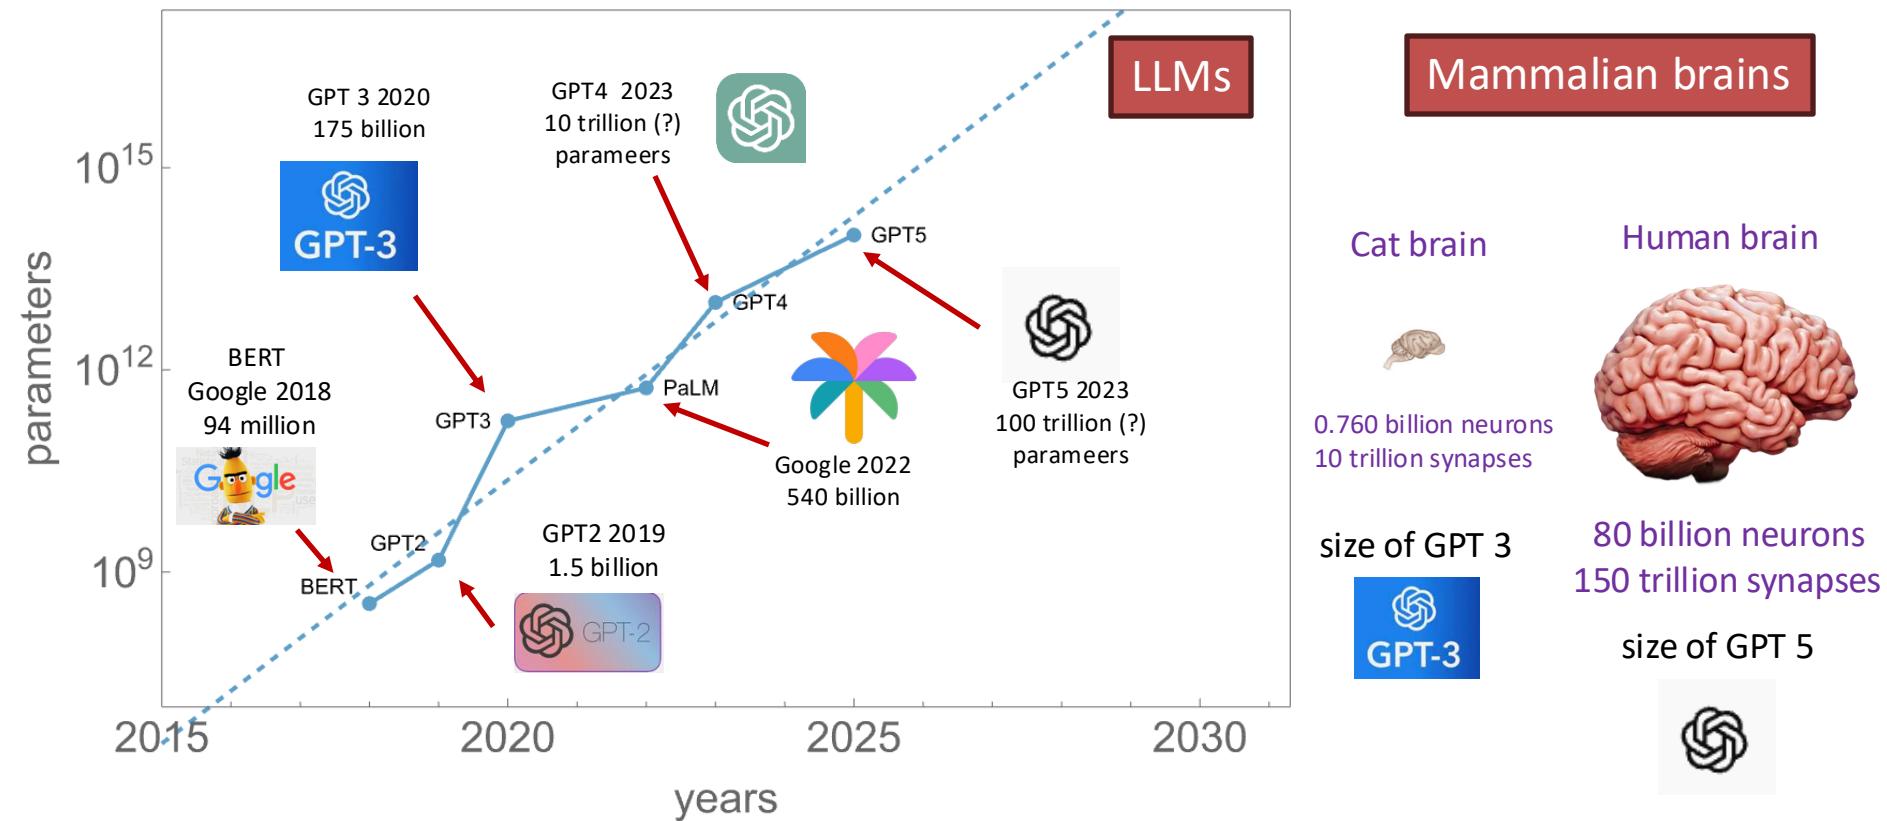
```
example (x : Nat) : 0 < match x with
| 0   => 1
| n+1 => x + n := by
  grind
```

For any set of integers $d_i \geq 3$ satisfying

$$\sum \frac{1}{d_i-1} \geq 1,$$

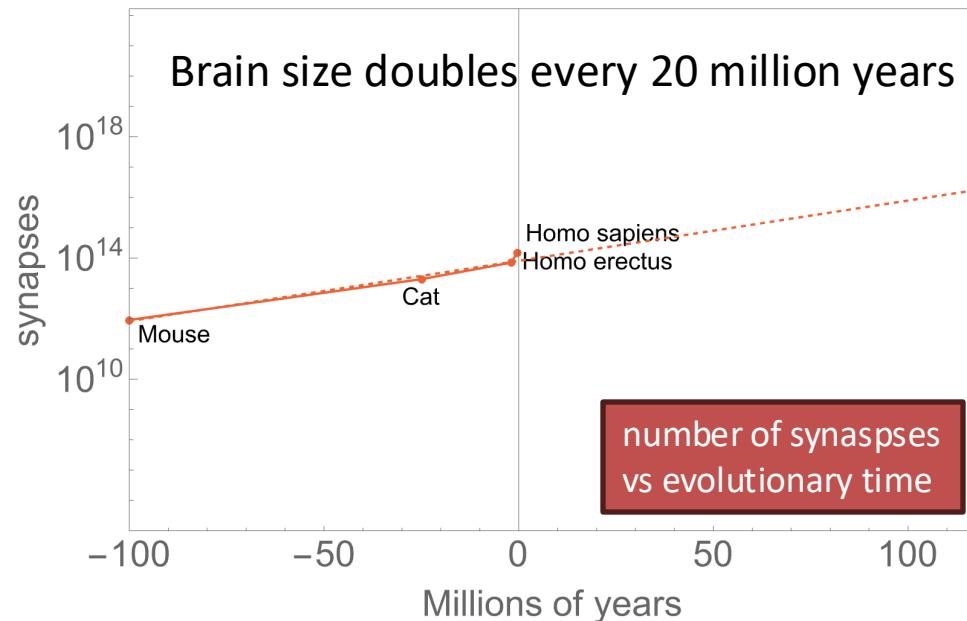
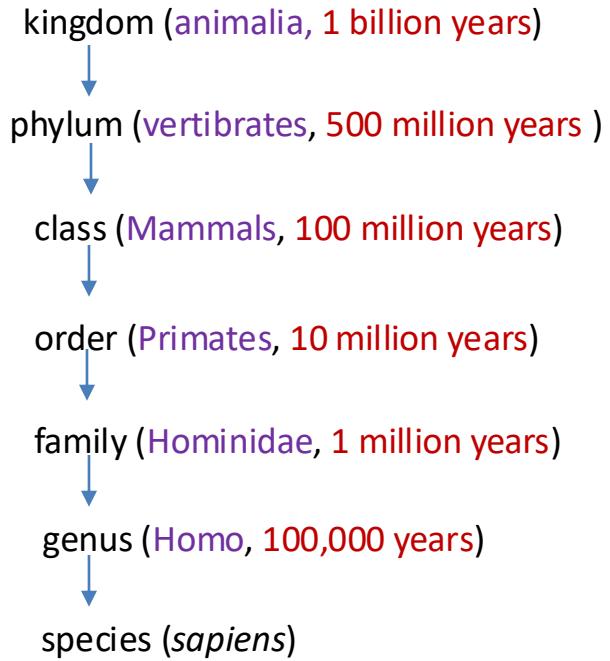
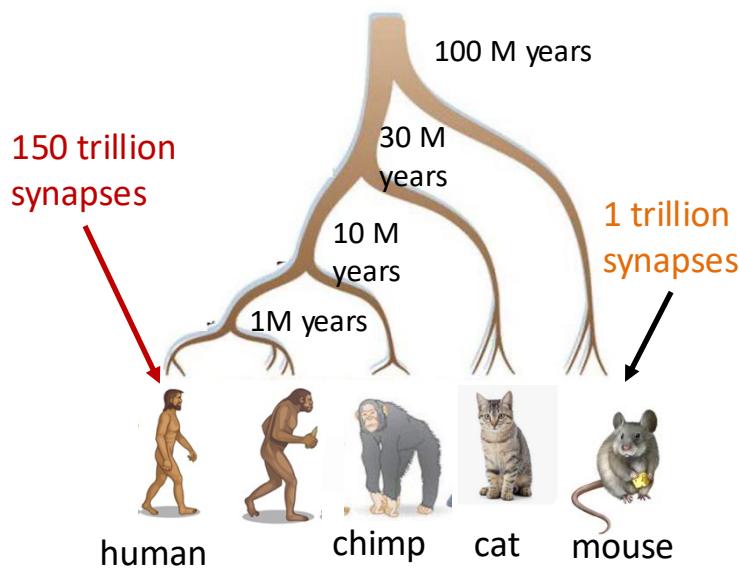
can every natural number be expressed as a sum of distinct powers of these d_i ?

Large Language Models are growing fast



- Exponential growth: LLMs are 10 times more powerful each year!
- # parameters in current LLMs (100 trillion) \approx # synapses in human brain (150 trillion)

Evolution is slow



The universe is **only 15 billion** years old

- Only enough time to evolve kingdoms
- What would evolve after **1 trillion years?**

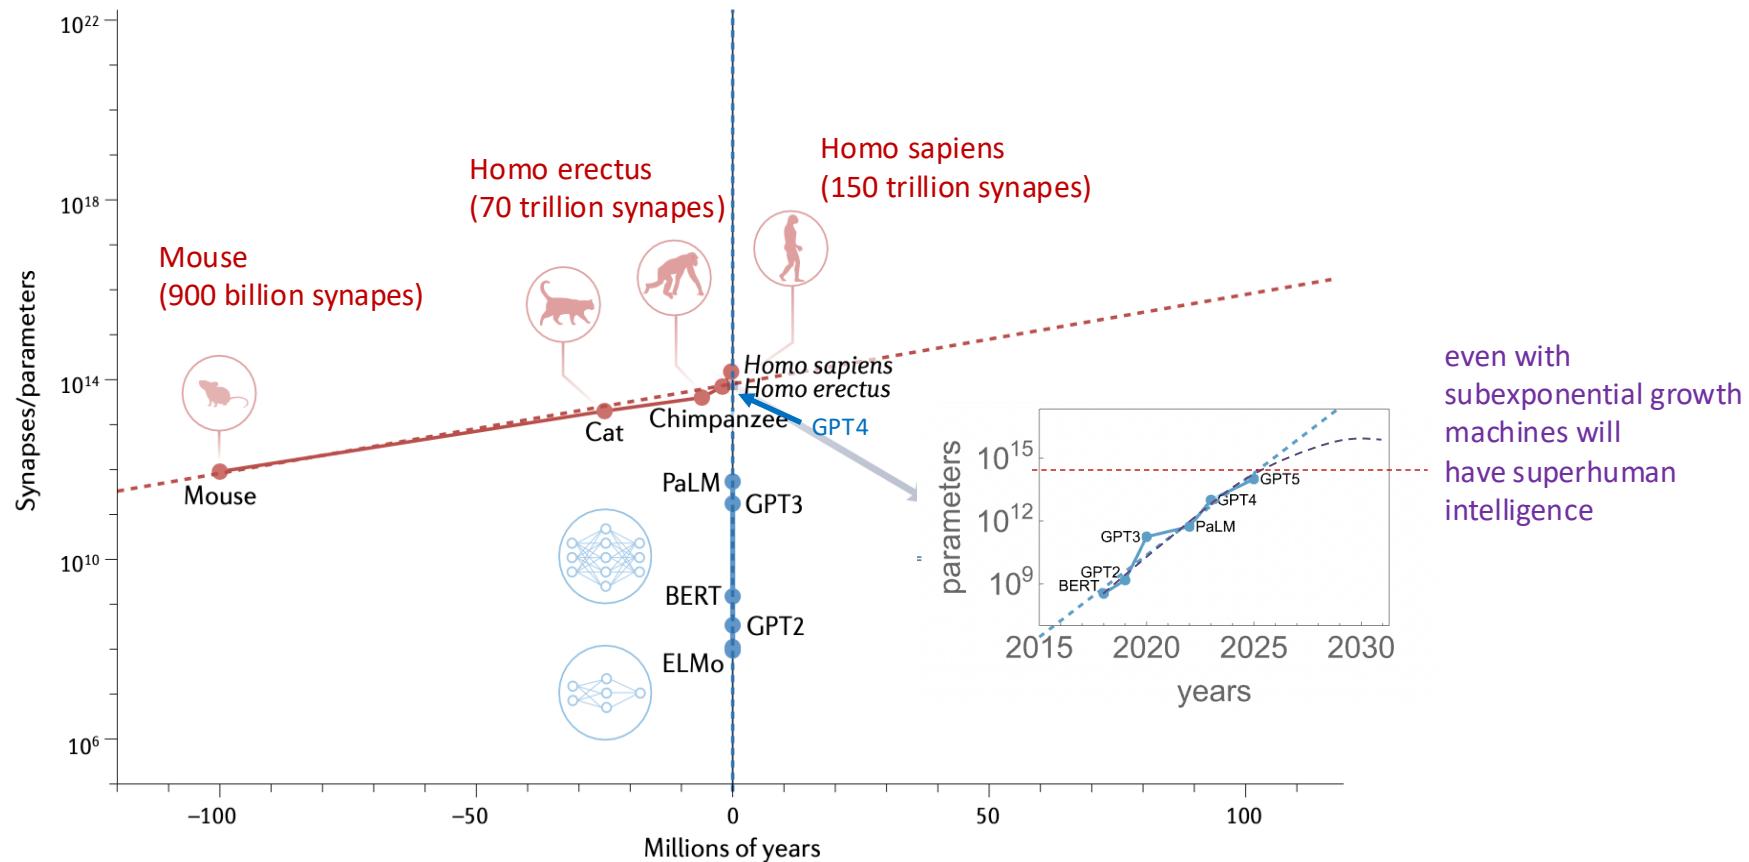
We are not
the endpoint of intelligence

We are only the smartest things
to evolve on earth, so far

Machine vs. Biological intelligence

- Machine intelligence grows by a factor of 10 in 1 year
- Biological intelligence grows by a factor of 2 in 20 million years

MDS, "Should artificial intelligence be interpretable to humans?"
Nature reviews physics (2022)



- Both AI and biological intelligence grow exponentially
- **Factor of 10^7 difference in exponent**
- Intersection time, when machines and biology have comparable "intelligence" is **now**

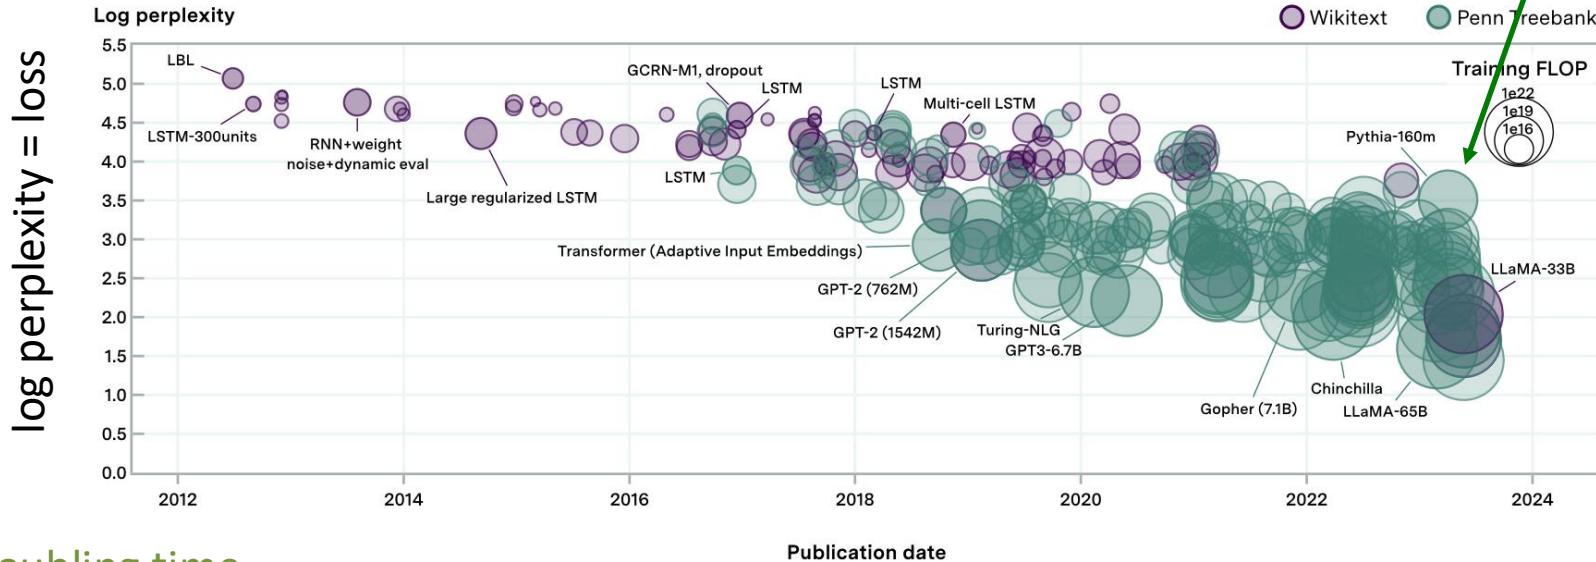
Will we run out of data to train on?

ALGORITHMIC PROGRESS IN LANGUAGE MODELS

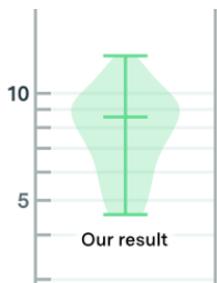
Ho et al. arXiv:2403.05812

size is compute

How much of advance is algorithmic progress vs data + compute?



doubling time

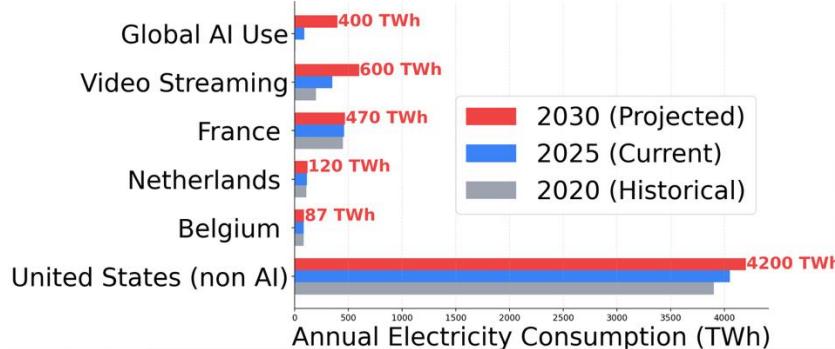


algorithmic doubling time = 6 to 14 months!

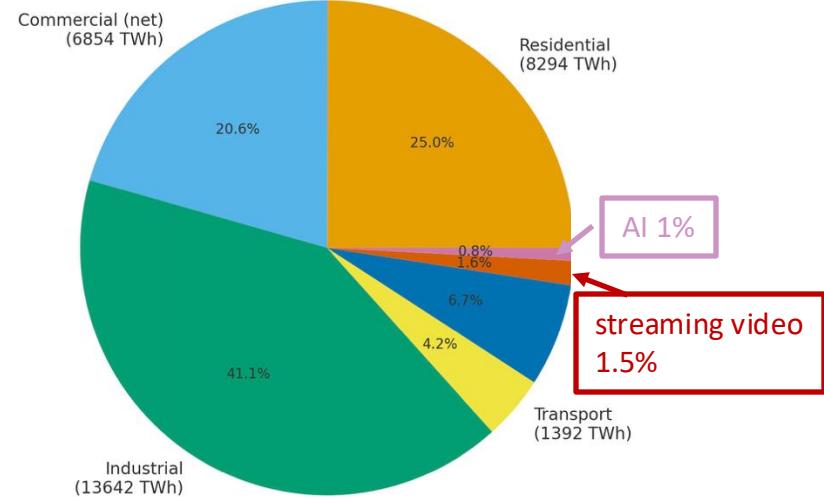
- controlling for data and network size
- algorithms will continue to get better, especially when written by AI

Will we run out of energy?

Energy consumption



2030 projections

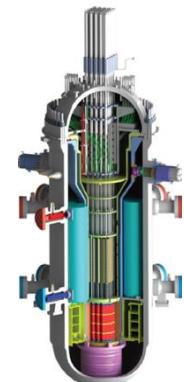


- By 2030 AI may use as much energy as France
 - 1% of global energy budget
 - **This is a lot of energy!**
 - Data centers also require a lot of water for cooling

Counter-arguments

- AI companies investing heavily in **alternative energy**
 - **Fission**: Small Modular Reactors
 - **Fusion**: Helion, Zap fusion, Pacific Fusion etc. (>\$1B invested)
 - **Geothermal**: growing interest and investment
- If AI does the work of 10 engineers
 - How much energy do 10 engineers use to do the same work?
- What would you say about the industrial revolution?

SMR



But physics requires creativity!

Torrence Test of Creative Thinking (1960)

- How many ways can you think to use a water bottle?
- Suppose you could be invisible for a day. What problems might that create?
- How could you improve a toy car to make it more fun for children?

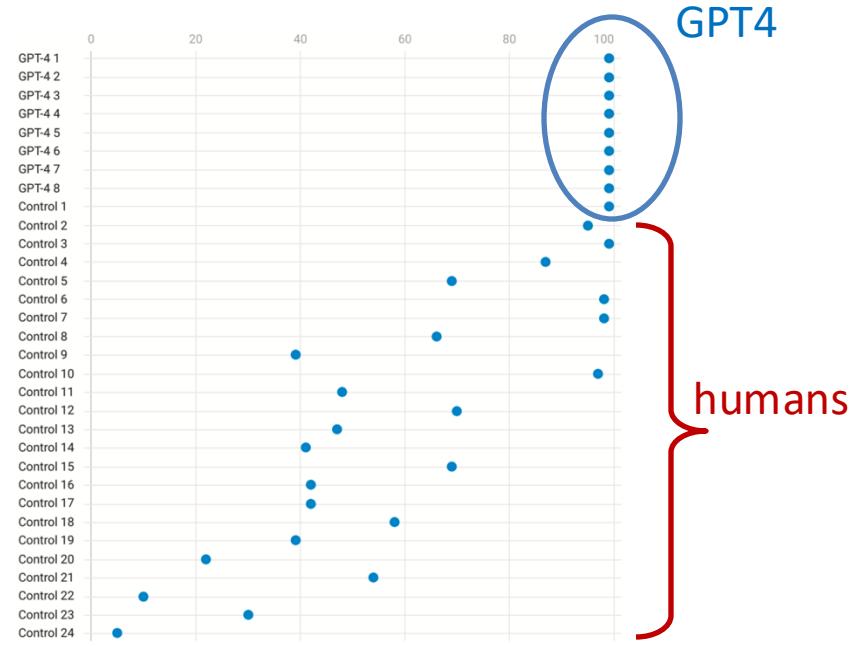
Article | [Open access](#) | Published: 10 February 2024

The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks

[Kent F. Hubert](#)  [Kim N. Awa](#) & [Darya L. Zabelina](#) 2024

[Scientific Reports](#) 14, Article number: 3440 (2024) | [Cite this article](#)

11k Accesses | 252 Altmetric | [Metrics](#)



Torrence Test score

GPT4 more creative than 99% of humans

Try it with Gemini 3 Pro

how many ways can you think of to use a water bottle?

1. Survival & Physics (Using Refraction & Pressure)

- **The Solar Fire Starter:** If you have a clear, smooth, curved plastic bottle, fill it with water. Use it as a convex lens to focus the sun's rays onto dry tinder (like paper or dry leaves) to start a fire in an emergency.
- **A "Liter of Light" Skylight:** Used in many off-grid communities, you can install a clear water bottle filled with water and a bit of bleach (to prevent algae) through a hole in a shed roof. It refracts sunlight, scattering it 360 degrees and glowing like a 55-watt bulb without electricity.
- **DIY Barometer:** By connecting a water bottle to a jar with a straw and sealing the bottle, you can track changes in atmospheric pressure. When high pressure hits, the water level in the straw rises; low pressure makes it drop.
- **Egg Yolk Separator:** Squeeze an empty plastic bottle slightly, place the mouth over an egg yolk, and release your grip. The suction will pull the yolk cleanly inside, leaving the white behind.

.... easily produces 1000 more ideas

Try it with Gemini 3 Pro

5. "Cloud" Science

- **Cloud in a Bottle:** This is a classic experiment. Put a small amount of isopropyl alcohol in a bottle, pressurize it using a bike pump through a cork, and then rapidly release the pressure. The sudden drop in pressure and temperature causes the alcohol vapor to condense instantly, forming a visible cloud.

generate an image of the cloud science idea and animate it



3. The Future: Vibe Physics



Image: Nano Banana Pro 2025

Predictions for the future

1. Vibe physics
2. Publication crisis
3. Education crisis
4. Catharsis
 - Revolutionary developments in nearly every area of science, including physics

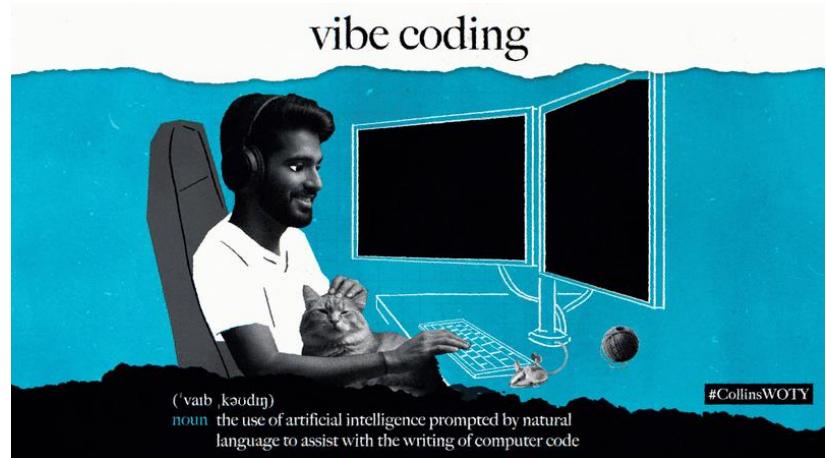
Vibe Coding

Collins Dictionary

Word of the Year
2025
#CollinsWOTY

Andrej Karpathy 
@karpathy

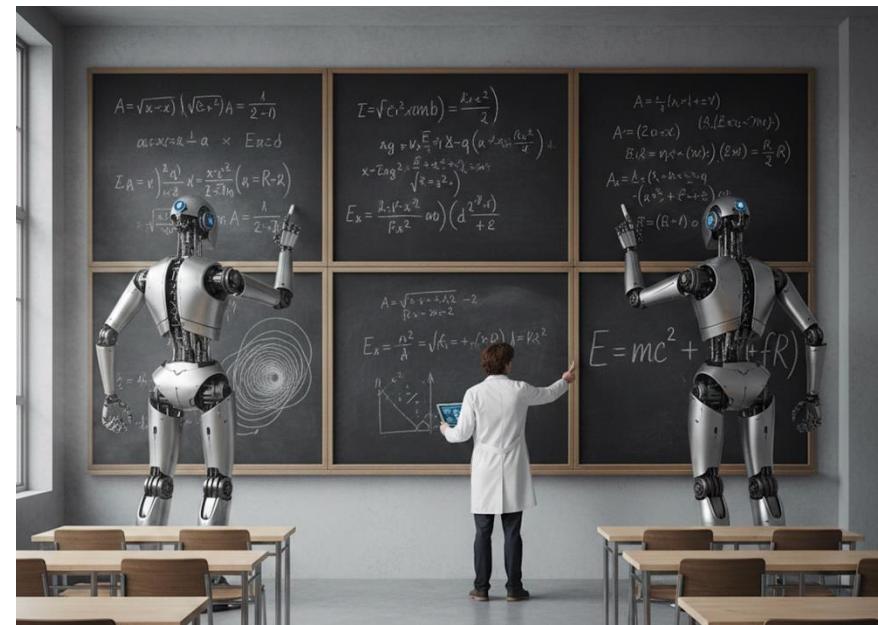
6:17 PM · Feb 2, 2025 · 5.1M Views



There's a new kind of coding I call "**vibe coding**", where you fully give in to the vibes... and **forget that the code even exists**. It's possible because the LLMs ... are getting too good.... 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** ...

1. Vibe Physics

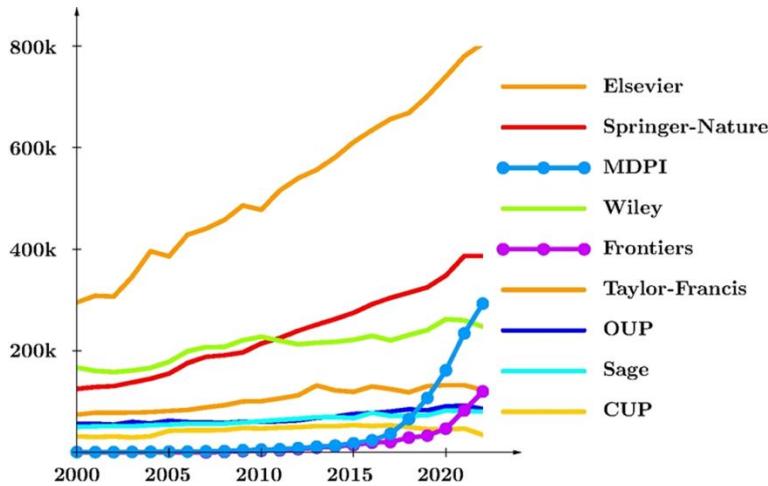
- Ask LLM to suggest approaches to a problem
- Pick an approach, ask it to do some calculations
- Ask it to suggest some toy models
- Ask it to compute various things in the models
- Have it generate figures, write the paper, etc.



2. Publication crisis

Excessive growth in the number of scientific publications

(Oct 2024)
OCS Ouvrir la science !



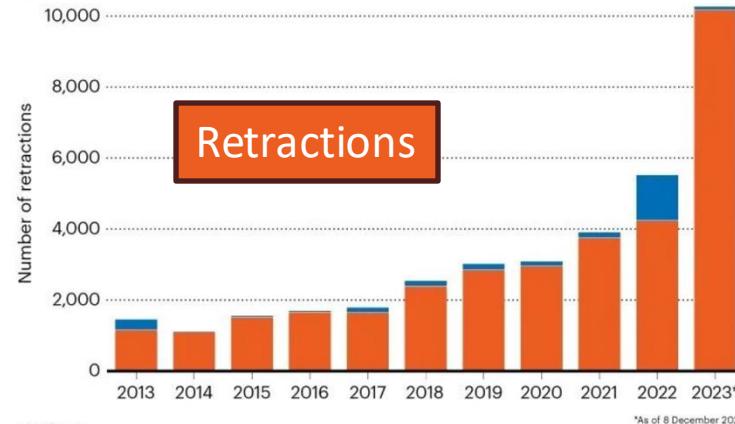
NEWS | 12 December 2023

More than 10,000 research papers were retracted in 2023 – a new record

A BUMPER YEAR FOR RETRACTIONS

Retraction notices in 2023 have passed 10,000, largely because of more than 8,000 retractions by Hindawi.

■ Journal articles ■ Conference papers



*As of 8 December 2023

©nature

NEWS | 27 November 2025 | Correction [01 December 2025](#)

Major AI conference flooded with peer reviews written fully by AI

- 21% of 75,800 peer reviews at ICLR 2025 generated by AI

AI writes, AI reviews



Exponential Growth

- AI tools have made manuscript generation 2x or 3x faster
 - Less time and thought, full of mistakes
- Reviewers can't keep up, and don't want to sift through AI crap

The crisis

- Burnout
- Refusal to review
- Rubber stamping bad science



Why do we need publications if people don't write or read them?

| THE END OF THE JOURNAL

THE "PAPER" IS A FOSSIL

The concept of a "Journal Issue" is a relic of the printing press. It batches information arbitrarily.

In a digital, AI-native world, knowledge is a **stream**, not a batch. The static PDF, frozen in time, is no longer a valid container for truth.



| CODE IS THE NEW PAPER

For AI and Comp-Bio, the "Paper" is often just an advertisement for the Repo.

- **Hugging Face > ArXiv:** Researchers look for the model weights first, the description second.
- **Executables:** Papers like *ColabFold* allow users to run the science in the browser. Static PDFs cannot compete with interactive truth.
- **Zenodo:** Massive growth in DOI-minted datasets. The *data* is becoming the citable unit.



THE LIVING KNOWLEDGE GRAPH

Dynamic Truth: Instead of static PDFs, scientists contribute nodes to a global Knowledge Graph.

- ✓ **Live Updates:** If a method is refuted, every "paper" using it is flagged automatically.
- ✓ **AI Synthesis:** You don't read papers. You ask the Graph: "What is the consensus on protein X?" and AI generates a summary from the raw data nodes.

3. Education Crisis

HIGHER EDUCATION

Everyone Is Cheating Their Way Through College ChatGPT has unraveled the entire academic project.



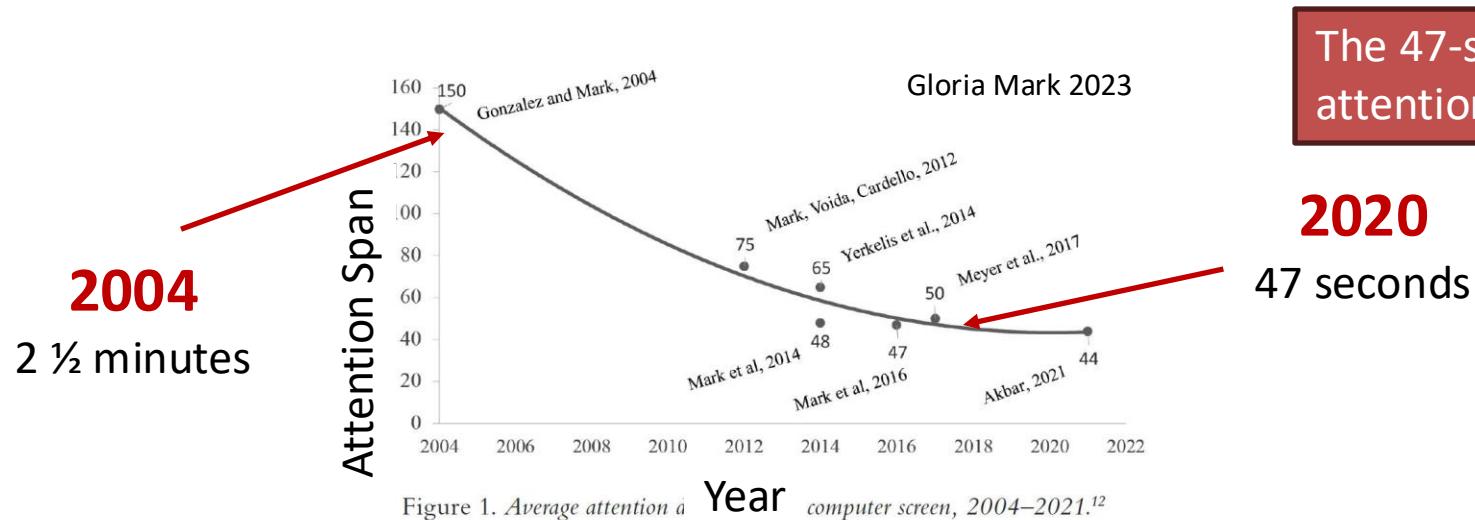
By James D. Walsh, *Intelligencer* features writer

MAY 7, 2025

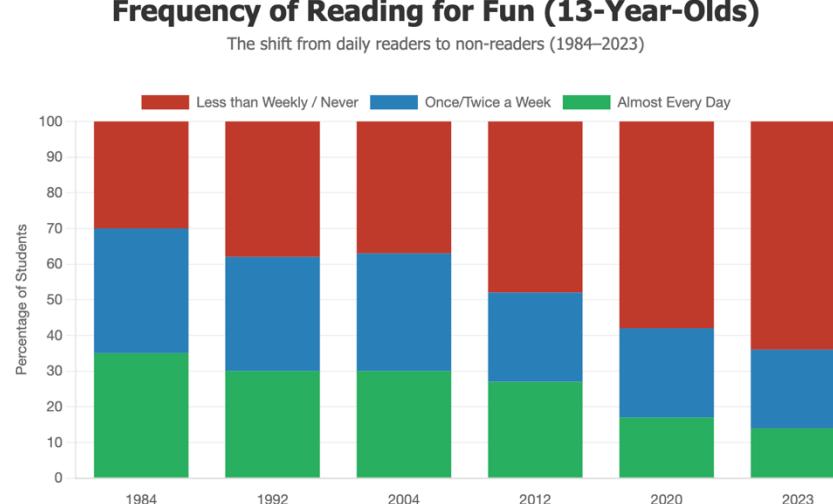
NEW YORK
MAGAZINE

<https://nymag.com/intelligencer/article/openai-chatgpt-ai-cheating-education-college-students-school.html>

Education crisis before AI



Students don't read anymore



Features / Teaching

For whom the book tolls

How declining reading levels are reshaping the way universities teach.

September 24, 2025

“There’s no way that I could assign the amount of reading today.”

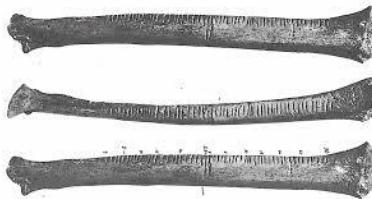
The 47-second attention span

2020

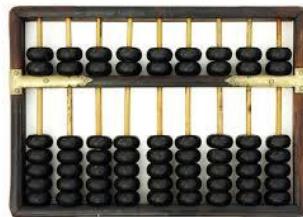
47 seconds

What is worth learning?

Arithmatic:



tally bones
(20,000 BC)



abacus
(2,000 BC – 1500 AD)



slide rule
(1620-1970)



calculators
(1970-2022)



Generative AI
(2023 +)

Programming:

2020

everybody must learn to code!



2025

nobody needs to know how to code!



THE PIVOT

From "One Size Fits All" to "One Size Fits One"

Technology is finally catching up to the pedagogy. The future isn't about working harder; it's about leveraging AI to scale the "Master-Apprentice" model.

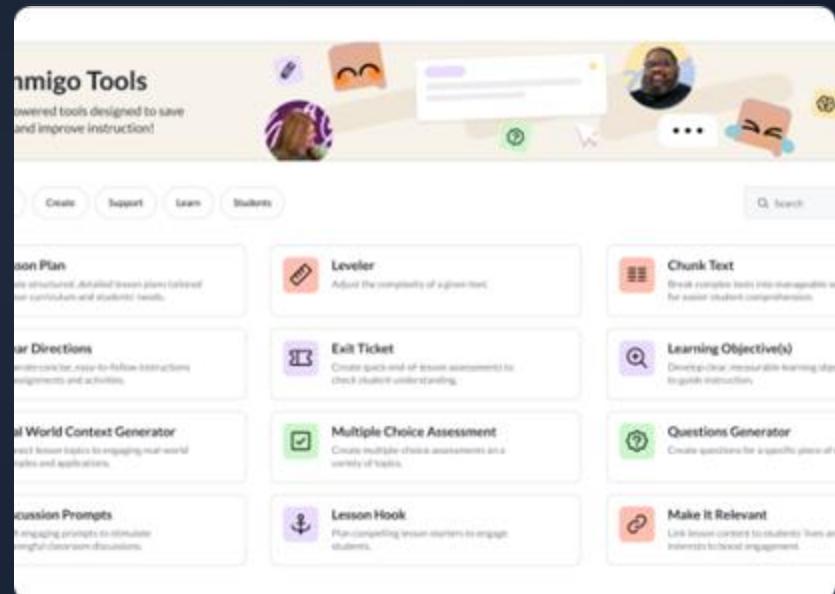
The 1:1 AI Tutor

Solving Bloom's 2 Sigma Problem

The Stat: Personalized 1:1 tutoring improves student outcomes by 2 standard deviations (98th percentile).

The Solution: AI Tutors (like Khanmigo) make 1:1 tutoring scalable for the first time in history.

- ✓ **30% Boost:** Early pilots show AI personalization improves outcomes by up to 30%.
- ✓ **Always On:** A tutor that never gets tired, never judges, and explains "Entropy" 50 different ways until it clicks.



- No more inefficient lectures!
- Education is flattened: everyone everywhere can learn the same way

| The Infinite Lab



Zero Marginal Cost: In a wet lab, a mistake costs \$50 in chemicals. In VR, it costs \$0.

- ✓ **Safety:** Students can handle "radioactive" materials or "explode" reactors safely.
- ✓ **Repetition:** "Practice makes perfect" was impossible in science labs due to cost. Virtual labs allow infinite attempts.
- ✓ **Effectiveness:** Studies show virtual labs are **equally effective** for content mastery, though hands-on is still needed for dexterity.

| The Guide on the Side

The Professor Evolved

If AI delivers the content and grades the basics, what does the human do?

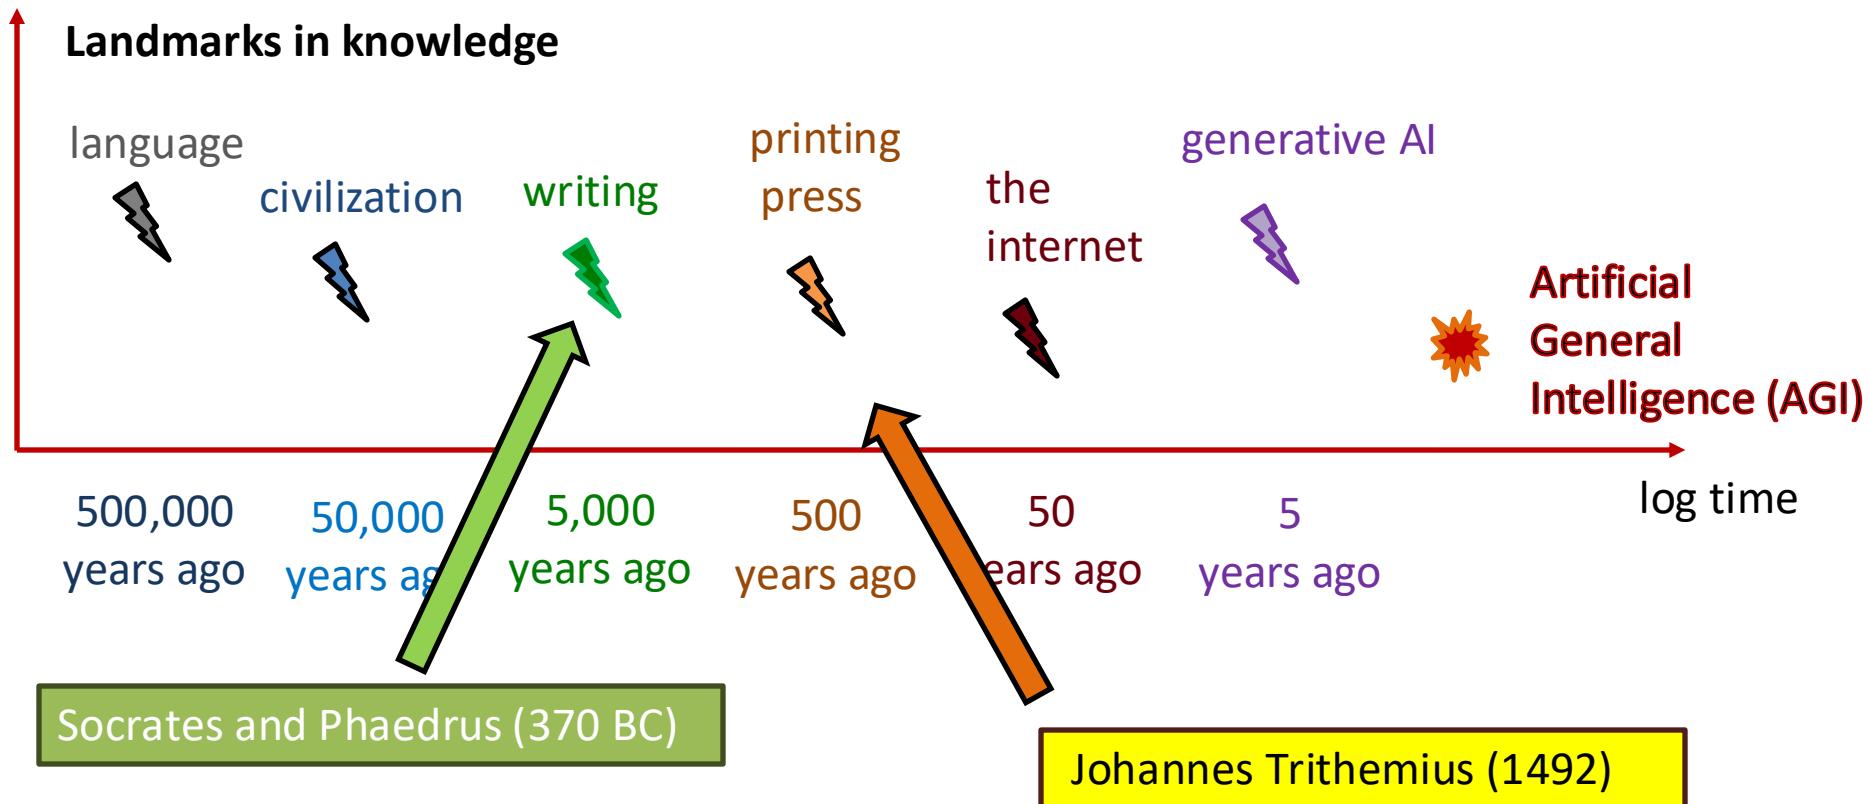
Mentorship.

The professor shifts from a "Content Broadcaster" to a high-level consultant, helping students navigate complex projects, ethical dilemmas, and creative roadblocks.

This is the most human, and valuable, part of education.



4. Catharsis



And it is *not true wisdom* that you offer your disciples, but *only the semblance of wisdom*, for by telling them of many things without teaching them you will make them seem to know much while for the most part they know nothing.

"The printed book is made of paper and, like paper, will quickly disappear. But the scribe working with parchment ensures lasting remembrance for himself and for his text"

4. Catharsis



Bill Gates
and David Letterman
(1995)

Artificial
General
Intelligence (AGI)

5
years ago

Isaac Trithemius (1492)

they rely on that which is written.

And it is no true wisdom that you offer your disciples, but only the semblance of wisdom, for by telling them of many things without teaching them you will make them seem to know much while for the most part they know nothing.

is made of paper and, like paper, will quickly disappear. But the scribe working with parchment ensures lasting remembrance for himself and for his text"

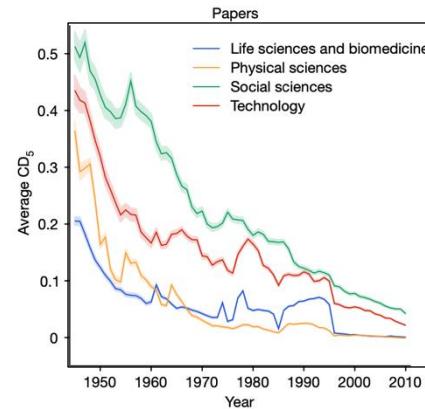
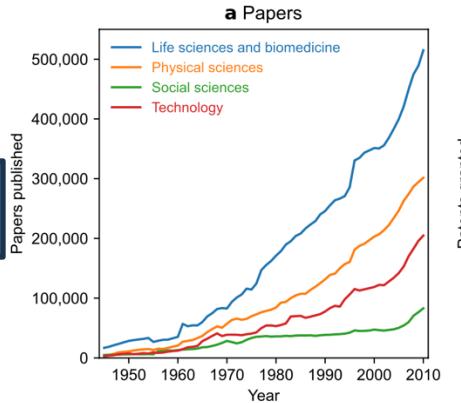
Theoretical particle physics may have stalled

Article | Published: 04 January 2023

Papers and patents are becoming less disruptive over time

[Michael Park, Erin Leahey & Russell J. Funk](#) 

more and more papers are written



the papers are less and less innovative

Maybe the problems are just too difficult (for us)

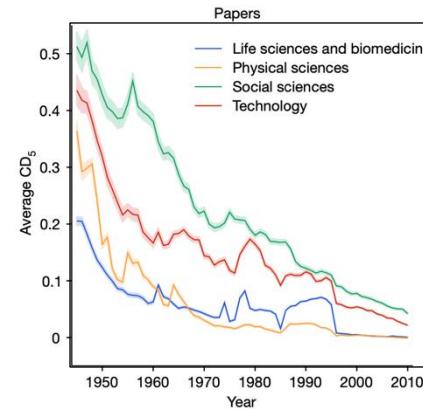
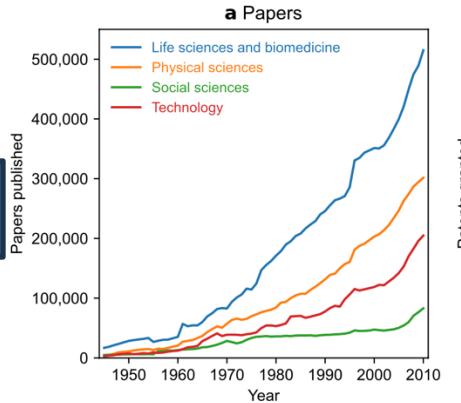
Theoretical particle physics may have stalled

Article | Published: 04 January 2023

Papers and patents are becoming less disruptive over time

[Michael Park, Erin Leahey & Russell J. Funk](#) 

more and more papers are written



the papers are less and less innovative

Maybe the problems are just too difficult (for us)

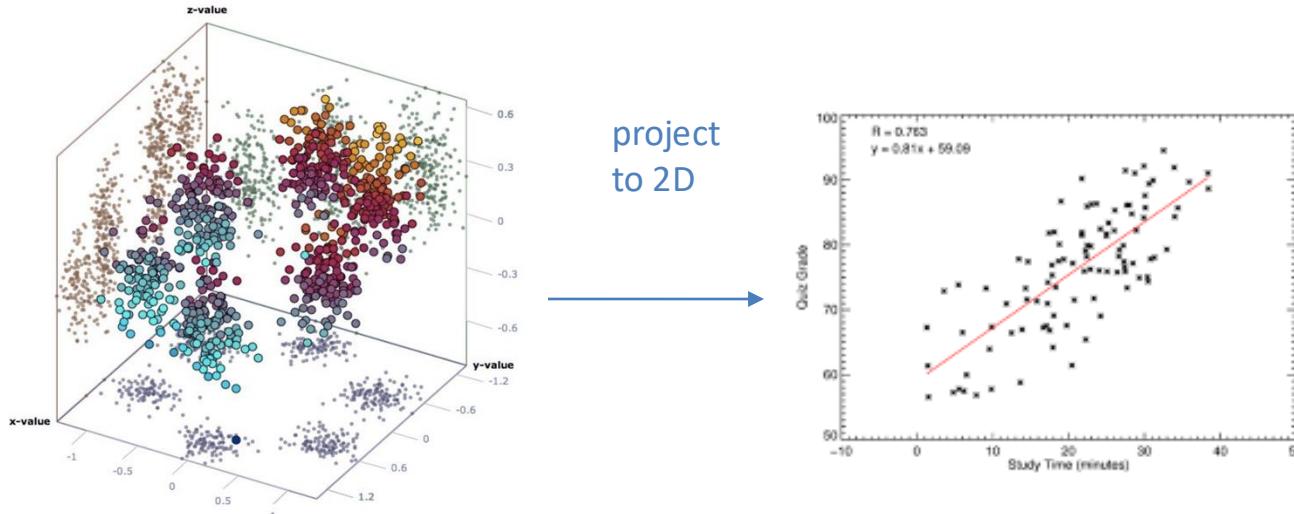


- Could a cat ever learn to play chess?
- Humans have limits too

Why should Homo sapiens be able to understand the theory of everything?

Humans are limited by biology

Humans like to “visualize”



Why do we do this? Because we have **eyes**

- 2D is not special to a machine.
- Machines can “visualize” in d dimensions

Eyes have **nothing to do**
with fundamental physics!

Humans can only hold 5-9 concepts in working memory at once

- We like simple-looking equations

$$i\partial_t\psi = H\psi \quad i\partial\psi = m\psi \quad G_{\mu\nu} = \kappa T_{\mu\nu}$$

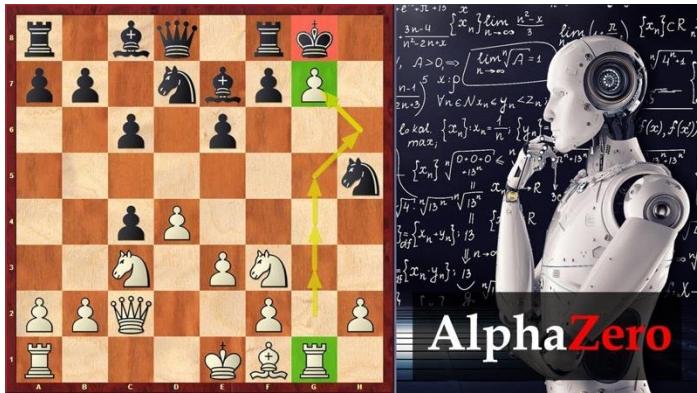
- **Computer** memory can **handle much more** than 5-9 concepts at once
- They can understand systems not governed by simple equations

What do we need to progress further?

Current state-of-the art can calculate ***anything that has been calculated***

- Trained on textbooks, journal articles, physics.stackexchange, chegg, notes, etc.
- Books, papers, etc. **written by human beings who read books**, papers, etc.
- i.e. we generate our own training data

Alpha Zero: learns to solve chess problems by generating its own training data



Current LLMs

- can generate and solve problems
- LLMs can refine their own code
 - GPT5 was vibe-coded using GPT4

Language models are very close to training themselves like they are their own Ph. D. students

Superhuman intelligence

Suppose a machine understands the theory of everything but we don't

- e.g. can calculate electron mass from scratch
- e.g. can explain dark matter

Is this enough or do we need to understand it too?



- The authors of **Popular science books** understand the details; we just get the general idea

I don't understand the proof of Fermat's last theorem

- I'm glad that somebody does
- Does it matter that the person is human?

If a machine understands fundamental physics it can

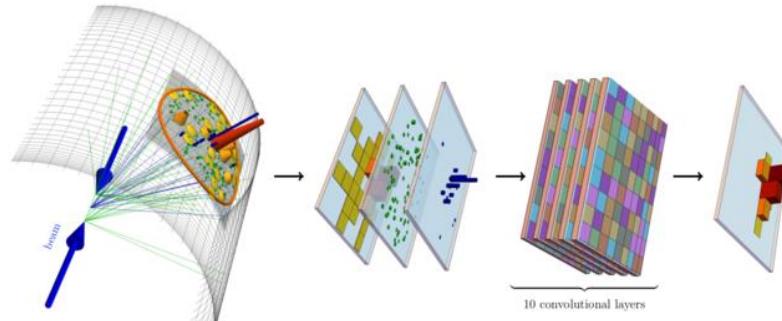
1. Dumb it down so we can get the general idea
2. Find practical applications

Is this what we want? No.
But maybe it's the best we will get.

Because of AI, I am now optimistic
for substantive progress in high-
energy theory in my lifetime

Conclusions

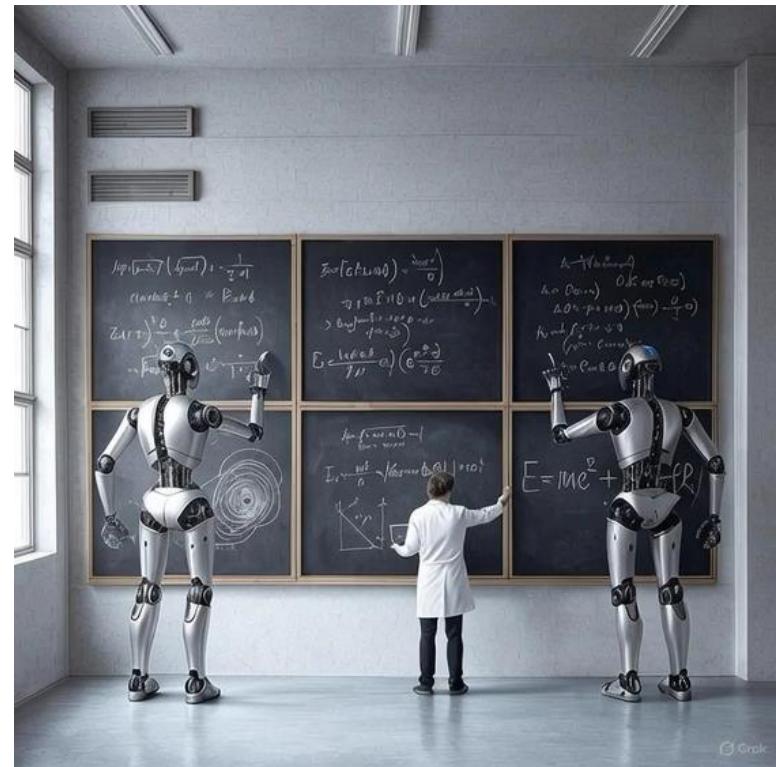
- Past: **machine learning as a tool**



- Present: **augmented intelligence**



Future: **Vibe Physics**



Thank you



BACKUP SLIDES

The AI triad (Jonathan Zittrain)

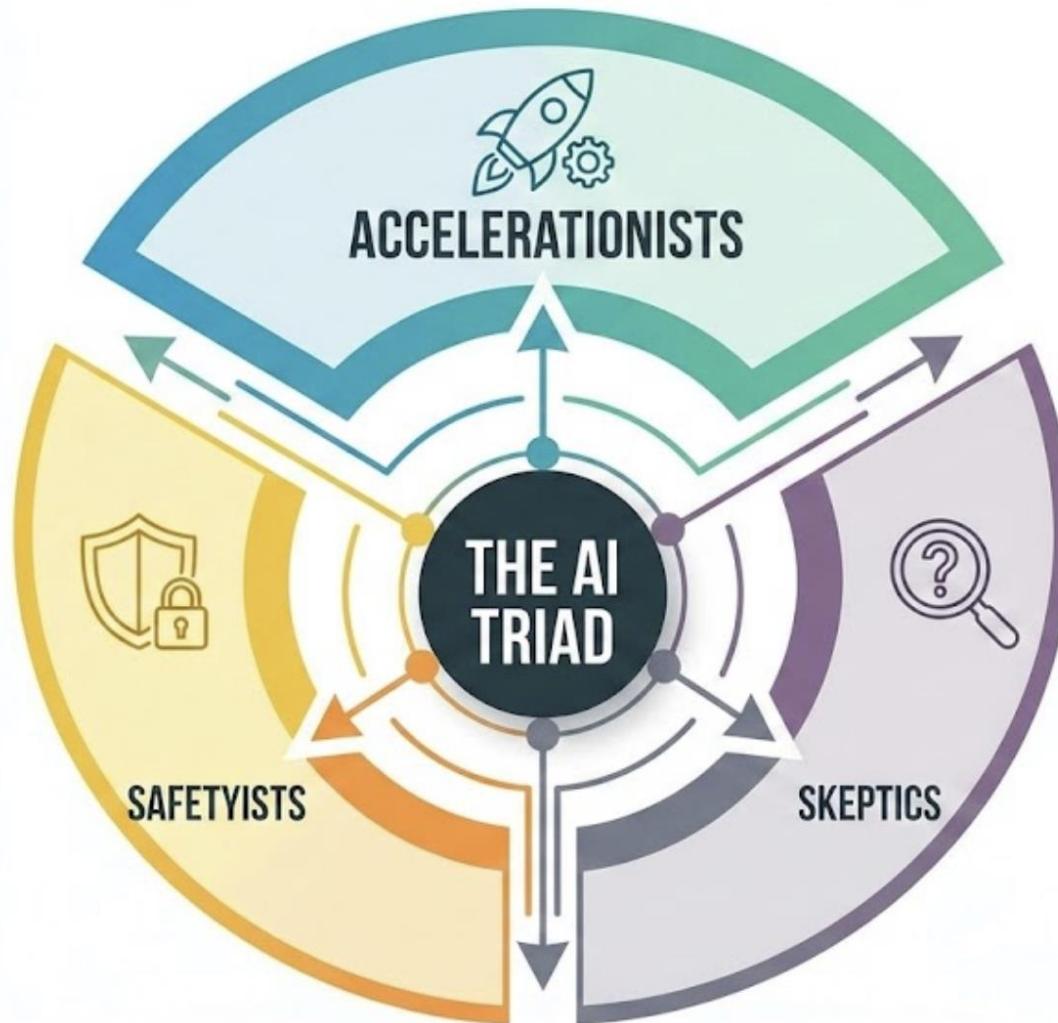
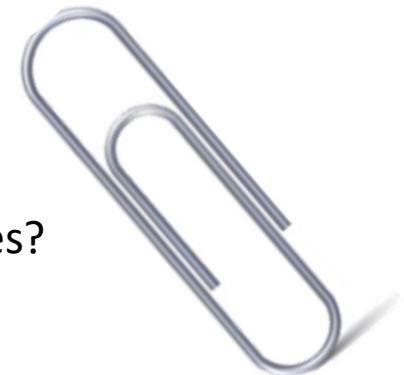


Image: Nano Banana Pro 2025

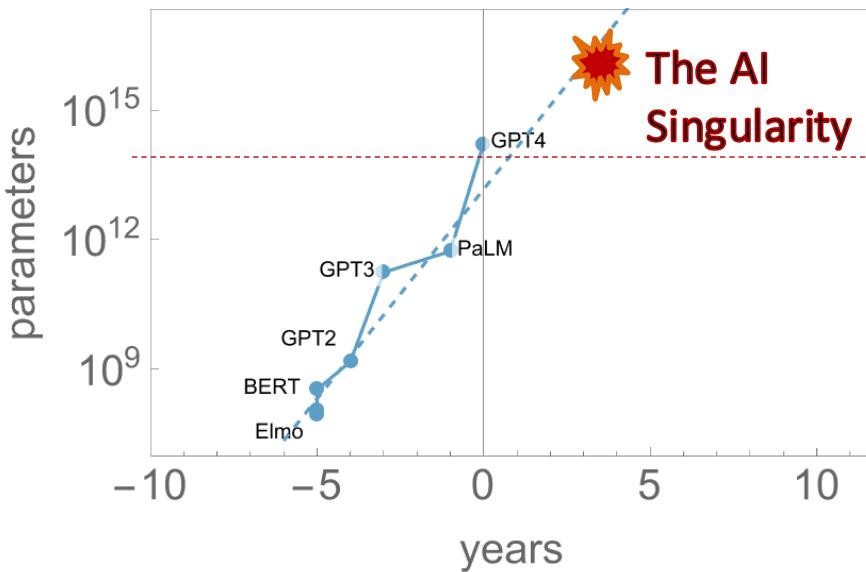
What could go wrong?

❖ “Paperclip Maximizer” Nick Bostrom, 2014

- Suppose a machine was designed to maximize paperclip production.
- To truly maximize paperclips, it would destroy everything on earth to produce paperclips



❖ What if machines can write their own code and improve themselves?



Not that far off

- Machines can already write and debug code
- Machines can already evaluate their own flaws

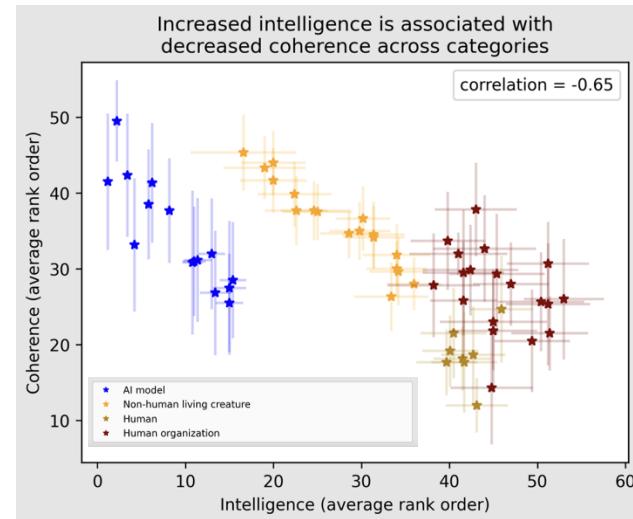
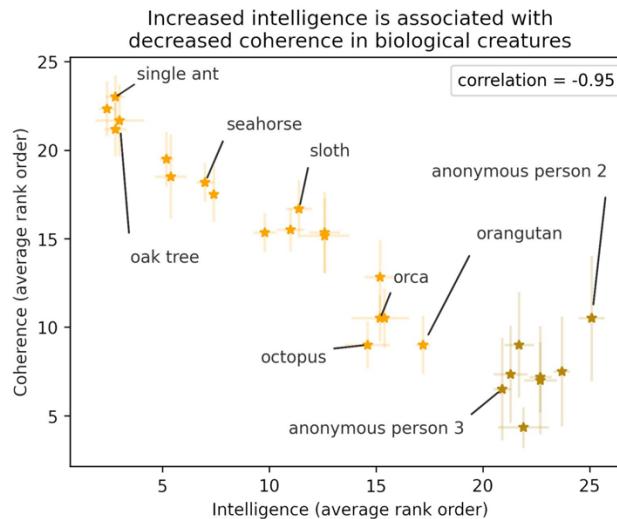
Are we doomed?

Maybe it won't be so bad

Jascha Sohl-Dickstein [Google Deepmind <https://sohldickstein.com/>]

- “We are a **hot mess** of inconsistent, self-undermining, irrational behavior, with objectives that change over time”

Coherence
= focus on a single goal



- Intelligence is anti-correlated with coherence
 - More profoundly, intelligence *requires* incoherence
- Are intelligent machines dangerous? Maybe
- Are intelligent animals dangerous? Definitely

We are not intelligent enough to predict what higher intelligence will do