Theoretical High-Energy Physics and Artificial Intelligence

EuCAIFCon

April 30, 2024

Matthew Schwartz

Harvard University

The NSF AI Institute for Artificial Intelligence and Fundamental Interactions (IAIFI)

Boston, MA, USA

Outline

Part 1: The Present

- Symbolic ML for high energy theory

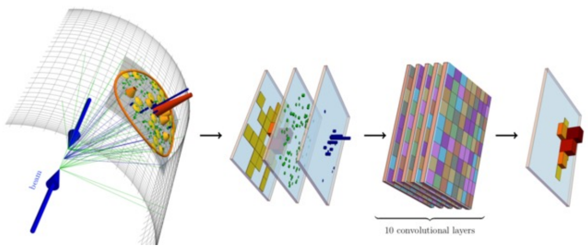

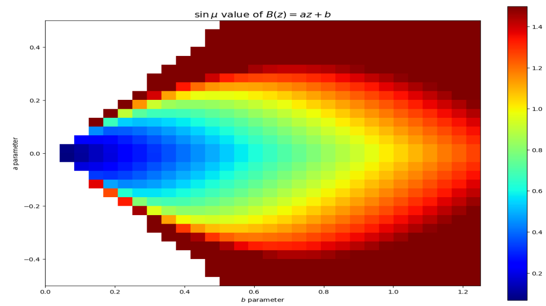

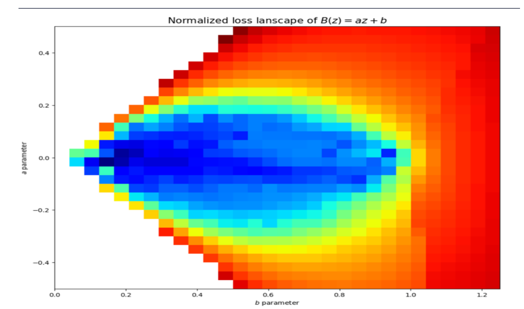

Part 2: The Future

- Can machines do theoretical physics?

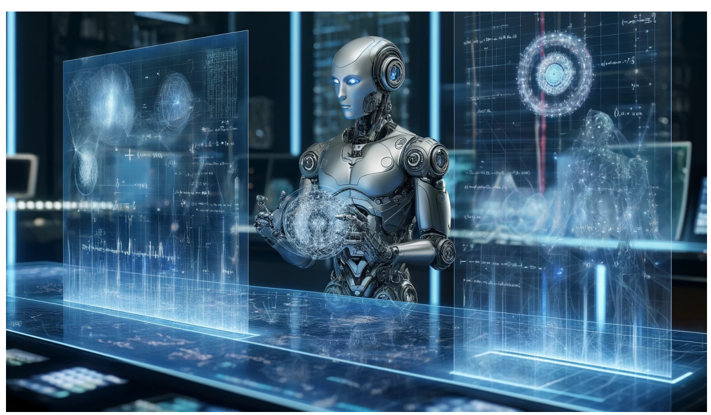

0. The Past

Björn Karlsson, MidJourney January, 2024

The Past: Collider Physics

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

Convolutional networks for facial recognition

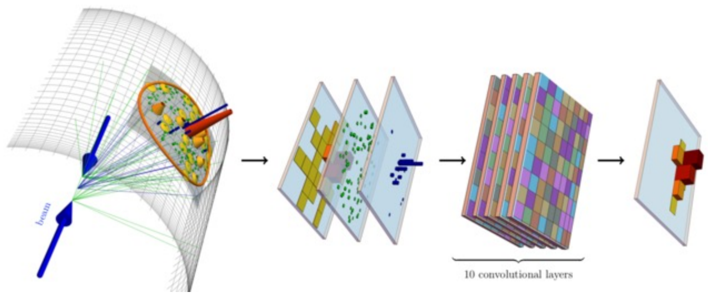

Point clouds/deep sets for self-driving cars

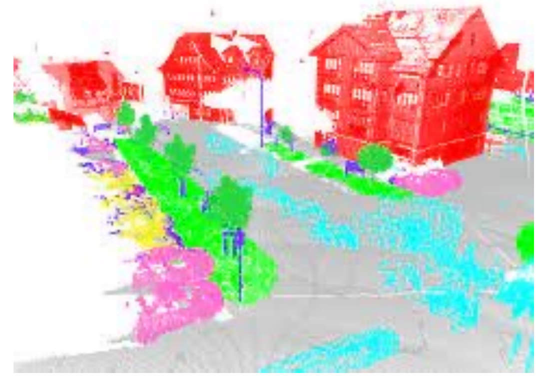

Top Tagging (2008 – 2022)

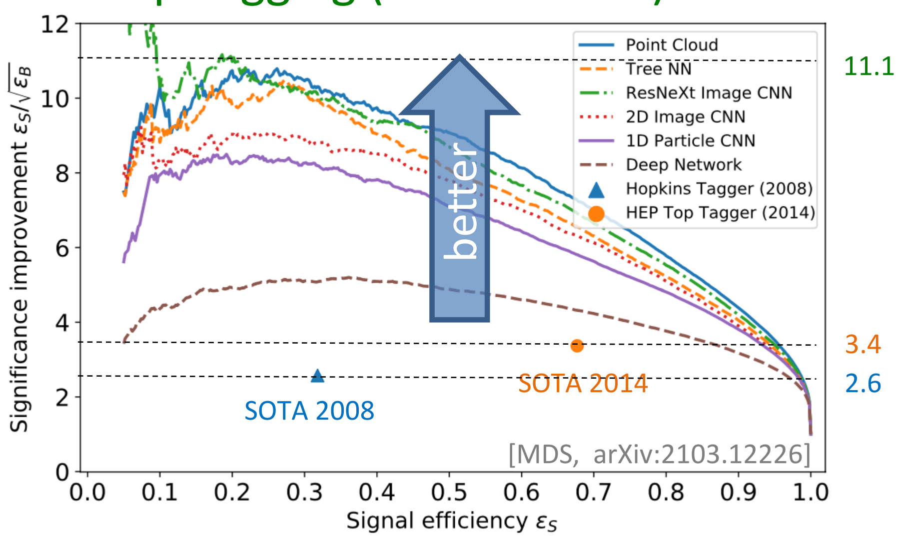

Lessons

- Machine learning is extremely powerful for characterizing numerical data
- Lower level inputs work better
- ML trades physical insight for performance

1. The present

Björn Karlsson, MidJourney January, 2024

What do high-energy theorists actually do?

1. Most papers on hep-th and hep-ph are largely symbolic

- Calculate something
- Find/establish some relationship
- Solve some toy model
- Extract behavior of some theory in some limit

Can ML do these things?

- It's starting to...
- ML is good at helping solve well-defined problems

2. What makes a question interesting?

- It connects to nature
- You can make progress on it
- Someone else thought it was interesting
- It is related to something someone else thought was interesting

Can ML answer this?

Not yet. But soon.

- Much harder problem

This conference

Most uses of ML in physics so far are **data science**

- Machine learning is great at characterizing numerical data
- Has led to revolutionary progress on a great variety of physical questions
- Much much more still to be done

This talk

How do we transition from **data science** to **symbolic problems?**

- Large Langauge Models show that ML good for symbolic problems
  - Potential is there
- A lot of what physicists do is study examples, look for patterns, and generalize
- A place to start is with *discrete* symbolic data
  - Find some theoretical physics problems where ML can help
  - Get a feel for what it can (and currently cannot) do

1. Mathematics

Pure mathematics is almont entirely focused on symbolic problems

- Automated proof assistants are improving fast

Lean [de Moura et al., 2015]

Coq [Barras et al., 1997]

Isabelle [Nipkow et al., 2002]

https://machine-learning-for-theorem-proving.github.io/ (NeurIPS 2023)

- Progress in **autoformalization**: using LLMs to translate natural language into formal mathematics (e.g. input to LEAN)

- FunSearch: use LLMs to write code to solve math problems

Article | Open access | Published: 14 December 2023

Mathematical discoveries from program search with large language models

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan

- These tools haven't seen much application in physics, but could soon...

2. Simplifying polylogarithms

[Dersy, MDS, Zhang, 2206.04115]

An important and challenging step in the computation of Feynman diagrams is to simplify polylogarithmic expressions

$$f(x) = 9\left(-\mathrm{Li}_3(x) - \mathrm{Li}_3\left(\frac{2ix}{-i+\sqrt{3}}\right) - \mathrm{Li}_3\left(-\frac{2ix}{i+\sqrt{3}}\right)\right)$$
$$+4\left(-\mathrm{Li}_3(x) + \mathrm{Li}_3\left(\frac{x}{x+1}\right) + \mathrm{Li}_3(x+1) - \mathrm{Li}_2(-x)\ln(x+1)\right)$$
$$-4\left(\mathrm{Li}_2(x+1)\ln(x+1) + \frac{1}{6}\ln^3(x+1) + \frac{1}{2}\ln(-x)\ln^2(x+1)\right)$$

1. What is its simplest form?
2. Does it simplify to zero?
3. What identities do we apply in what order to simplify it?

- No classical algorithm for this problem
- Easy to make more complicated expressions from simple ones by scrambling
  - Like cooking an egg: easy to scramble, hard to unscramble

We generate symbolic data and then look for patterns to simplify

Simplifying polylogarithms

Two machine learning approaches

Reflection	*Inversion*	*Duplication*
$-7\mathrm{Li}_2(x) \to 7\mathrm{Li}_2(1-x)$	$-7\mathrm{Li}_2(x) \to 7\mathrm{Li}_2\left(\frac{1}{x}\right)$	$-7\mathrm{Li}_2(x) \to 7\mathrm{Li}_2(-x) - \frac{7}{2}\mathrm{Li}_2(x^2)$

1. Reinforcement learning

Basic idea:

- apply known polylog identities like moves in a game
- train by learning to reverse scrambling steps

2. Transformer networks

- Used by large langauge models
- Learn to ``guess" answer
  translate from complicated to simple

Dutch: naamsveranderingsdocumentenbriefgeheel

$$f(x) = 9\left(-\mathrm{Li}_3(x) - \mathrm{Li}_3\left(\frac{2ix}{-i+\sqrt{3}}\right) - \mathrm{Li}_3\left(-\frac{2ix}{i+\sqrt{3}}\right)\right)$$
$$+4\left(-\mathrm{Li}_3(x) + \mathrm{Li}_3\left(\frac{x}{x+1}\right) + \mathrm{Li}_3(x+1) - \mathrm{Li}_2(-x)\ln(x+1)\right)$$
$$-4\left(\mathrm{Li}_2(x+1)\ln(x+1) + \frac{1}{6}\ln^3(x+1) + \frac{1}{2}\ln(-x)\ln^2(x+1)\right)$$

translate

English: dossier

$$f(x) = -\mathrm{Li}_3(x^3) - \mathrm{Li}_3(x^2) + 4\zeta_3$$

Results

- Both methods work well (>80% success)
- Transformers do better: **91% success** up to transcendental weight 4

3. Simplifying spinor-helicity amplitudes

[Cheung, Dersy, MDS, in perparation]

The same basic RL or transformer approach can be applied to many problems

- Need to be able to generate training data
- Have identities to apply

e.g: spinor helicity amplitudes

simplify this:

$$-\frac{[23]\,[35]\,[46]^4}{\begin{array}{l}(\langle 12\rangle\langle 15\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[23]\,[26]\,[35]\,[56]^2 + \langle 12\rangle\langle 15\rangle\langle 24\rangle^2\langle 26\rangle\langle 34\rangle\langle 35\rangle\,[23]\,[35]\,[36]\,[56]^2 \\ \quad + \langle 14\rangle\langle 15\rangle\langle 23\rangle\langle 24\rangle^2\langle 26\rangle\langle 35\rangle\,[23]\,[35]\,[36]\,[56]^2 + \langle 14\rangle\langle 15\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[13]\,[23]\,[45]\,[56] \\ \quad - \langle 14\rangle\langle 15\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[13]\,[24]\,[35]\,[56] + \langle 14\rangle\langle 15\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[15]\,[23]\,[34]\,[56] \\ \quad + \langle 14\rangle\langle 15\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[23]\,[35]\,[46]\,[56]^2 + \langle 14\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle^2\,[23]\,[34]\,[35]\,[56] \\ \quad + \langle 14\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\langle 45\rangle\,[24]\,[34]\,[35]\,[56] - \langle 14\rangle\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\langle 56\rangle\,[24]\,[35]\,[36]\,[56] \\ \quad + \langle 15\rangle^2\langle 24\rangle^3\langle 26\rangle\langle 35\rangle\,[23]\,[35]\,[56]^3)\end{array}}$$

Possible "moves":

Schouten identity

$$\langle ij\rangle \to \frac{\langle il\rangle\langle kj\rangle}{\langle kl\rangle} + \frac{\langle ik\rangle\langle jl\rangle}{\langle kl\rangle}$$

Multiply by 1

$$\frac{\langle p_1p_2\rangle[p_3p_4]}{\langle p_1p_3\rangle} \to \frac{\langle p_1p_2\rangle[p_3p_4]}{\langle p_1p_3\rangle}\,\frac{\langle ij\rangle}{\frac{\langle il\rangle\langle kj\rangle}{\langle kl\rangle} + \frac{\langle ik\rangle\langle jl\rangle}{\langle kl\rangle}}$$

Momentum conservation

$$\langle ij\rangle \to -\sum_{t\neq j}\frac{\langle it\rangle[tk]}{[jk]}$$

Add 0

$$\langle mk\rangle + [ln] \to \langle mk\rangle + [ln] + \left(\langle ij\rangle - \frac{\langle il\rangle\langle kj\rangle}{\langle kl\rangle} + \frac{\langle ik\rangle\langle jl\rangle}{\langle kl\rangle}\right)$$

Generate training data by scrambling

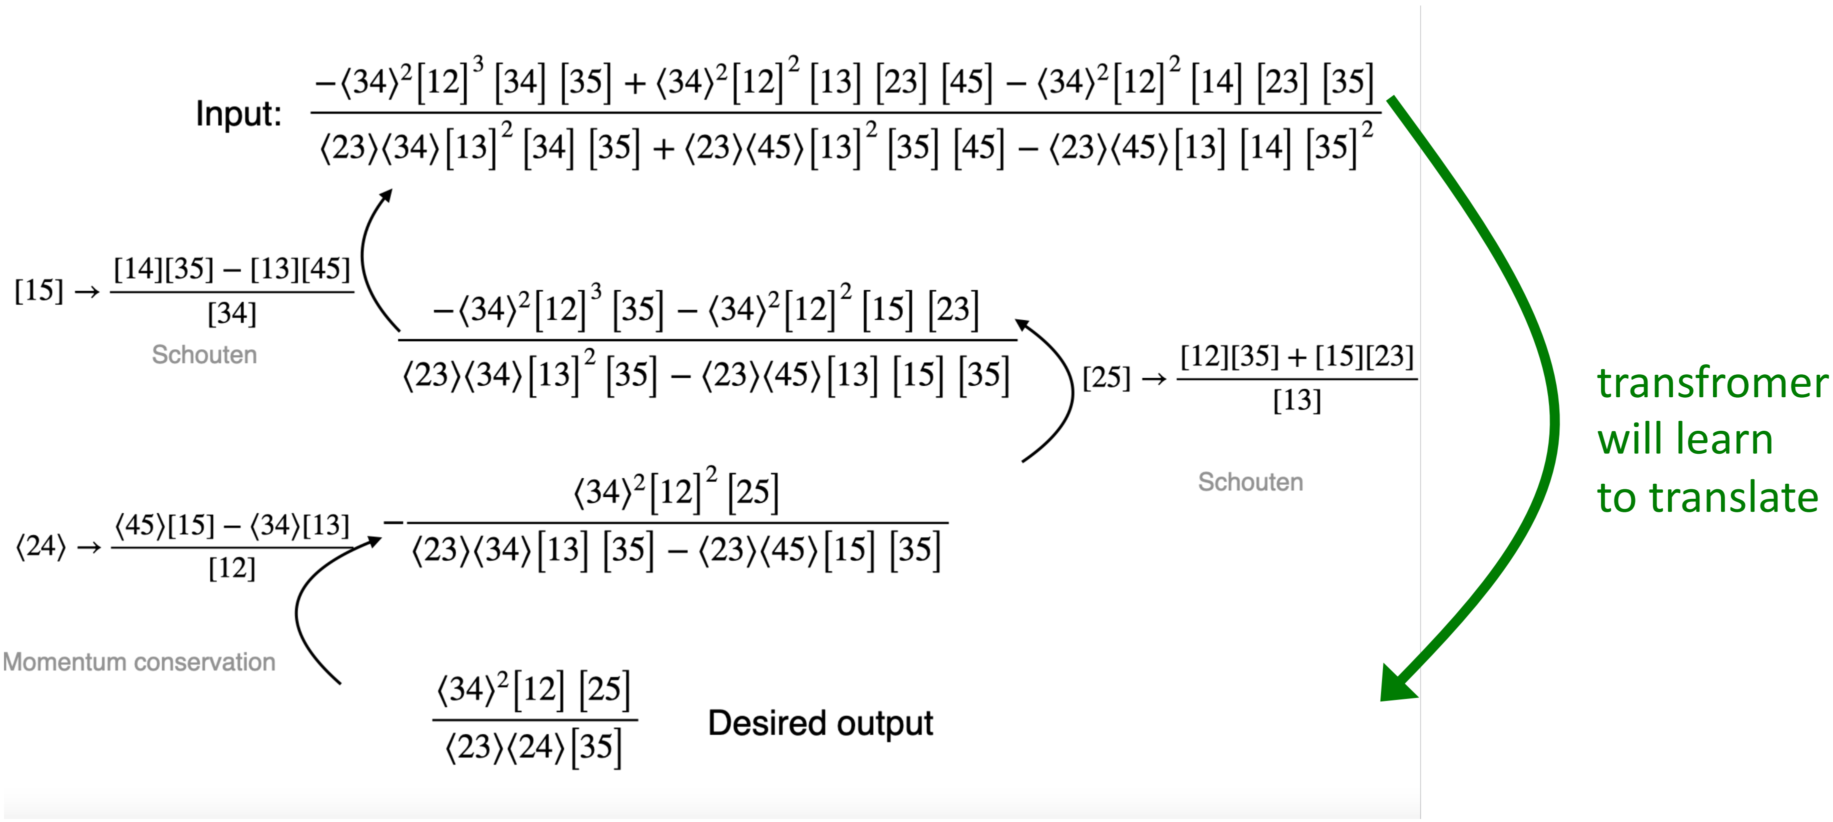

- Expressions can be very long: 50 terms or more
- Need new techniques for organizing transformer
  - We use **contrastive learning**

Contrastive learning

Learn an embedding so that terms that are similar are close

- similar = appear in some identity
- close = metric on embedding space

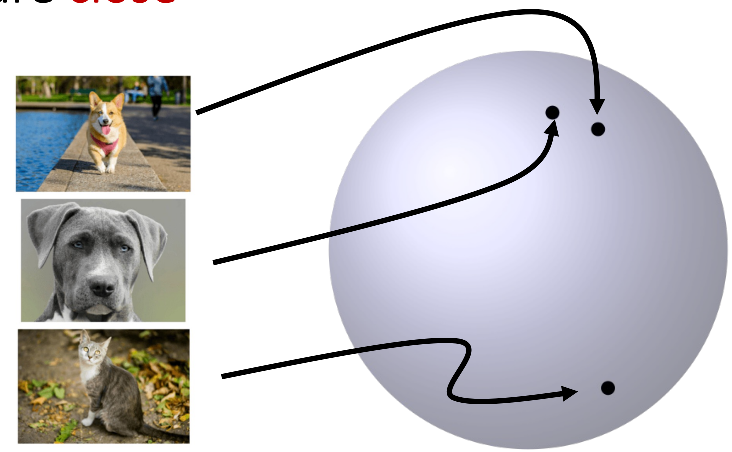

learned mapping

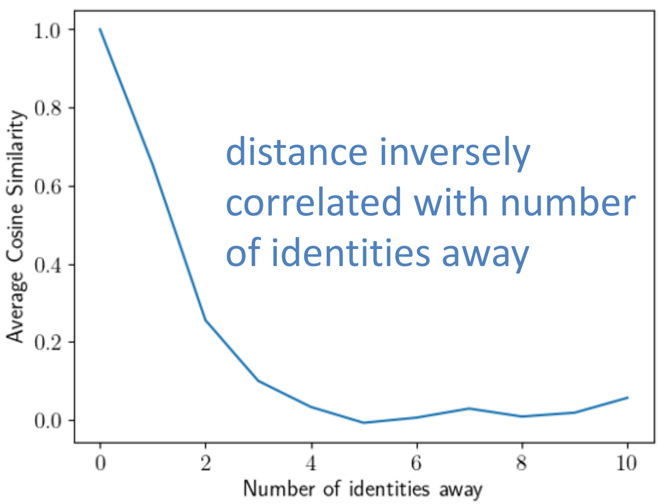

t-SNE on latent space

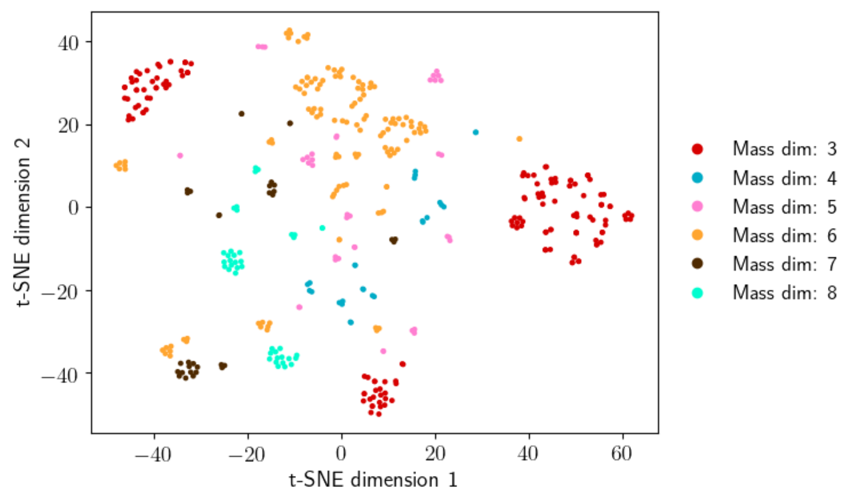

- Learns dimensional analysis
- Learns other features as well

Example application

- Compute 5-point MHV amplitude with Feynman diagrams: **79 terms (1990 tokens)**

$$\frac{\langle 12\rangle^2\langle 15\rangle^2\langle 24\rangle\langle 34\rangle[12][14][15][23][25]+\langle 12\rangle^2\langle 15\rangle\langle 23\rangle\langle 34\rangle\langle 45\rangle[12][15][23][25][34]+\cdots 77\ \text{terms}}{\langle 15\rangle^2\langle 23\rangle\langle 34\rangle^2\langle 45\rangle^2[12]^2[15][23][45]}$$

- Feed to network

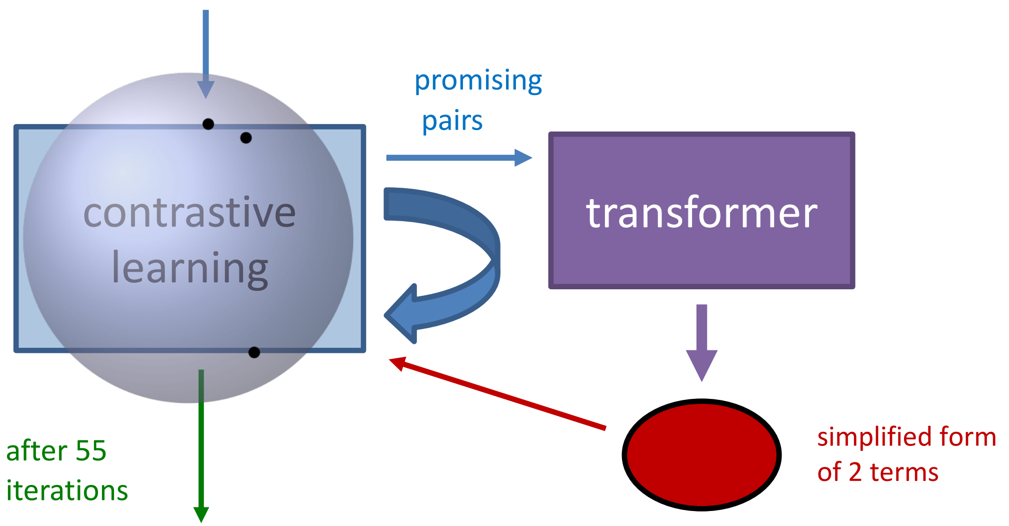

$$-\frac{\langle 12\rangle^3}{\langle 15\rangle\langle 23\rangle\langle 34\rangle\langle 45\rangle}$$

- Output is 1 term (27 tokens)

4. S-Matrix bootstrap

The S-matrix is the fundamental object of Quantum Field Theory

- A lot is known about it from pertubation theory (Feynman diagrams)
- Some things are known/conjectured about it non-perturbatively
  - e.g. it should be *unitary* which implies the optical theorem

Optical theorem

$$\mathrm{Im}\langle a|\mathcal{M}|b\rangle = \sum_{|c\rangle} \langle a|\mathcal{M}^\dagger|c\rangle\langle c|\mathcal{M}|b\rangle$$

- Non-perturbative constraint
- Relates complex scattering amplitude M to real cross section $\sigma = M^\dagger M$

Given cross section σ=|M|² can the phase of M be uniquely determined?

proof:

$$S^\dagger S = 1$$

$$S = 1 + i\mathcal{M}$$

$$S^\dagger S = 1 + i\mathcal{M} + i\mathcal{M}^\dagger + \mathcal{M}^\dagger\mathcal{M} = 1$$

Elastic scattering

Unitarity constraint simplifies in the "elastic regime"

- $4\,m^2 < s < 9\,m^2$
- Only 2-particle states are relevant
- energy conserved, kinematics described by scattering angle $z = \cos\theta$

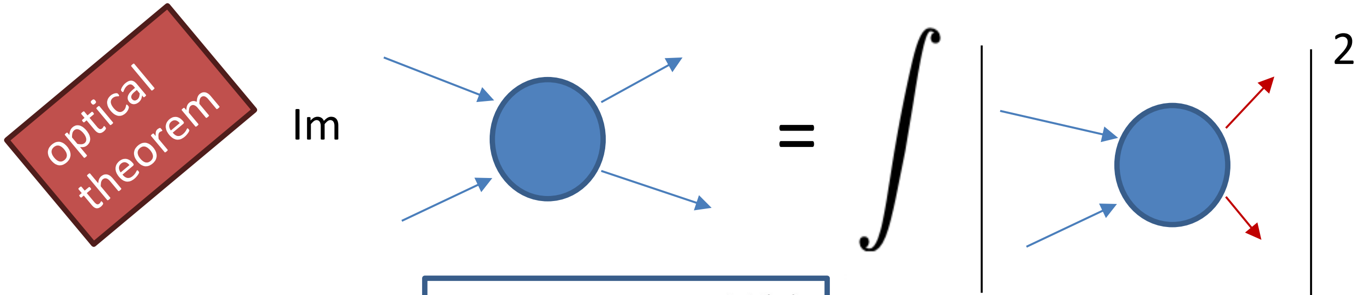

Write the amplitude as $\mathcal{M}(z) = B(z)e^{i\phi(z)}$

Unitarity constraint $\mathrm{Im}\mathcal{M}(z) = \int d\Pi\,|B|^2$

$$\Rightarrow \quad B(z_{12})\sin\phi(z_{12}) = \frac{1}{4\pi}\int d\Omega_3 B(z_{13})B(z_{23})\cos(\phi(z_{13}) - \phi(z_{23}))$$

$$z_2 = z z_1 + \sqrt{1-z^2}\sqrt{1-z_1^2}\cos\phi_1$$

$$\Rightarrow \quad \sin\phi(z) = \frac{\int_{-1}^{1} dz_1 \int_0^{2\pi} d\phi_1 B(z_1)B(z_2)e^{i\phi(z_1)-i\phi(z_2)}}{4\pi B(z)}$$

optical theorem

For what B(z) does ϕ(z) exist satisfying this equation?
For what B(z) is ϕ(z) unique or not-unique?

Can we find ϕ(z) given B(z) with ML? ... Yes!

[Dersy, MDS, Zhiboedov, 2308.09451]

- Parametrize ϕ(z) as a neural network

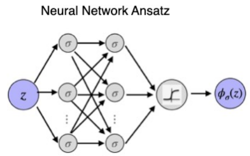

- Loss function is unitarity condition

$$\mathscr{L} = \mathbb{E}\left|\left|B(z)\sin\phi(z) - \frac{1}{4\pi}\int_{-1}^{1} dz_1 \int_0^{2\pi} d\phi_1 B(z_1)B(z_2)\cos\left(\phi(z_1)-\phi(z_2)\right)\right|\right|^2$$

some known examples

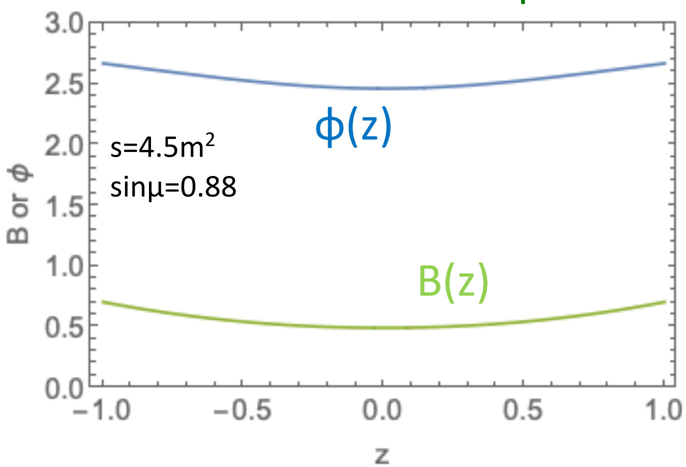

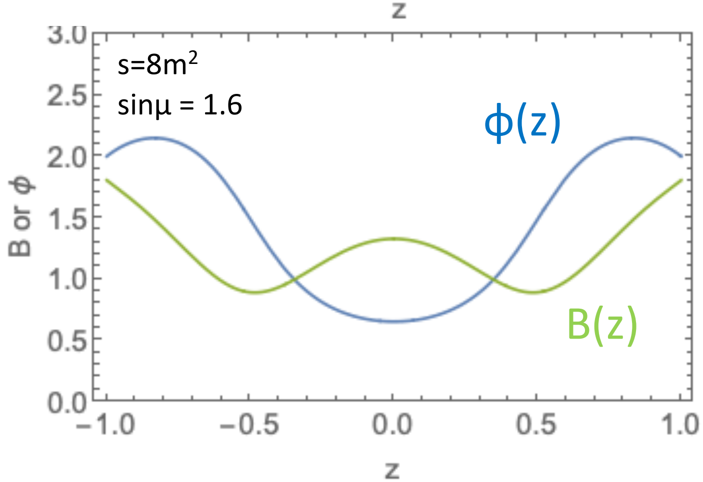

find ϕ given B using ML

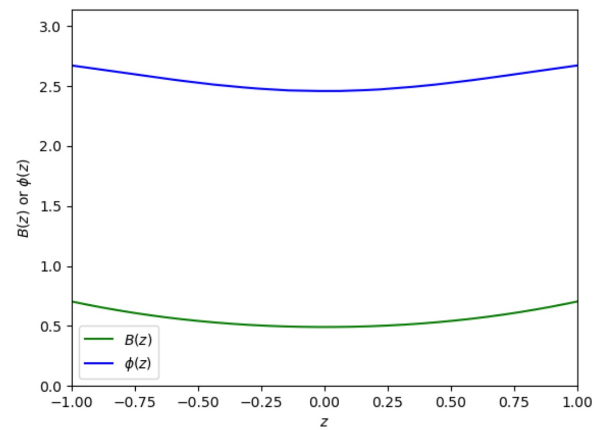

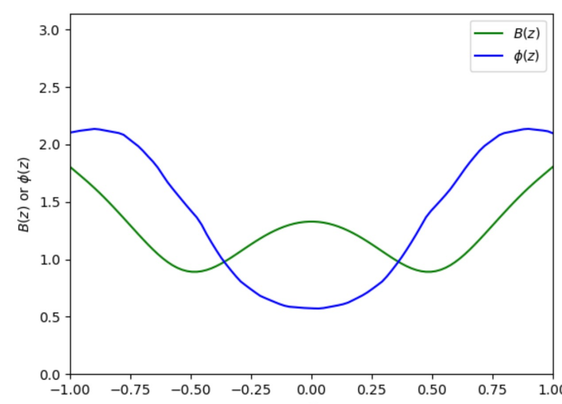

excellent agreement with known results

When does a phase exist?

In 1967 Andre Martin proposed

$$\sin\mu = \max_z \frac{\int_{-1}^{1} dz_1 \int_0^{2\pi} d\phi_1 B(z_1) B(z_2)}{4\pi B(z)}$$

as an indicator of phase determination

- Martin proved:
  - If sinμ < 1 for a given B(z) then there always exists a phase ϕ(z)

What is special about sinμ?

Contours of sin μ

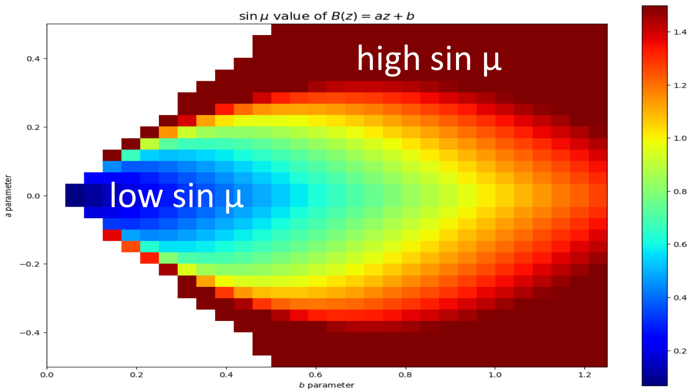

Loss landscape from ML search for ϕ

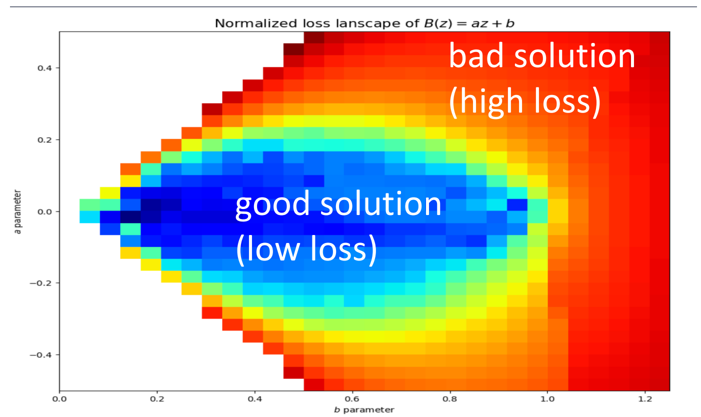

- Loss landscape correlates with sin μ
  - sin μ measures how hard these solutions are to find
  - Don’t need exact solutions to learn this lesson

Can there be more than one ϕ given B?

Crichton (1966): yes

Atkinson (1977) found two phases $\phi_1(z)$ and $\phi_2(z)$ for the same B(z) with $\sin\mu > 2.15$

Open question since 1977
Are there phase-ambiguous amplitudes with $\sin\mu < 2.15$?

ML approach: two NNs for two phases $\phi_1(z)$ and $\phi_2(z)$

$\phi_1(z)$ =

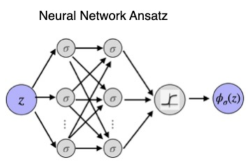

$\phi_2(z)$ =

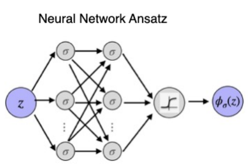

- Impose unitarity condition loss for each ϕ:

$$\mathscr{L} = \mathbb{E}\left|\left|B(z)\sin\phi(z) - \frac{1}{4\pi}\int_{-1}^{1} dz_1 \int_0^{2\pi} d\phi_1 B(z_1)B(z_2)\cos\left(\phi(z_1) - \phi(z_2)\right)\right|\right|^2$$

- Add repulsive loss to keep solutions apart

$$\mathcal{L}_R = \mathbb{E}_z\,||d\left(\phi_1(z), \phi_2(z)\right)||^{-p} + \mathbb{E}_z\,||d\left(\pi - \phi_1(z), \phi_2(z)\right)||^{-p}$$

Gradient descent in sin μ leads to new phase-ambiguous cross sections

First new phase-ambiguous solution in 50 years!

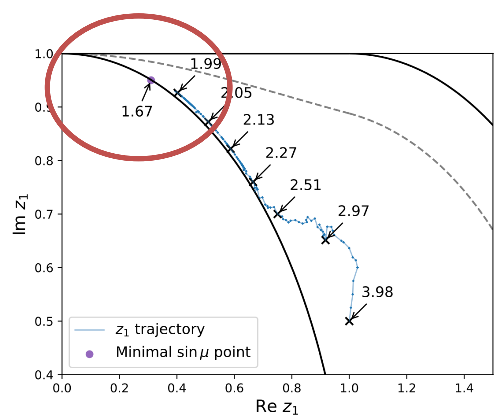

New lowest sinμ phase ambuity

sinμ = 1.67

5. String theory

String theory is still trying to find the Standard Model among its infinite vacuua

Calabi proved metrics exist and are unique
- Not constructive: no non-trivial analytic Calabi-Yau metrics are known!

one approach:

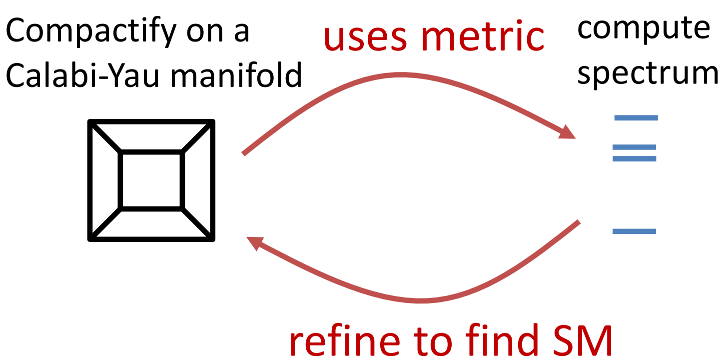

Ashmore, He, Ovrut arXiv:1910.08605
- use ML to find metrics numerically
- Loss function is Ricci-flatness

Halverson and Ruehle arXiv:2310.19870
- NNs to look for fixed points of metric flows

$$\frac{\mathrm{d}g_{ij}(x)}{\mathrm{d}t} = -2\,R_{ij}(x),$$

Carifio et al arXiv:1707.00655
- Trained models to predict rank of gauge group in F-theory compactifications
- Patterns led to conjecture, which was then proven by humans

Recent review

nature reviews physics

April 2024

https://doi.org/10.1038/s42254-024-00709-0

Sergei Gukov[1,2], James Halverson[3,4] & Fabian Ruehle[3,4,5]

3. The future

Markus Graf, July, 2023

Large Language Models are the (immediate) future

LLMs

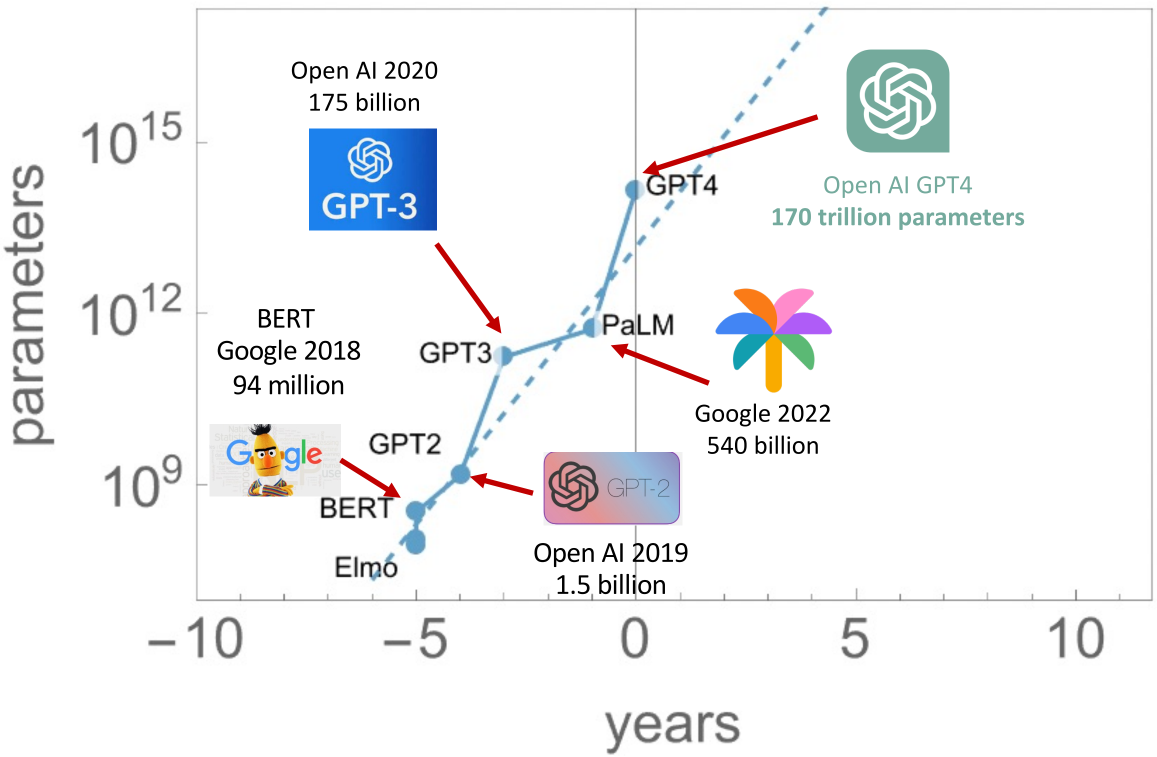

Mammalian brains

Cat brain

0.760 billion neurons
10 trillion synapses

size of GPT 3.5

Human brain

80 billion neurons
150 trillion synapses

size of GPT 4

- Current LLMs have roughly the same number of parameters (10^{14}) as the human brain
- And more compute:
  - Brain (10^{16} FLOPS) over a lifetime (100 years) = 10^{22} operations
  - LLM training time = 10^{25} operations

Machine vs. Biological intelligence

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

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

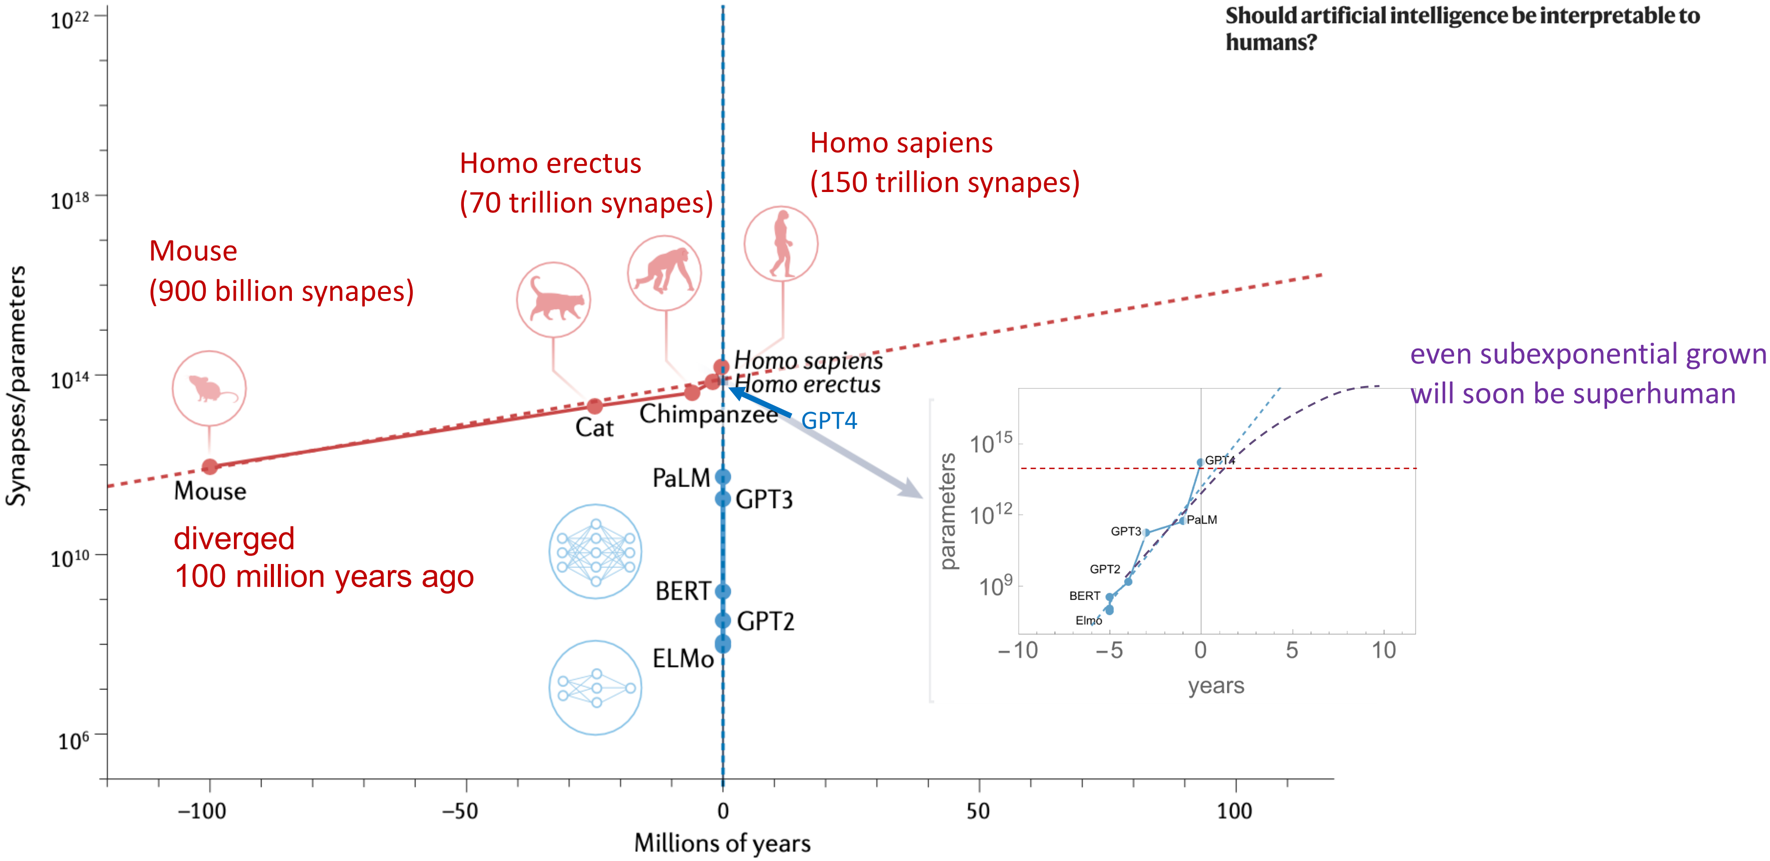

- Both AI and biological intelligence grow exponentially
- **Factor of 10^6 difference in *exponent***
- Intersection, when machines and biology have comparable "intellegence" is now

Will we run out of data/compute/energy?

ALGORITHMIC PROGRESS IN LANGUAGE MODELS

Ho et al. arXiv:2403.05812

size is compute

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

log perplexity = loss

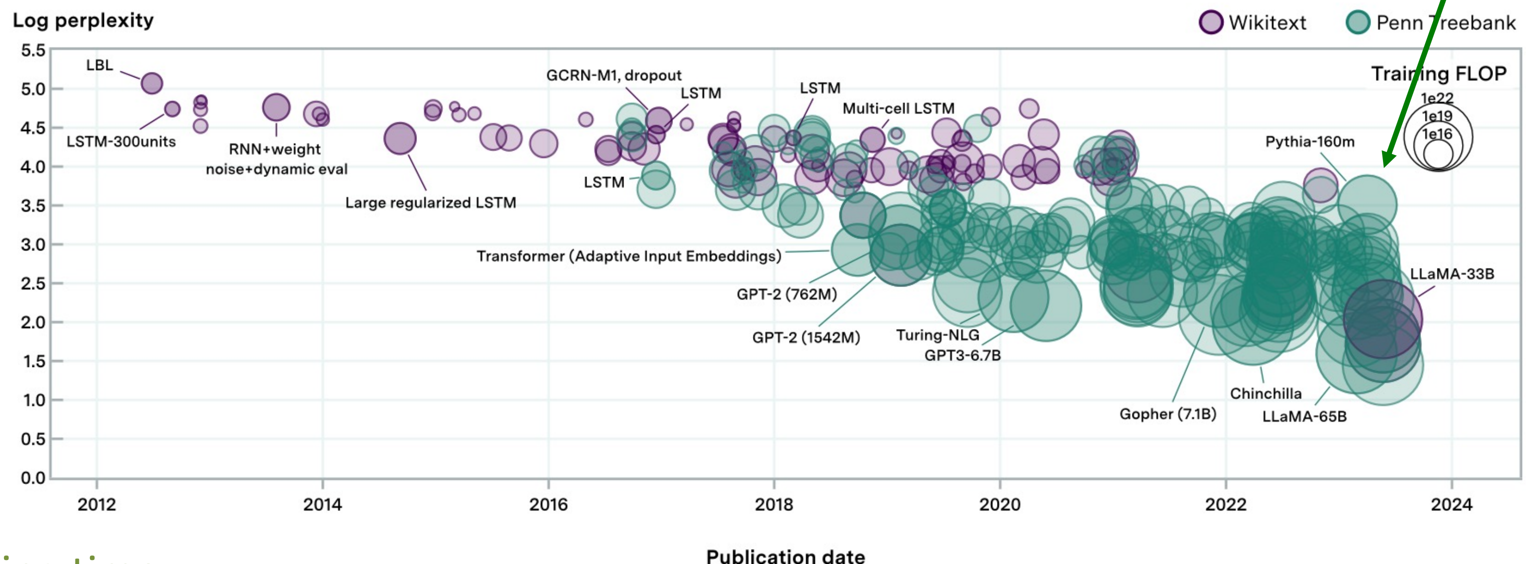

doubling time

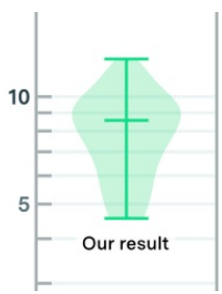

algorithmic doubling time = 6 to 14 months!

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

But physics requires creativity!?

The originality of machines: AI takes the Torrance Test

Erik E. Guzik [a,*], Christian Byrge [b], Christian Gilde [c]

[a] University of Montana, College of Business, 32 Campus Drive, Missoula, MT 59812, United States
[b] Vilnius University Business School, Lithuania
[c] University of Montana Western, Business Department, United States

Torrence Test of Creative Thinking (1960)

- How many new ways can you think to use a water bottle?
- Suppose you could be invisible for a day. What problems might that create? What would the benefits of being invisible be?

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

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

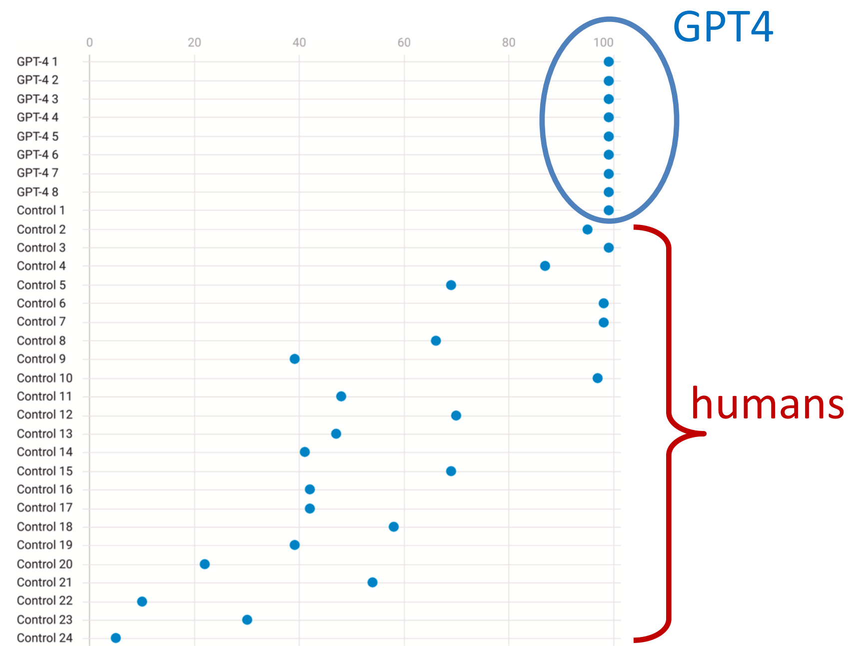

Fig. 3. Originality National Percentile Ranks (GPT-4 and Control Group).

Torrence Test score

GPT4 more creative than 99% of humans

Try it with GPT 4

MA **You**

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

ChatGPT

A water bottle can be versatile! Here are some creative ways to use one:

1. **Drinking Container:** The most obvious use, for carrying water or other beverages.
2. **Plant Waterer:** Use it to water plants. The bottle can be especially handy for targeting small plants or hard-to-reach areas.
3. **Weight for Exercise:** Fill it with water, sand, or pebbles and use it as a makeshift dumbbell.
4. **Ice Pack:** Fill it with water and freeze it. You can use it to keep food cool in a cooler or as an ice pack for injuries.
5. **Bird Feeder:** Cut it into a bird feeder and fill it with bird seed.
6. **Storage Container:** Use it to store dry goods, nuts, bolts, or coins.
7. **Funnel:** Cut the top off and use it as a funnel for liquids or small grains like rice or birdseed.
8. **Piggy Bank:** Decorate it and cut a slot to make a homemade piggy bank.
9. **Sensory Bottle:** Fill it with glitter, beads, and water (or baby oil) to make a calming sensory bottle.
10. **DIY Sprinkler:** Poke holes in it, attach to a hose, and use it as a lawn or garden sprinkler.
11. **Lamp or Light:** Insert a string of LED lights inside a frosted bottle to create a makeshift lamp.
12. **Time Capsule:** Fill it with memorable items and bury it or store it as a time capsule.
13. **Shaker for Music:** Fill with beads or beans to make a DIY musical shaker.

Come up with 20 more

1. **Mobile Phone Holder:** Cut it to the right size to make a holder or stand for your mobile phone.
2. **Yarn Dispenser:** Use it to keep yarn tangle-free while knitting or crocheting; just feed the end through a hole in the cap.
3. **Pet Toy:** Fill with small treats and cut holes just big enough for them to fall out as your pet plays with it.
4. **Vase:** Decorate it and use it as a vase for flowers.
5. **Rice Roller:** Fill with rice or sand and use it to flatten or compact soil in trays.
6. **Sand Art Container:** Layer colored sand inside for a decorative sand ar
7. **Drip Irrigation System:** Pierce small holes and use it to slowly irrigate p
8. **Protective Packaging:** Cut into rings and use as cushioning for packin
9. **Craft Supply Organizer:** Sort and store beads, buttons, or other small
10. **Mini Greenhouse:** Cut it in half and use the top as a cover over seedli greenhouse.
11. **Water Filtration Experiment:** Demonstrate water filtration by layering charcoal inside.
12. **Kitchen Dispensers:** Use for storing and dispensing bulk spices or gr
13. **Pencil Holder:** Decorate and use it to store pens, pencils, and other o
14. **Portable Shower:** Poke holes in the cap, fill with water, and use it as a camping.
15. **Boot Shaper:** Fill with water and freeze to use as a shaper to keep boo
16. **Hand Washing Station:** Set up with soapy water and a towel for outdo
17. **Art Palette:** Cut and use the bottom as alette for mixing paints.
18. **Bubble Blower:** Fill with a soap solution and use for blowing bubbles.
19. **Salt Spreader:** Fill with salt and poke holes in the lid to use for de-icing dri sidewalks in winter.
20. **Dog Poop Bag Dispenser:** Store and dispense dog poop bags convenientl

MA

You

10 more. be very creative

1. **Homemade Barometer:** Use it to create a simple weather instrument. Seal partway filled with colored water, and observe the water level changes to predict weather changes.
2. **Solar Water Heater:** Paint it black, fill with water, and leave it in the sun. It can heat up water for washing dishes or hands when camping.
3. **Hologram Projector:** Cut the top off and use a smartphone to project a hologram. There are plenty of online tutorials to turn a cut bottle into a DIY 3D hologram projector.
4. **DIY Kayak Outrigger:** Connect two large water bottles on either side of a small kayak or canoe with rods for added stability.
5. **Miniature Snow Globe:** Fill it with glycerin, water, glitter, and tiny waterproof figures to create a personalized snow globe.
6. **Interactive Cat Toy:** Secure a feather or toy to the top of the bottle. Cats can bat it around and enjoy the spinning motion.
7. **Refillable Paint Roller:** Fill with paint and attach a modified roller on top; it can simplify painting by integrating the paint container and applicator.
8. **Zero Gravity Coffee Brewer:** Inspired by space technology, use it to design a coffee brewer that can operate in zero gravity, using a system of internal chambers for brewing and dispensing.
9. **Underwater Camera Housing:** Waterproof and modify it to hold a small camera for underwater photography in shallow waters.
10. **Wind-Powered Light:** Attach blades to a water bottle and mount it on a pole with a small LED light inside. The wind rotates the bottle, generating enough power to light the LED.

Augmented intelligence: LLMs help us thrive

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

Can LLMs help consultants?

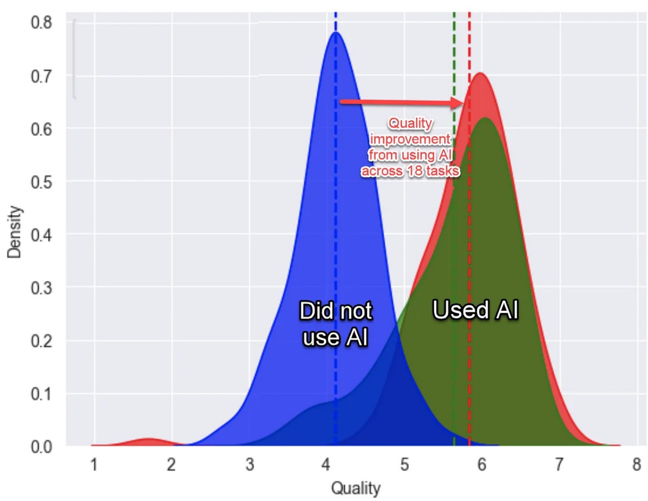

Distribution of output quality across all the tasks. The blue group did not use AI, the green and red groups used AI, the red group got some additional training on how to use AI.

Yes!
By a lot!

AI is a “skill-leveler”

bottom half improved 43%

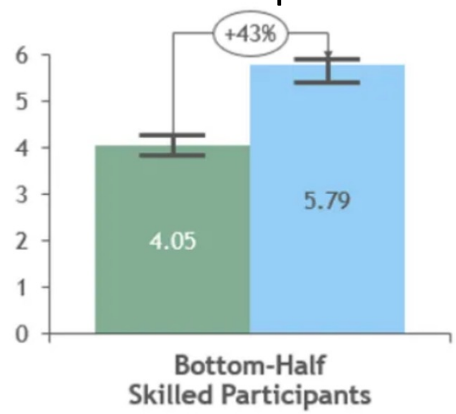

top half improved 17%

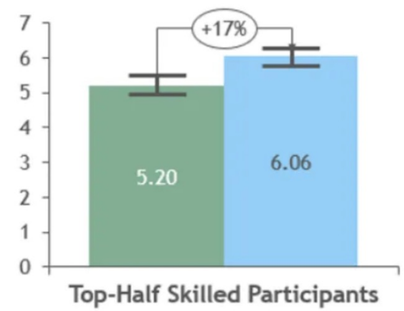

- Average consultants using AI become as good as the best consultants

10,000 Einsteins

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.

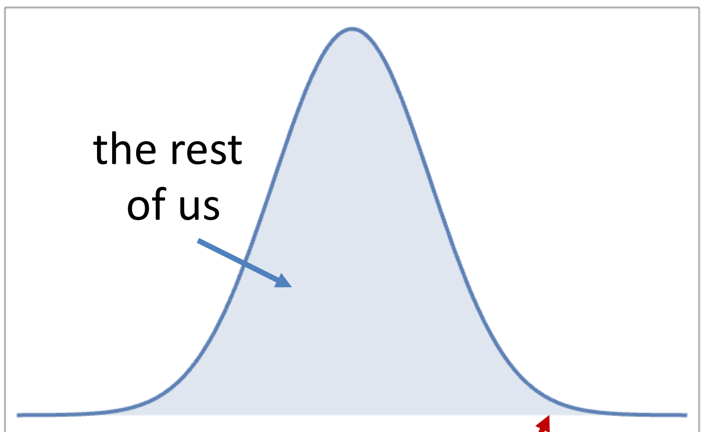

exceptional
theoretical physicsts

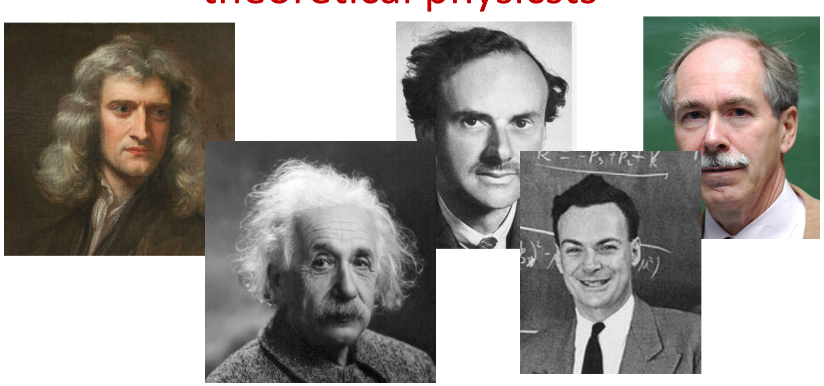

Can AI be a skill-leveler for high-energy theory?

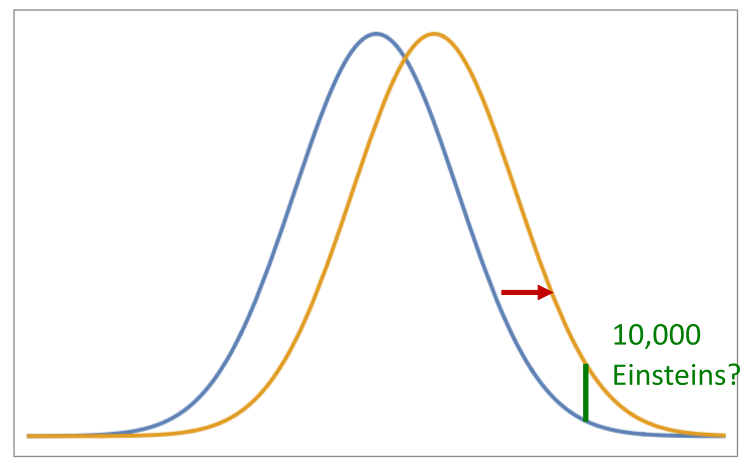

- LLMs are already helping

Theoretical High Energy Physics may have stalled

In the past, we made progress depsite many dead ends

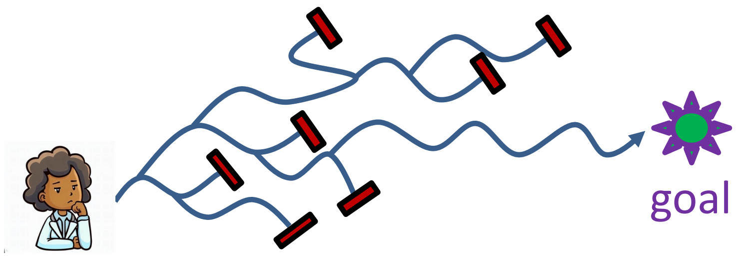

Are we even making forward progress anymore?

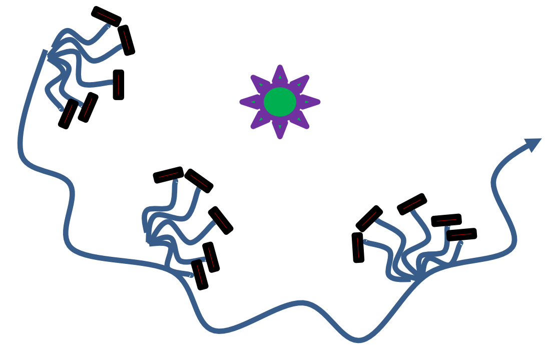

Maybe the problems are just too difficult (for us)

Could a cat ever learn to play chess?

- Humans have limits too

Humans are limited by biology

Humans like to "visualize"

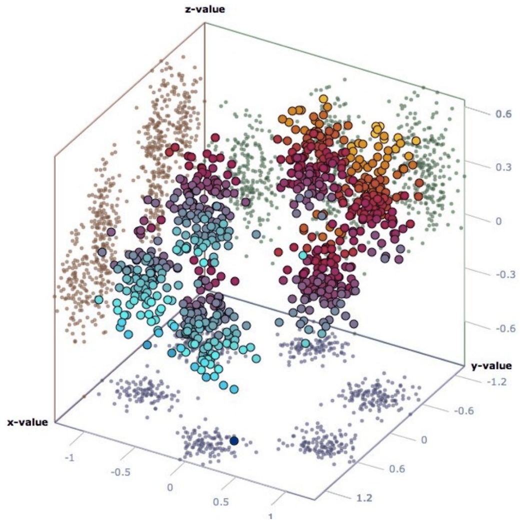

project
to 2D

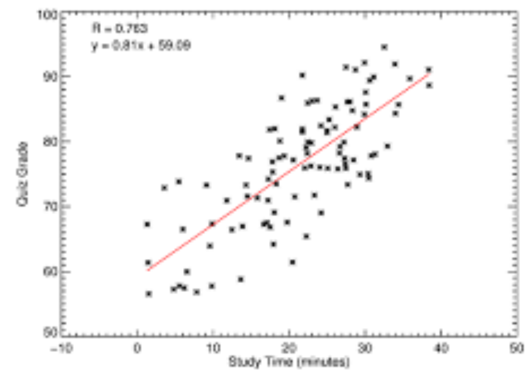

Why do we do this? Because we have eyes

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

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

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

- We like simple-looking equations

$$i\partial_t\psi = H\psi \qquad i\,\not{\partial}\psi = m\psi \qquad G_{\mu\nu} = \kappa T_{\mu\nu}$$

- **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 get there?

Current state-of-the art can solve textbook physics problems

- Trained on solved problems from books, physics.stackexchange, chegg etc.
- Books, chegg, etc. written by human beings who read books, chegg, etc.
- i.e. we generate our own training data

Graduate school is largely about learning how to train yourself

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

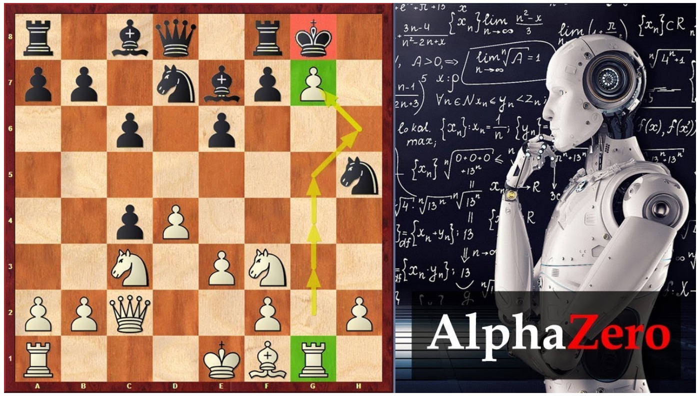

Current LLMs

- can generate and solve problems
- user (human) feedback helps refine model
- it can refine its own model!
- G Ph. D?

Language models are vey close to training themselves to be better physicists

Beyond augmented intelligence

Suppose a machine understands the theory of everything but we don't
- e.g. can calculate the fine-structure constant from scratch
- e.g. can preduct the endpoint of black-hole evaporation

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

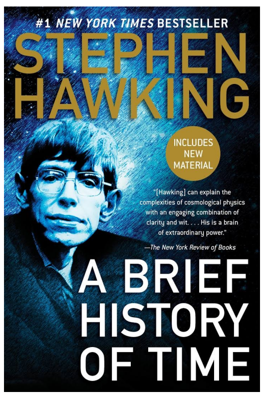

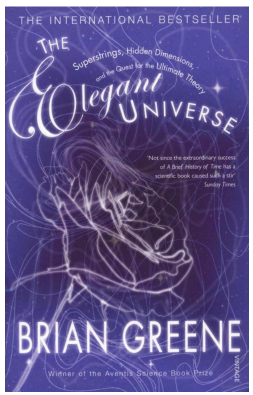

- The authors of **Popular science books** understand the details; we just get the general idea

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

- Machine learning is **rapidly tranforming high energy physics**
  - Current revolution in applications and advances are in “data science”
  - In hep-th and hep-ph problems are largely symbolic

1. How do we transition from data science to symbolic theoretical physics?

- It will get easier once we get started
  - Symbolic search problems (polylogarithms, spinor helicity)
  - Properties of the S-matrix (unitarity)
  - String Theory Vacuua

  } searching for simplicity

2. Generative AI is the future

- Short term: augmented intelligence
  - Machines help us organize information
  - Smooth transition to arXAIv: more and more AI input into arXiv papers
- Long term: artificial intelligence
  - Machines will suggest problems, solve problems: G Ph. T
  - Machines will dumb things down, so we can appreciate their work
  - Superhard problem in theoretical physics may finally be solved