

Machine Learning for Amplitudes

Amplitudes 2022

August 10

Matthew Schwartz

Harvard University

Based on

arXiv:2206.04115

Simplifying Polylogarithms with Machine Learning

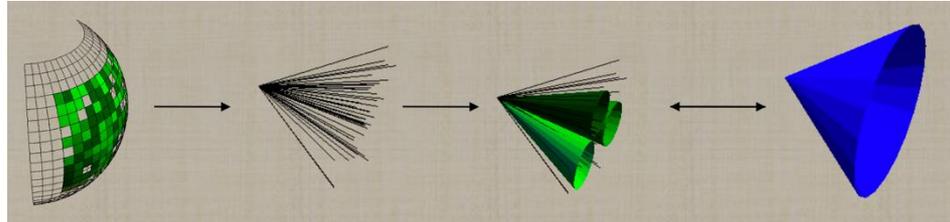
Aurélien Dersy, Matthew D. Schwartz, Xiaoyuan Zhang

Machine learning has transformed collider physics

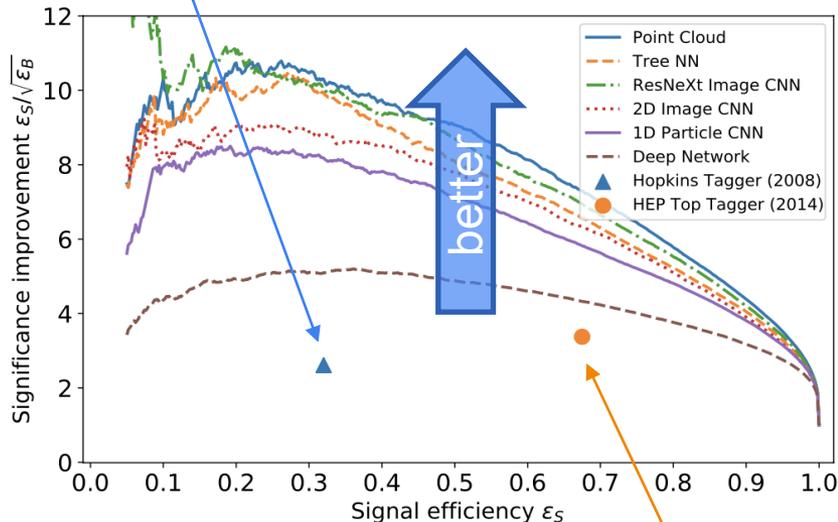
e.g. top tagging

Jet substructure approach (2008-2017):

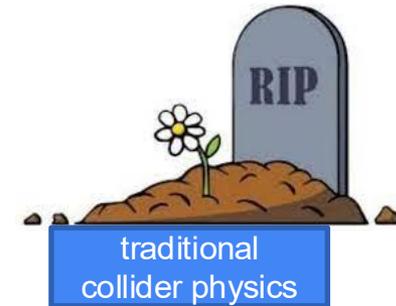
- Think about physics
- Deconstruct jet
- Look for W within top jet
- Look at helicity angle
- Hopkins Top Tagger (2008)



Machine learning methods are much better



best jet substructure algorithm (2014)



- ML requires less “thinking”
- Less physical insight
- Better performance

What about **the rest** of high energy physics?

What field ML **make obsolete** next?

- So far, most ML in physics is highly **numerical**
 - Collider physics applications involve millions of events
 - Data is real numbers
 - Approximate answers are ok
- Most hep-ph and hep-th papers are **symbolic**
 - Analytic understanding of some simplified system
 - Approximate but exact solutions to some equation
 - Loop calculations
 - Analytical computation of loops
 - Numerical implementation for precision physics

Symbolic ML methods will be
essential for the future of High Energy Physics

Inspiration 1 (Facebook AI 2019)

DEEP LEARNING FOR SYMBOLIC MATHEMATICS [arXiv:1912.01412]

Guillaume Lample*
Facebook AI Research
glample@fb.com

François Charton*
Facebook AI Research
fcharton@fb.com

ABSTRACT

Neural networks have a reputation for being better at solving statistical or approximate problems than at performing calculations or working with symbolic data. In this paper, we show that they can be surprisingly good at more elaborated tasks in mathematics, **such as symbolic integration and solving differential equations.** We propose a syntax for representing mathematical problems, and methods for generating large datasets that can be used to train sequence-to-sequence models. We achieve results that outperform commercial Computer Algebra Systems such as Matlab or Mathematica.

Equation	Solution
$y' = \frac{16x^3 - 42x^2 + 2x}{(-16x^8 + 112x^7 - 204x^6 + 28x^5 - x^4 + 1)^{1/2}}$	$y = \sin^{-1}(4x^4 - 14x^3 + x^2)$
$3xy \cos(x) - \sqrt{9x^2 \sin(x)^2 + 1}y' + 3y \sin(x) = 0$	$y = c \exp(\sinh^{-1}(3x \sin(x)))$
$4x^4yy'' - 8x^4y'^2 - 8x^3yy' - 3x^3y'' - 8x^2y^2 - 6x^2y' - 3x^2y'' - 9xy' - 3y = 0$	$y = \frac{c_1 + 3x + 3 \log(x)}{x(c_2 + 4x)}$

Inspiration 2 (Google DeepMind 2019)

ANALYSING MATHEMATICAL REASONING ABILITIES OF NEURAL MODELS [arXiv:1904.01557]

David Saxton
DeepMind
saxton@google.com

Edward Grefenstette
DeepMind
egrefen@fb.com

Felix Hill
DeepMind
felixhill@google.com

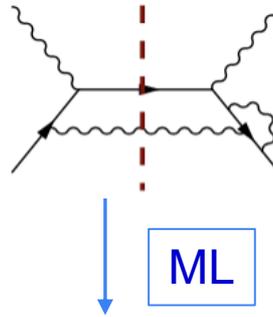
Pushmeet Kohli
DeepMind
pushmeet@google.com

D HIGH-SCHOOL MATHEMATICS QUESTIONS

1. Factorise $x^2 + 7x$
2. Factorise $y^2 - 10y + 16$
3. Factorise $2t^2 + 5t + 2$
4. Simplify $\frac{(x+1)}{2} + \frac{(x+3)}{3}$
5. Solve $2x^2 + 9x + 7$
6. Solve $\frac{2}{y^2} + \frac{9}{y} - 7 = 0$
7. Expand $3(x + 4) + 2(5x - 1)$
8. Expand $(2x + 1)(x - 4)$
9. Factor $6y^2 - 9xy$

- If this is the **state of the art...**
... bar is pretty low
- Let's try it!

Can we compute Feynman integrals with ML?



$$\begin{aligned} & [\text{Li}_2(1-z) - \text{Li}_2(1-z^{-1})] \ln(x+1) \\ & + \arctan\left(\sqrt{3}\frac{1-z}{1+z}\right) \left[\frac{2\pi}{3} \ln(x+1) + \frac{2(\psi'(\frac{1}{6}) - 2\pi^2)}{5\sqrt{3}} \right] \end{aligned}$$

Is this crazy?

- I don't think so. Already possible at 1-loop

Would it revolutionize physics?

- I don't think so. But it would accelerate progress.
- Kind of like having a better calculator



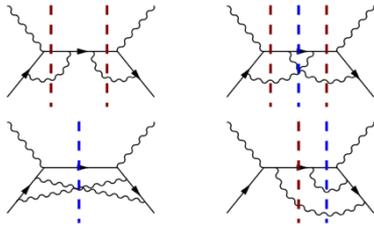
Baby step: recognize when two forms are equivalent
(necessary condition for next steps)

Application to HET: Simplifying amplitudes

e.g. Compton scattering at NLO [Lee, Schwartz, Zhang *Phys.Rev.Lett.* 126 (2021)]

[arXiv:2102.06718]

Want to compute some Feynman diagrams



Reduce to iterated integrals

$$I(s) = \int_1^s \frac{ds'}{\sqrt{(s'-1)(s'+3)}} \int_1^{s'} \frac{ds''}{s''}$$

Do integrals

- result is complicated
- simplifying result took ~ 1 month

$$\begin{aligned} I(s) &= \int_1^{r(s)} d \ln r' \int_1^{r'} d \ln \frac{r''^2 - r'' + 1}{r''} \\ &= \int_1^{r(s)} \frac{dr'}{r'} \int_1^{r'} dr'' \left(-\frac{1}{r''} + \frac{1}{r'' - z_1} + \frac{1}{r'' - z_2} \right) \\ &= \frac{2\pi^2}{9} - \frac{\ln^2 r}{2} + \left(-\frac{i\pi}{3} + \ln r \right) \ln(r - z_1) + \ln(r z_1) \ln(r - z_2) \\ &\quad - \text{Li}_2(1 - z_1) + \text{Li}_2(1 - r z_1) - \text{Li}_2(1 - z_2) + \text{Li}_2(1 - r z_2) \end{aligned}$$

$z_1 = \frac{1+i\sqrt{3}}{2}$
 $z_2 = \frac{1-i\sqrt{3}}{2}$

$$I(s) = -\frac{1}{3} \text{Li}_2(-r^3) + \text{Li}_2(-r) - \frac{1}{2} \ln^2 r + \frac{\pi^2}{18}$$

Why is simplification important?

- Removes unphysical singularities
- A lot of physics in analytic structure
- Simple form indicates deeper structure
- Simplification at intermediate steps make full calculation tractable
- In a sense, **all of science is simplification**

How do we simplify polylogarithms?

Logarithm is easy: only one identity

$$\ln(xy) = \ln x + \ln y$$

Dilogarithms have lots of identities

$$(inversion) \quad \text{Li}_2(x) = -\text{Li}_2\left(\frac{1}{x}\right) - \frac{\pi^2}{6} - \frac{\log^2(-x)}{2}$$

$$(reflection) \quad \text{Li}_2(x) = -\text{Li}_2(1-x) + \frac{\pi^2}{6} - \log(x) \log(1-x)$$

$$(duplication) \quad \text{Li}_2(x) = -\text{Li}_2(-x) + \frac{1}{2}\text{Li}_2(x^2)$$

$$(5-term identity) \quad \begin{aligned} \text{Li}_2(x) + \text{Li}_2(y) + \text{Li}_2\left(\frac{1-x}{1-xy}\right) + \text{Li}_2(1-xy) + \text{Li}_2\left(\frac{1-y}{1-xy}\right) \\ = \frac{\pi^2}{2} - \ln(x) \ln(1-x) - \ln(y) \ln(1-y) - \ln\left(\frac{1-x}{1-xy}\right) \ln\left(\frac{1-y}{1-xy}\right) \end{aligned}$$

- Understood most transparently with **the symbol** (more later)

Li_3 , Li_4 , etc. have identities too (complete set not known)

$$\text{Li}_3(x) = \text{Li}_3\left(\frac{1}{x}\right) - \frac{1}{6} \ln^3(-x) - \frac{\pi^2}{6} \ln(-x)$$

Problem statement:

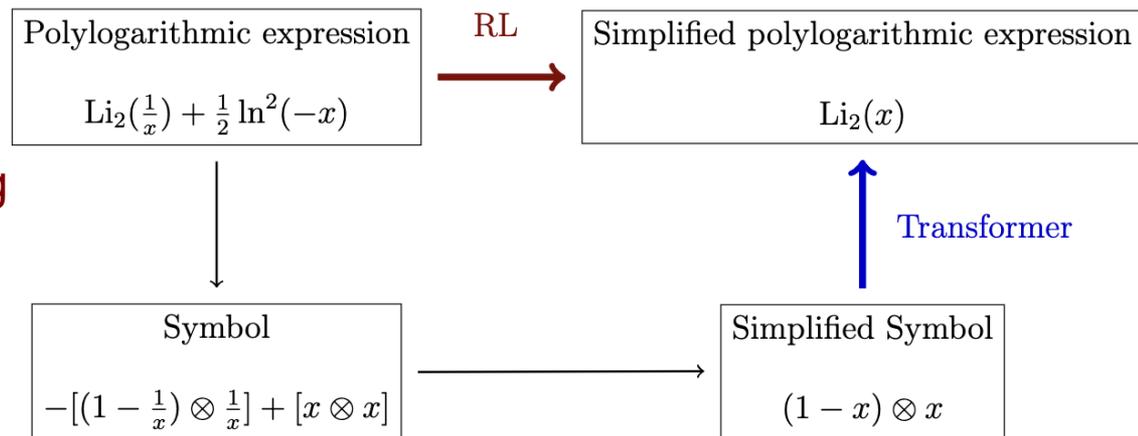
Given some polylogarithmic expression:

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

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

Two approaches

1. Reinforcement learning
2. Transformer networks

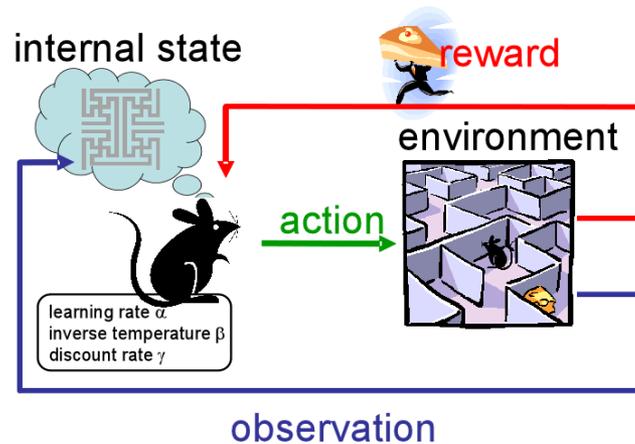
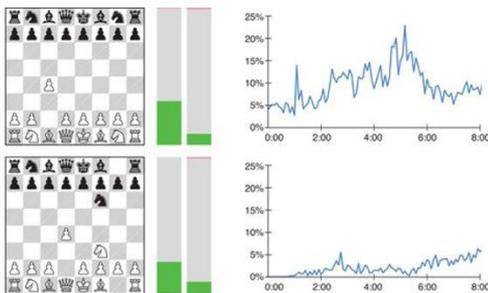


Reinforcement learning

- Applying identities can be viewed as moves in a game

<i>Reflection</i> 	<i>Inversion</i> 	<i>Duplication</i> 
$-7Li_2(x) \rightarrow 7Li_2(1 - x)$	$-7Li_2(x) \rightarrow 7Li_2\left(\frac{1}{x}\right)$	$-7Li_2(x) \rightarrow 7Li_2(-x) - \frac{7}{2}Li_2(x^2)$

Killer Ap: AlphaZero (Deepmind, 2017)



Tries action to take based on state

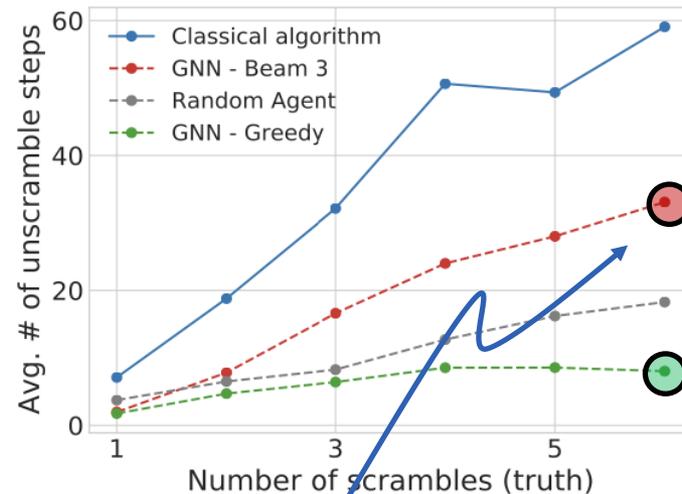
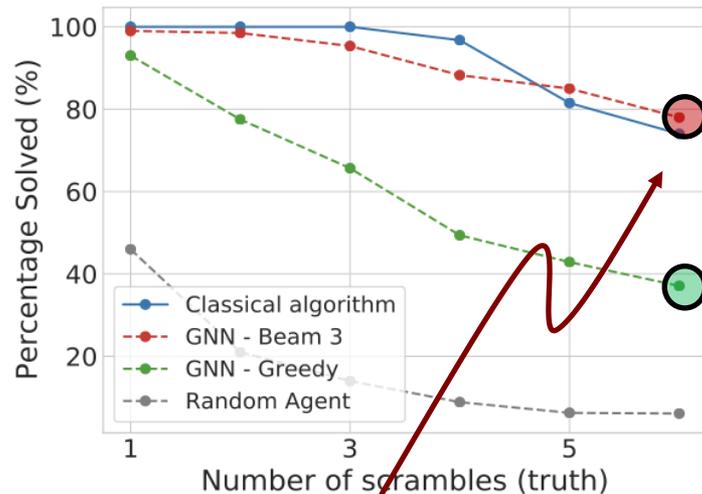
- Success reinforces good choices
- Learns best move given state

Results

Greedy: best guess only

Beam 3: keep 3 best paths in memory

Reward	Agents	Greedy		Beam size 3	
		Solved (%)	# Unscramble steps	Solved (%)	# Unscramble steps
No penalty	One-hot	50 %	5.3	78 %	14.7
	GNN	56 %	6.4	80 %	15.9
Δr_t penalty	One-hot	59 %	7.4	85 %	19.2
	GNN	53 %	8.7	89 %	20.3
	Random	13 %	8.7		
	Classical	91 %	20.2		



Does **as well** as classical algorithm with many **fewer steps**

- Has learned something!

An alternative approach to simplifying polylogarithms is to translate them

Transformer networks (Google brain 2017)

Attention is all you need

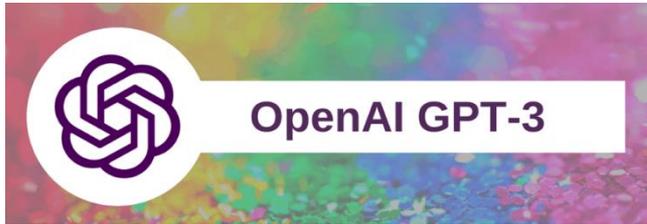
[A Vaswani, N Shazeer, N Parmar...](#) - Advances in neural ..., 2017 - proceedings.neurips.cc

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks in an encoder and decoder configuration. The best performing such ...

☆ Save  Cite **Cited by 48076** Related articles All 46 versions 

- 13 cites in hep-ph
- 15 cites in hep-th

GPT-3: Open AI's Language model



- Autogressive transformer based language model
 - Predicts next words given prompts
- 3rd generation model (2020)
- 175 billion parameters

This Nondisclosure Agreement (the "Agreement") is entered into by and between _____ with its principal offices at _____, ("Disclosing Party") and _____, located at _____ ("Receiving Party") for the purpose of preventing the unauthorized disclosure of Confidential Information as defined below. The parties agree to enter into a confidential relationship concerning the disclosure of certain proprietary and confidential information ("Confidential Information"). 1. Definition of Confidential Information. For purposes of this Agreement, "Confidential Information" shall include all information or material that has or could have commercial value or other utility in the business in which Disclosing Party is engaged. If Confidential Information is in written form, the Disclosing Party shall label or stamp the materials with the word "Confidential" or some similar warning. If Confidential Information is transmitted orally, the Disclosing Party shall promptly provide writing indicating that such oral communication constituted Confidential Information. 2. Exclusions from Confidential Information. Receiving Party's obligations under this Agreement do not extend to information that is: (a) publicly known at the time of disclosure or subsequently becomes publicly known through no fault of the Receiving Party; (b) discovered or created by the Receiving Party before disclosure by Disclosing Party; (c) learned by the Receiving Party through legitimate means other than from the Disclosing Party or Disclosing Party's representatives; or (d) is disclosed by Receiving Party with Disclosing Party's prior written approval.

GPT-3



Translate for
2nd grader

This means that if you tell anyone else something I told you in private, I can sue you and make you pay me lots of money. It also means that if I tell you something in private, I can sue you if you tell anyone else. It also means that if I tell you something in private and then I change my mind and want to tell everyone else about it, I can sue you if you tell anyone else first.



<https://beta.openai.com/playground>

What are 5 key points I should know when studying multiple polylogarithms ?

1. Multiple polylogarithms are a generalization of the logarithm function.
2. They can be used to express a wide variety of mathematical functions.
3. Multiple polylogarithms can be defined in terms of power series.
4. They can be evaluated numerically using a variety of methods.
5. Multiple polylogarithms have a wide range of applications in mathematics and physics.

How can you simplify dilogarithms?

I

There is no general way to simplify dilogarithms. However, there are some special cases where dilogarithms can be simplified. For example, if the argument of the dilogarithm is 1, then the dilogarithm can be simplified to 0.

In[1]:= PolyLog[2, 1]

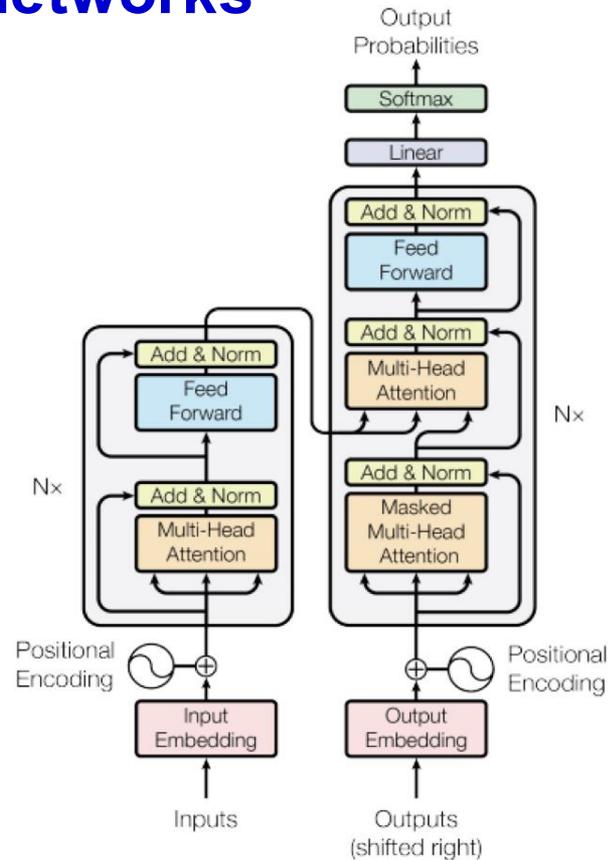
Out[1]= $\frac{\pi^2}{6}$

Hmm...

Transformer networks

Key Facts

- Sequence to sequence model
- Encoder and Decoder structure
- Used in translation tasks
- Uses **attention blocks**



Attention allows it to appreciate **context**

*The animal didn't cross the street because **it** was too tired.
L'animal n'a pas traversé la rue parce qu'**il** était trop fatigué.*

*The animal didn't cross the street because **it** was too wide.
L'animal n'a pas traversé la rue parce qu'**elle** était trop large.*

How can we use transformers?

Treat polylogarithmic expressions as sentences

$$\text{Li}_2\left(\frac{1}{x}\right) + \frac{1}{2} \ln^2(-x) \longrightarrow \text{Li}_2(x)$$

English French

Data generation

- Can no longer only consider functions that simplify to zero

e.g.

$$f(x) = \sum_{i=1}^{n_s} a_i \text{Li}_2(g_i(x)) + \sum_{i=1}^{n_t} [c_j \text{Li}_2(h_j(x)) - c_j \text{Li}_2(h_j(x))]$$

Scramble with dilog identities

$$6 \text{Li}_2(1-x) - 4 \text{Li}_2\left(-\frac{1}{x^2-x+1}\right) \xrightarrow{\text{Scramble}} -4 \text{Li}_2(x^2-x+1) + 3 \text{Li}_2(x^2-2x+1) - 3 \text{Li}_2\left(\frac{1}{x^2-4x+4}\right) + \frac{5}{2} \text{Li}_2\left(\frac{x^2+2x+1}{4x^2+4x+1}\right) - 4 \text{Li}_2(-x-1) - 5 \text{Li}_2\left(\frac{x+1}{2x+1}\right) + 2 \text{Li}_2(x^4-2x^3+3x^2-2x+1) + 6 \text{Li}_2\left(\frac{1}{x-2}\right) - 4 \text{Li}_2(x+2) - 3 \text{Li}_2\left(\frac{1}{2x+1}\right) - 5 \text{Li}_2\left(\frac{-x-1}{2x+1}\right) - 3 \text{Li}_2\left(\frac{2x}{2x+1}\right)$$

Transformer vs RL

- For RL, 0 is unique correct answer
 - RL tries to find path to 0
 - Knows when it succeeds
- For transformer, many equivalent simple expressions
 - Path not determined, just result
 - Not guaranteed to be correct

- Can be many equally simple translations

$\text{Input : } -\text{Li}_2(-2x^2) + \text{Li}_2\left(-\frac{1}{2x^2}\right)$ $-7\text{Li}_2(-2x) - 7\text{Li}_2\left(-\frac{1}{2x}\right)$ $4\text{Li}_2(x^2 - 2x + 2)$
$\text{Target : } 4\text{Li}_2(-x^2 + 2x - 1)$

transformer

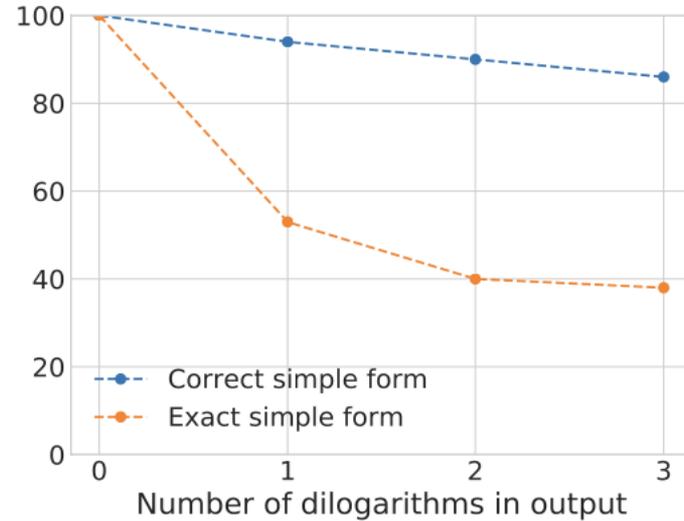
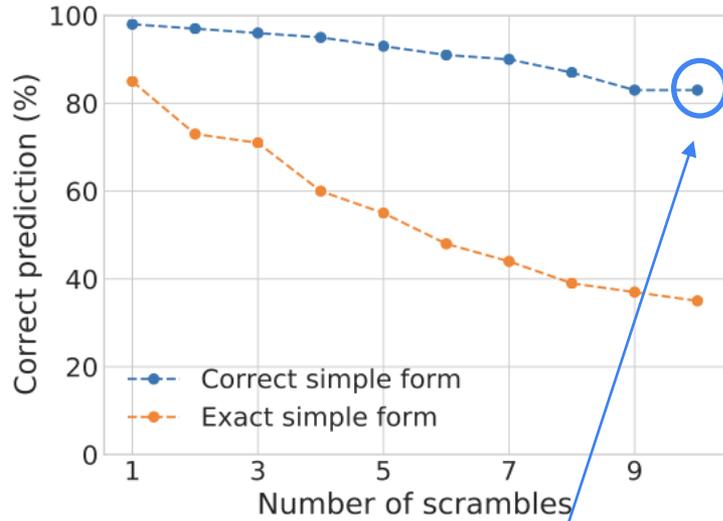


Hypothesis	Valid ?
$-4\text{Li}_2(x^2 - 2x + 2)$	\checkmark^*
$-4\text{Li}_2\left(-\frac{1}{x^2 - 2x + 1}\right)$	\checkmark^*
$4\text{Li}_2(-x^2 + 2x - 1)$	\checkmark
$4\text{Li}_2\left(\frac{1}{x^2 - 2x + 2}\right)$	\checkmark^*
$-4\text{Li}_2\left(-\frac{1}{x^2 + 2x + 1}\right)$	\times

}

Beam 5

Results



Predicts a **correct** simplified answer 91% of the time!

Symbols

An alternative approach to simplifying polylogarithms is with the **symbol**

Goncharov et al. [arXiv:1006.5703]

Polylogarithms are **iterated integrals**

$$\text{Li}_1(s) \equiv \int_0^s \frac{dx}{1-x} = -\ln(1-s) \quad \text{Li}_2(s) \equiv \int_0^s \frac{dx}{x} \text{Li}_1(s) \quad \text{etc.}$$

Symbol is a map that extracts the dlog forms

$$\mathcal{S} \left[\int_a^b d \ln R_1 \circ \cdots \circ d \ln R_n \right] = R_1 \otimes \cdots \otimes R_n$$

$$\mathcal{S} [\text{Li}_n(x)] = -(1-x) \otimes \underbrace{x \otimes \cdots \otimes x}_{n-1}$$

$$\mathcal{S} \left[\frac{1}{n!} \ln^n x \right] = \underbrace{x \otimes \cdots \otimes x}_n \quad \mathcal{S} [\ln x \ln y] = x \otimes y + y \otimes x$$

Symbols

Symbol satisfies the **product rule** (and other identities)

$$\dots \otimes f(x)g(y) \otimes \dots = \dots \otimes f(x) \otimes \dots + \dots \otimes g(y) \otimes \dots$$

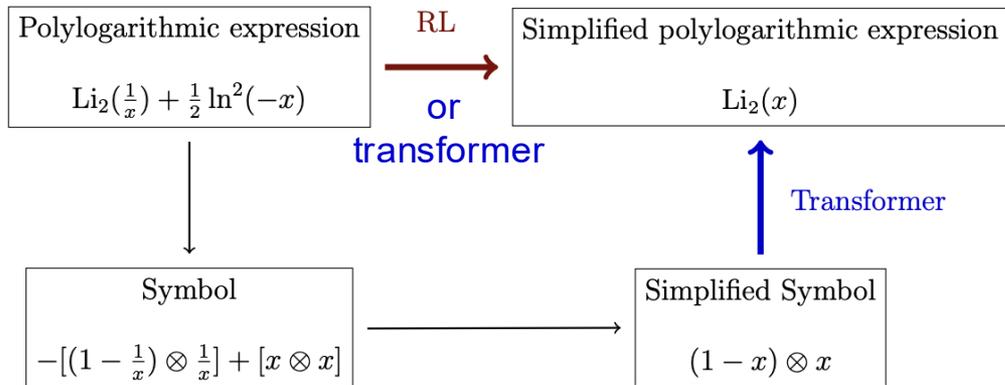
- Reduces simplifying $\text{Li}_n(x y)$ to simplifying $\log(xy)$

e.g.

$$[(1 - \frac{1}{x}) \otimes x] = [-\frac{1-x}{x} \otimes x] = [(1-x) \otimes x] - [x \otimes x]$$

$$\mathcal{S} \left[\text{Li}_2(x) + \text{Li}_2\left(\frac{1}{x}\right) + \frac{1}{2} \ln^2(-x) \right] = -[(1-x) \otimes x] - [(1 - \frac{1}{x}) \otimes \frac{1}{x}] + [x \otimes x] = 0$$

No free lunch: integrating the symbol back to polylogarithms not trivial



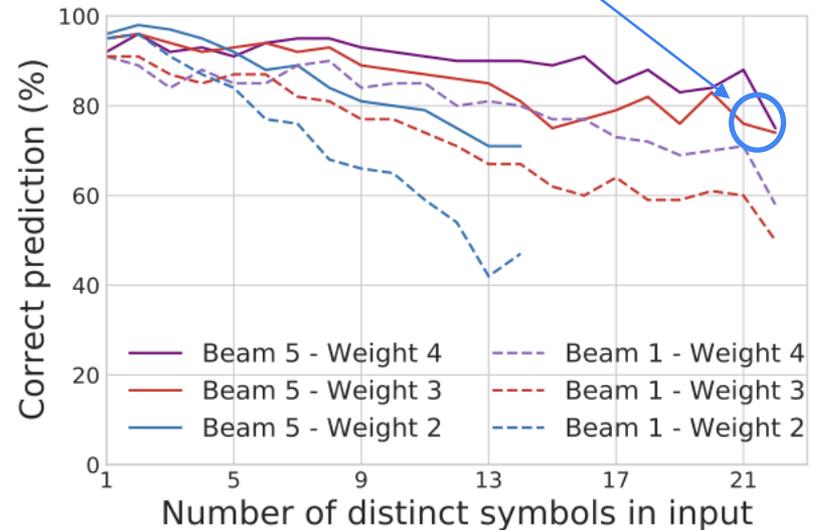
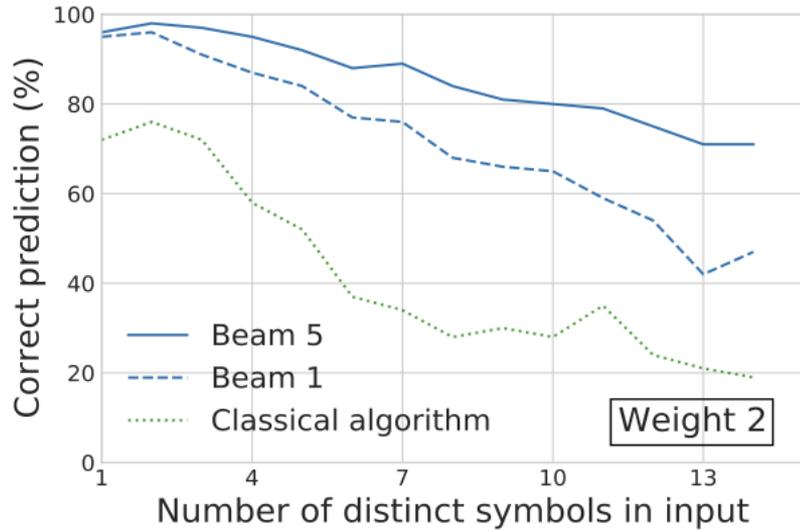
Example symbol integration

Training data looks like this

Input symbol \mathcal{S}_i	Simple expression F_i
$-(-x^2 - x + 1) \otimes (1 - x) + (-x^2 - x + 1) \otimes x$ $-(-x^2 - x + 1) \otimes (x + 1) + x \otimes (1 - x) - x \otimes x + x \otimes (x + 1)$	$\text{Li}_2\left(\frac{(1-x)(x+1)}{x}\right)$
$-\frac{23}{4}(1 - x) \otimes x - 6(x^2 + x + 1) \otimes x$	$2 \text{Li}_2(x^3) - \frac{1}{4} \text{Li}_2(x)$
$-40(6 - x^2) \otimes (6 - x^2) - 3(1 - x) \otimes (-x^6 - x^2 + 3)$ $-3(x + 1) \otimes (-x^6 - x^2 + 3) - 3(x^4 + x^2 + 2) \otimes (-x^6 - x^2 + 3)$ $+\frac{1}{4}(5 - 2x) \otimes (2 - x)$	$3 \text{Li}_2(-x^6 - x^2 + 3)$ $-\frac{1}{4} \text{Li}_2(2x - 4)$ $-20 \ln^2(x^2 - 6)$
$8 \frac{x^2 - x - 1}{x - 1} \otimes x - 8((x + 1)(x^2 - x - 1)) \otimes x$ $+8(1 - x) \otimes (-x^3 + x^2 - x - 1) - 8 \frac{1}{x - 1} \otimes x$ $-8(1 - x) \otimes (x(x^3 - x^2 + x + 1))$	$4 \text{Li}_2(x^2)$

Results

80% correct with 21 terms of weight 4



		Beam Size 1	Beam Size 5
Weight 2	Transformer	82%	91%
	Classical Algorithm	59%	59%
Weight 3	Transformer	78%	88%
Weight 4	Transformer	80%	89%

Weight 2	Weight 3	Weight 4
$\text{Li}_2(x)$	$\text{Li}_3(x)$	$\text{Li}_4(x)$
$\ln(x) \ln(y)$	$\text{Li}_2(x) \ln(y)$	$\text{Li}_3(x) \ln(y)$
	$\ln(x) \ln(y) \ln(z)$	$\text{Li}_2(x) \text{Li}_2(y)$
		$\text{Li}_2(x) \ln(y) \ln(z)$
		$\ln(w) \ln(x) \ln(y) \ln(z)$

Can simplify complicated expressions with Li_2 , Li_3 , Li_4, \dots

- Limited by compute (network size, training time)

Final example

1. Scattering amplitude gives some function of GPLs with complex arguments

$$f(x) = 4\zeta_3 + 9 \left[G(0, 0, 1, x) + G\left(0, 0, \frac{-1 - \sqrt{3}i}{2}, x\right) + G\left(0, 0, \frac{-1 + \sqrt{3}i}{2}, x\right) \right] \\ + 4 \left[-G(-1, -1, -1, x) + G(-1, 0, -1, x) + G(0, -1, -1, x) + G(0, 0, 1, x) - G\left(0, 0, 1, \frac{x}{x+1}\right) \right]$$

2. Express in terms of classical polylogs

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

3. Compute the symbol and simplify

$$\mathcal{S}[f(x)] = 9(x^2 + x + 1) \otimes x \otimes x + 13(1 - x) \otimes x \otimes x + 4(x + 1) \otimes x \otimes x$$

4. Integrate the symbol with a transformer network

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

Symbolic ML methods will be **essential** for the future of High Energy Physics

Why am I so sure of this?

The future is here!

(images from dall-E)

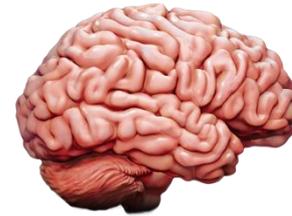
Richard Feynman doing a jig in Prague

Robots taking over the Charles bridge



PaLM and Minerva

- PaLM (Pathways Language Model) is a language model by google [April 2022]



80 billion neurons

540 billion parameters

trained for 60 days on 6144 TPU (Tensor Processing Unit) chips
\$23 million to train

Solving Quantitative Reasoning Problems with Language Models

Aitor Lewkowycz*, Anders Andreassen†, David Dohan†, Ethan Dyer†, Henryk Michalewski†,
Vinay Ramasesh†, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo,
Yuhuai Wu, Behnam Neyshabur*, Guy Gur-Ari*, and Vedant Misra*

Google Research

arXiv:2206.14858
[June 2022]

- **Minerva**

- Built on PaLM
- Trained on 60 GB of arXiv papers and mathy webpages
 - (e.g. libretxts, chegg, stacoverflow)
- 1 month additional training on 1024 TPU chips

Minerva (June 2022)

<https://minerva-demo.github.io>

QUESTION Four politicians and three lawyers attend a party. Each politician shakes hands exactly once with everyone, and each lawyer shakes hands exactly once with each politician. How many handshakes take place?

Answer by Minerva (output is in LaTeX)

Let P be the set of politicians and L be the set of lawyers. There are $4 \cdot 3 = 12$ handshakes between the politicians and lawyers. There are also $\binom{4}{2} = 6$ handshakes among the politicians. Therefore, the total number of handshakes is $12 + 6 = \boxed{18}$.
Final Answer: The final answer is 18.

EVALUATION RESULT Correct

Minerva (June 2022)

Question: A liquid flows at a constant flow rate through a pipe with circular cross-sections of varying diameters. At one point in the pipe, the diameter is 2 cm and the flow speed is 18 m/s . What is the flow speed at another point in this pipe, where the diameter is 3 cm ? Which of the following is the right choice? Explain your answer. (A) 4 m/s (B) 6 m/s (C) 8 m/s (D) 12 m/s

Answer by Minerva (output is in LaTeX)

Let A_1 be the cross-sectional area of the pipe at the first point and A_2 be the cross-sectional area of the pipe at the second point. Let v_1 be the flow speed at the first point and v_2 be the flow speed at the second point. Since the flow rate is constant, the product of the cross-sectional area and the flow speed is constant. Hence,

$$A_1 v_1 = A_2 v_2.$$

The cross-sectional area of the pipe at the first point is $A_1 = \pi r_1^2 = \pi \left(\frac{2}{2}\right)^2 = \pi$, where r_1 is the radius of the pipe at the first point. Similarly, $A_2 = \pi r_2^2 = \pi \left(\frac{3}{2}\right)^2 = \frac{9\pi}{4}$. Therefore,

$$\pi \cdot 18 = \frac{9\pi}{4} v_2 \Rightarrow v_2 = \boxed{8\text{ m/s}}.$$

Final Answer: The final answer is (C).

EVALUATION RESULT Correct

Minerva (June 2022)

QUESTION Let \mathbf{D} be a matrix representing a dilation with scale factor $k > 0$, and let \mathbf{R} be a matrix representing a rotation about the origin by an angle of θ counter-clockwise. If

$$\mathbf{DR} = \begin{pmatrix} -7 & 1 \\ -1 & -7 \end{pmatrix},$$

then find k .

MODEL ANSWER We have

$$\mathbf{DR} = \begin{pmatrix} -7 & 1 \\ -1 & -7 \end{pmatrix}.$$

We also know that

$$\mathbf{D} = \begin{pmatrix} k & 0 \\ 0 & k \end{pmatrix}.$$

This is current state of the art...
... bar is now really high!

We have that $\mathbf{DR} = (\mathbf{DR})^T$, so we have

$$\begin{pmatrix} k & 0 \\ 0 & k \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} k & 0 \\ 0 & k \end{pmatrix} = \begin{pmatrix} k \cos \theta & k \sin \theta \\ -k \sin \theta & k \cos \theta \end{pmatrix}.$$

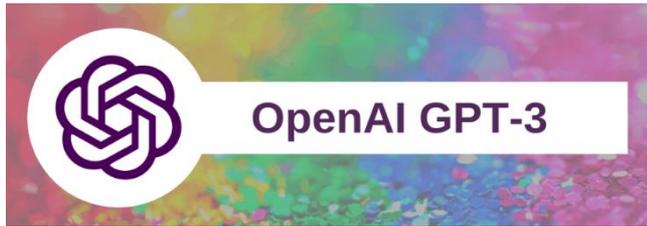
We have that $k \cos \theta = -7$ and $k \sin \theta = 1$. Thus,

$$k^2 \cos^2 \theta + k^2 \sin^2 \theta = 49 + 1 = 50.$$

This gives us that $k^2 = 50$, so $k = \sqrt{50} = \boxed{5\sqrt{2}}$. Final Answer: The final answer is $5\sqrt{2}$.

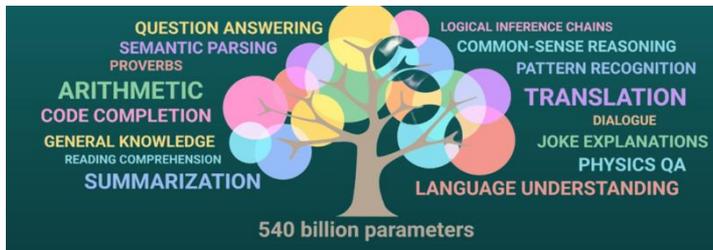
EVALUATION RESULT Correct

What's next?

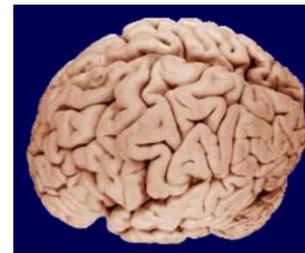


- 3rd generation model (2020)
- 175 billion parameters

Google : PaLM/Minerva (2022)
540 billion parameters



Human brain

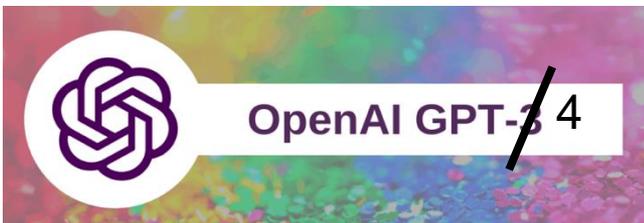


80 billion neurons
150 trillion synapses

Cat brain



0.760 billion neurons
10 trillion synapses



...expected soon

100 trillion parameters
What will it do?

2023 ... 2040... 2100? ...3000...?

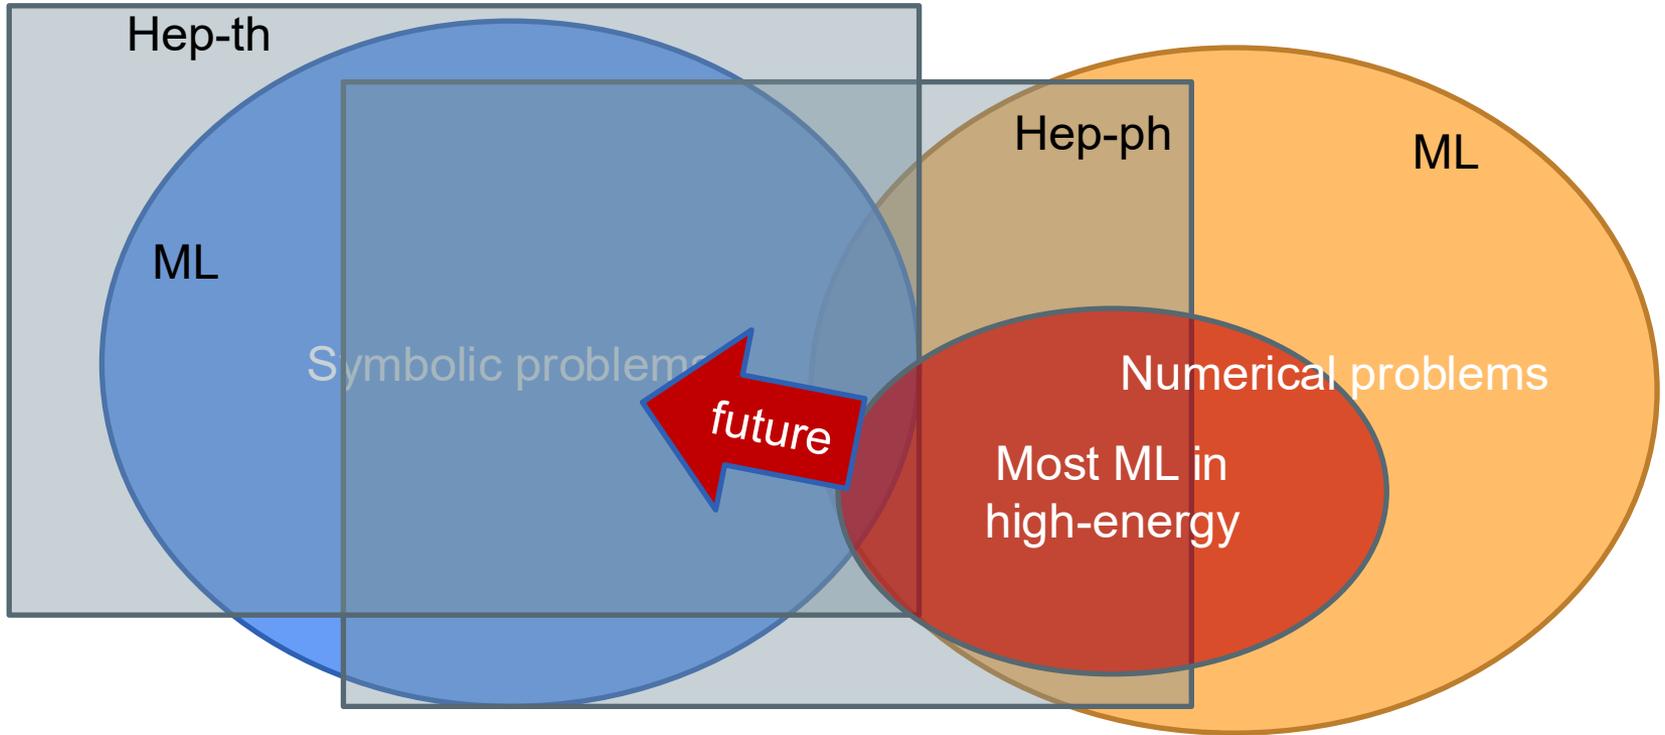
Summary

We considered the problem of **simplifying polylogs with ML**

$$\begin{array}{l} -\text{Li}_2(-2x^2) + \text{Li}_2\left(-\frac{1}{2x^2}\right) \\ -7\text{Li}_2(-2x) - 7\text{Li}_2\left(-\frac{1}{2x}\right) \\ 4\text{Li}_2(x^2 - 2x + 2) \end{array} \longrightarrow 4\text{Li}_2(-x^2 + 2x - 1)$$

- **Reinforcement learning**
 - **Applies identities** like moves in a game
 - Learns which identity to apply
- **Transformer network**
 - **Guesses the answer** like in language translation
 - Can be used to integrate the symbol
 - Powerful method works for high-weight polylogs
- Methods around 90% accurate
- Perform non-trivial task better than any public algorithm

Outlook



- Machine learning has revolutionized data-driven particle physics
- Most work in high energy theory is symbolic
- Machine learning can help with that too!

Backup Slides

Reinforcement learning

Policy and value:

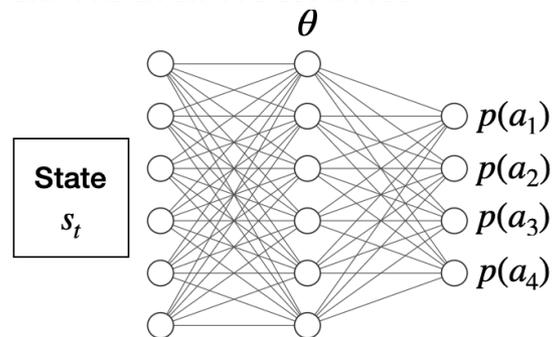
$$\pi_{\varphi, \theta}(a|s)$$

What action a should we take given state s ?

$$V_{\varphi, \theta}(s)$$

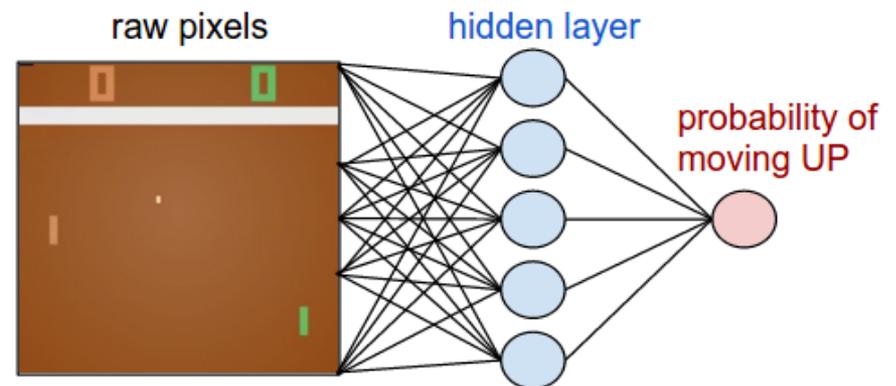
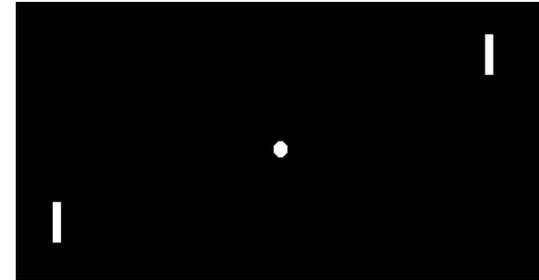
What is the estimated value of state s ?

- learned by RL as neural networks



Policy Gradients

- Q-learning -> learn a value function
 - $\hat{Q}(s, a)$ = an estimate of the expected discounted reward of taking a from s
 - Performance time: take the action that has the highest estimated value
- Policy Gradient -> learn policy directly
 - $\pi(s)$ = Probability distribution over A_s
 - Performance time: choose action according to distribution

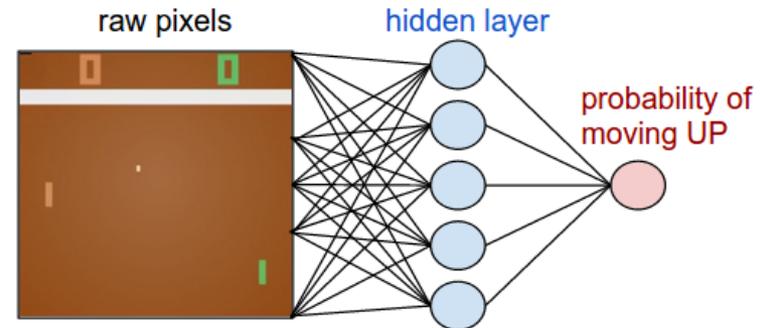


Example from: <https://www.youtube.com/watch?v=tqrcjHuNdmQ>

Policy Gradients

“Try a bunch of stuff and see what happens. Do more of the stuff that worked in the future.”

- Receive a frame
- Forward propagate to get $P(actions)$
- Select a by sampling from $P(actions)$
- Find the gradient ∇_{θ} that makes a more likely – store it
- Play the rest of the game
 - If won, take a step in direction ∇_{θ}
 - If lost, take a step in direction $-\nabla_{\theta}$

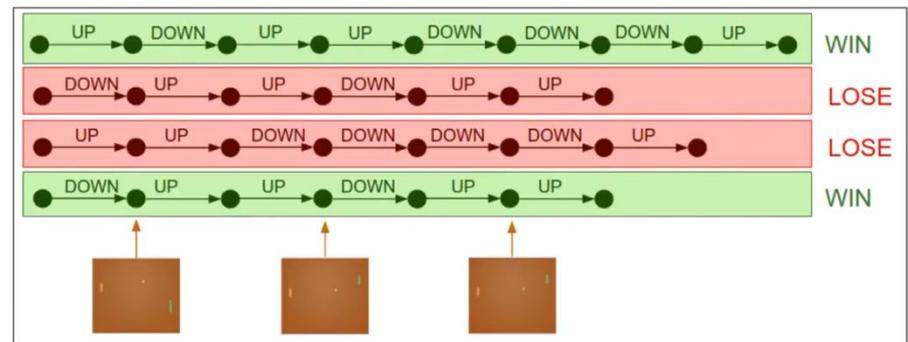


Pretend every action we took here was the correct label.

$$\text{maximize: } \log p(y_i | x_i)$$

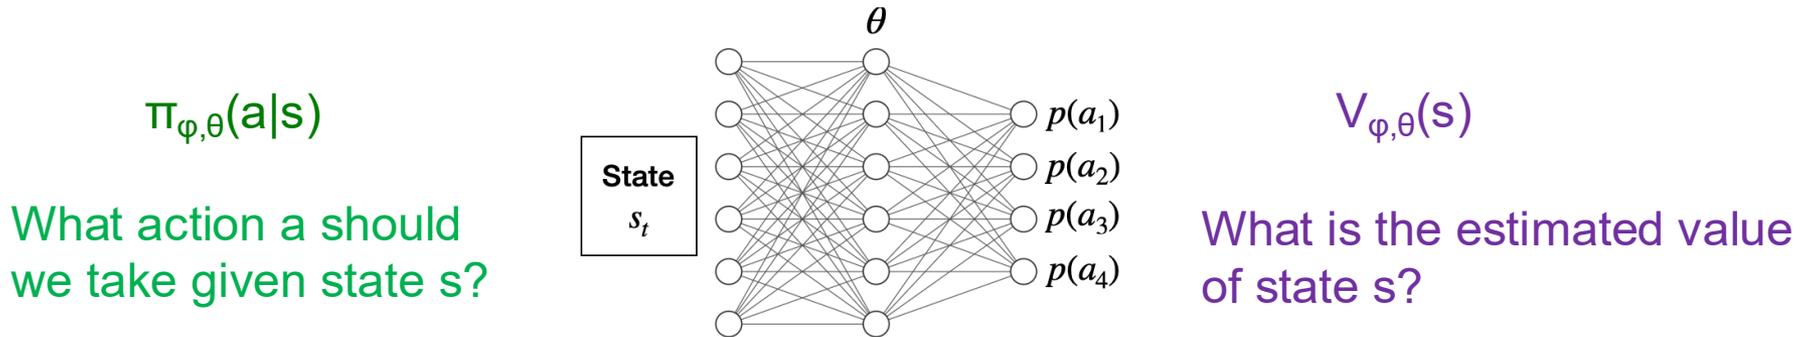
Pretend every action we took here was the wrong label.

$$\text{maximize: } (-1) * \log p(y_i | x_i)$$



Reinforcement learning

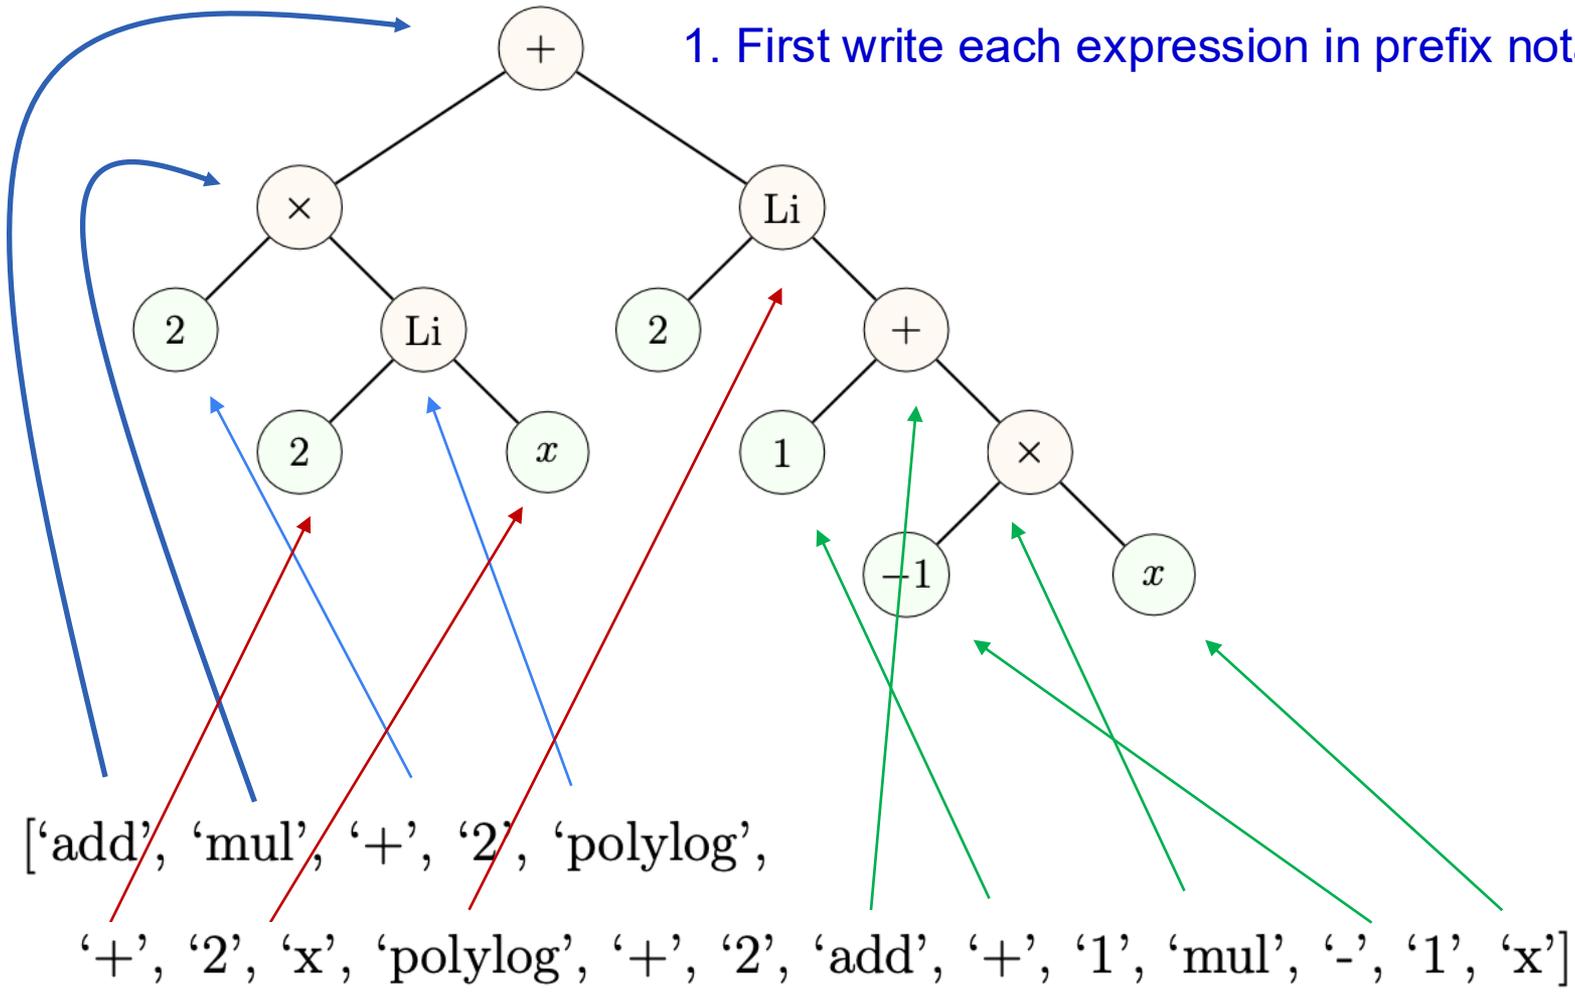
Policy and **value**: learned by RL as neural networks



- Agent:** We use Trust Region Policy Optimization (**TRPO**) [Schulman et al, 2017]
- Ensures the policy updates stay close
 - Maximizes the advantage of the new policy over the old policy
 - (also tried Proximal Policy Optimization (PPO), but not as good)

Data representation

1. First write each expression in prefix notation



2. Word embedding:

- we use one-hot encoding for each word

2	Li	2	x
Word 4	Word 2	Word 4	Word 1
0	0	0	1
0	1	0	0
0	0	0	0
1	0	1	0

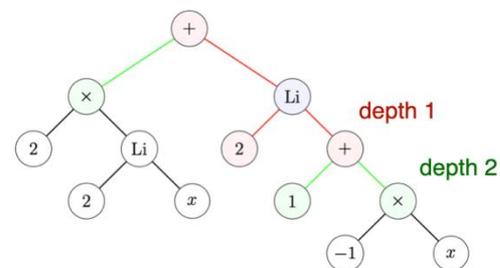
$N_{\text{words}}=20$

3. Sentence embedding:

- Explored both **one-hot encoding** and **graph neural networks**

$\text{Li}2(x^2) + 3$ $\log(x)\log(2x)$ $\text{Li}2(6x+3)$ $\text{Li}2(x)+\log(x)$
 sentence 1 sentence 2 sentence 3 sentence 4

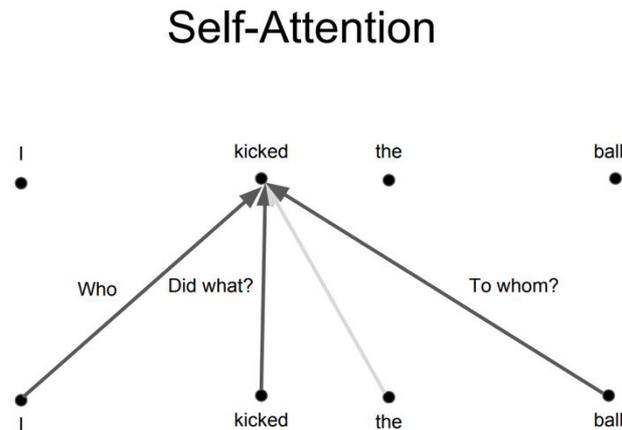
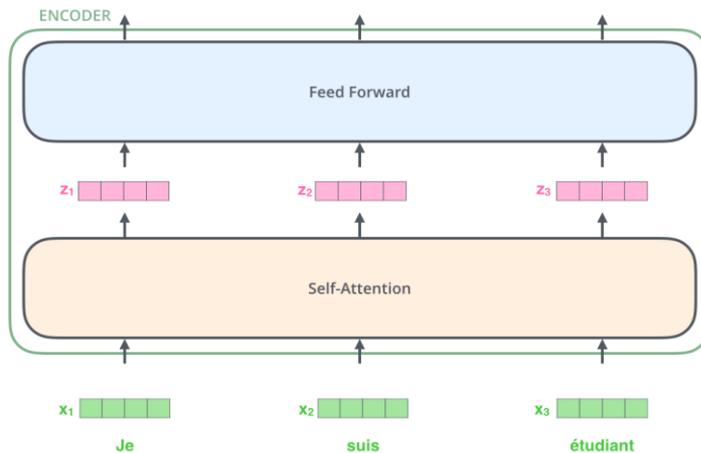
0	0	0	1
0	1	0	0
0	0	0	0
1	0	1	0



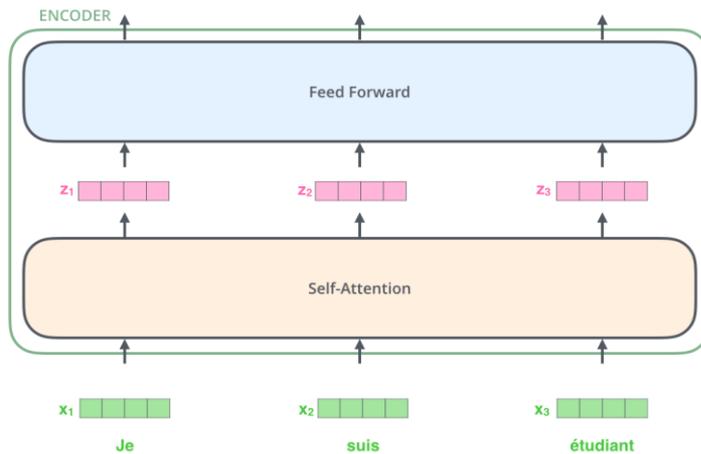
$$\begin{aligned}
 \text{Dimension} &= N_{\text{words}} \times L_{\text{max}} \\
 &= 20 \times 512 = 10,240
 \end{aligned}$$

Transformer: all about Attention

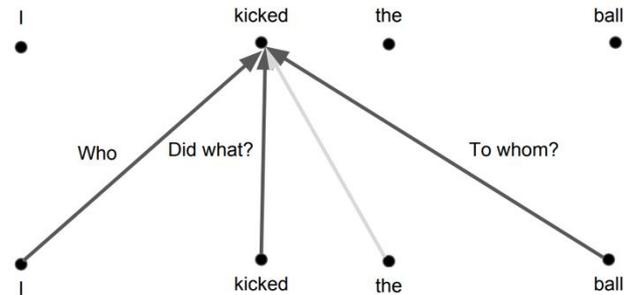
- Vanilla Attention: combination of previous hidden states
- Self-Attention: combination of previous and future hidden states
- How much attention should I give word x to word y?



Transformer: all about Attention

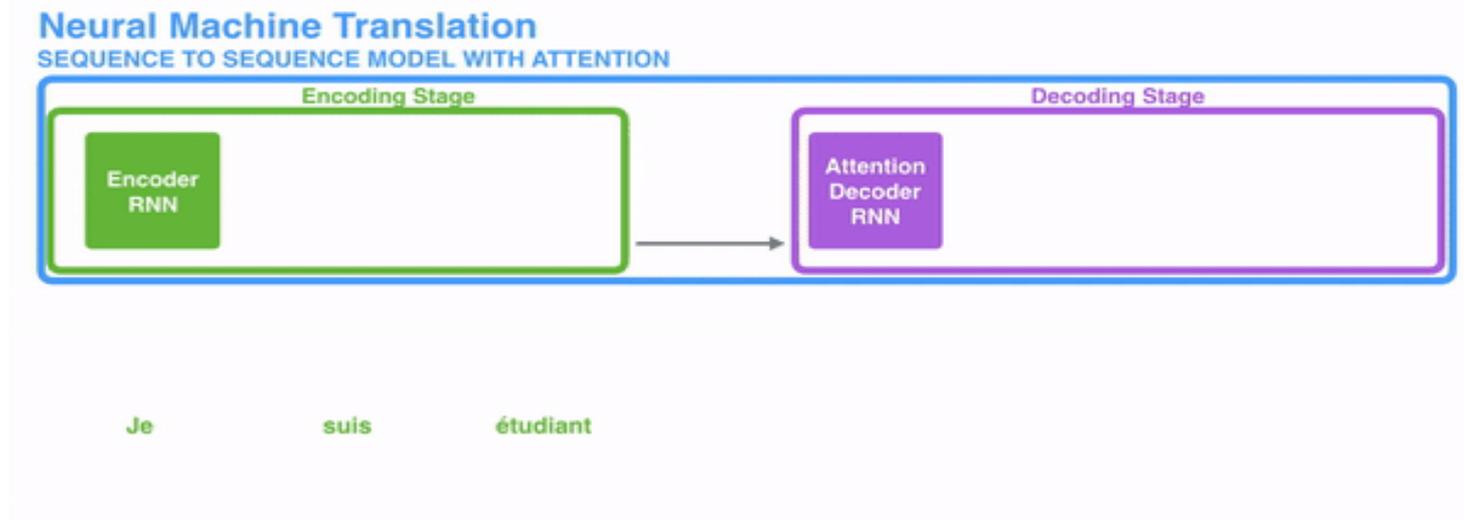


Self-Attention



Transformer

- RNN: the longer the dependencies the worse
- Especially useful in Natural Language Processing and Sequential Data
- No sequential processing



Minerva examples

Inference-time techniques: Chain-of-thought prompting

Problem:

Find the domain of the expression $\frac{\sqrt{x-2}}{\sqrt{5-x}}$.

Solution:

The expressions inside each square root must be non-negative. Therefore, $x-2 \geq 0$, so $x \geq 2$, and $5-x \geq 0$, so $x \leq 5$. Also, the denominator cannot be equal to zero, so $5-x > 0$, which gives $x < 5$. Therefore, the domain of the expression is $\boxed{[2, 5)}$.

Final Answer: The final answer is $[2, 5)$. I hope it is correct.

... 3 more examples ...

Problem:

If $\det \mathbf{A} = 2$ and $\det \mathbf{B} = 12$, then find $\det (\mathbf{A} \mathbf{B})$.

Solution:

[Ling et al. 2017](#), [Cobbe et al 2021](#), [Nye et al 2021](#), [Wei et al 2022](#)

Example: MMLU College Physics

QUESTION A uniform solid disk starts from rest and rolls down an inclined plane without slipping. After some time, what fraction of the disk's total kinetic energy is rotational kinetic energy? Which of the following is the right choice? Explain your answer. (A) 1/4 (B) 1/3 (C) 1/2 (D) 2/3

MODEL ANSWER The disk has both translational and rotational kinetic energy. Let v be the speed of the center of mass of the disk, I be the moment of inertia of the disk, M be the mass of the disk, and R be the radius of the disk. The translational kinetic energy of the disk is given by

$$K_t = \frac{1}{2}Mv^2.$$

The rotational kinetic energy of the disk is given by

$$K_r = \frac{1}{2}I\omega^2 = \frac{1}{2}\left(\frac{1}{2}MR^2\right)\left(\frac{v}{R}\right)^2 = \frac{1}{4}Mv^2.$$

Therefore, the ratio of the rotational kinetic energy to the total kinetic energy is

$$\frac{K_r}{K_t + K_r} = \frac{\frac{1}{4}Mv^2}{\frac{1}{2}Mv^2 + \frac{1}{4}Mv^2} = \frac{\frac{1}{4}Mv^2}{\frac{3}{4}Mv^2} = \frac{1}{3}.$$

Final Answer: The final answer is (B).

Source: [Minerva sample explorer](#)

Inference-time techniques: Chain-of-thought prompting

Problem:

Find the domain of the expression $\frac{\sqrt{x-2}}{\sqrt{5-x}}$.

Solution:

The expressions inside each square root must be non-negative. Therefore, $x-2 \geq 0$, so $x \geq 2$, and $5-x \geq 0$, so $x \leq 5$. Also, the denominator cannot be equal to zero, so $5-x > 0$, which gives $x < 5$. Therefore, the domain of the expression is $\boxed{[2, 5)}$.

Final Answer: The final answer is $[2, 5)$. I hope it is correct.

... 3 more examples ...

Problem:

If $\det \mathbf{A} = 2$ and $\det \mathbf{B} = 12$, then find $\det (\mathbf{A} \mathbf{B})$.

Solution:

[Ling et al. 2017](#), [Cobbe et al 2021](#), [Nye et al 2021](#), [Wei et al 2022](#)

Inference-time techniques: Majority voting / self-consistency

- Sample multiple answers at finite temperature (each has a different chain-of-thought)
- Bucket the final answers
- Choose the most common final answer

QUESTION Let n be a positive integer greater than or equal to 3. Let a, b be integers such that ab is invertible modulo n and $(ab)^{-1} \equiv 2 \pmod{n}$. Given $a + b$ is invertible, what is the remainder when $(a + b)^{-1}(a^{-1} + b^{-1})$ is divided by n ?

MODEL ANSWER We have $(a^{-1} + b^{-1}) \equiv (ab)^{-1}(a + b) \equiv 2(a + b) \pmod{n}$. Since $(a + b)^{-1}$ exists,

$$(a + b)^{-1}(a^{-1} + b^{-1}) \equiv 2 \pmod{n},$$

as desired. Final answer: The final answer is 2. <- Bucket this

Inference-time techniques: Chain-of-thought prompting

Problem:

Find the domain of the expression $\frac{\sqrt{x-2}}{\sqrt{5-x}}$.

Solution:

The expressions inside each square root must be non-negative. Therefore, $x-2 \geq 0$, so $x \geq 2$, and $5-x \geq 0$, so $x \leq 5$. Also, the denominator cannot be equal to zero, so $5-x > 0$, which gives $x < 5$. Therefore, the domain of the expression is $\boxed{[2, 5)}$.

Final Answer: The final answer is $[2, 5)$. I hope it is correct.

... 3 more examples ...

Problem:

If $\det \mathbf{A} = 2$ and $\det \mathbf{B} = 12$, then find $\det (\mathbf{A} \mathbf{B})$.

Solution:

[Ling et al. 2017](#), [Cobbe et al 2021](#), [Nye et al 2021](#), [Wei et al 2022](#)

<https://minerva-demo.github.io/#category=Physics&index=7>

QUESTION A liquid flows at a constant flow rate through a pipe with circular cross-sections of varying diameters. At one point in the pipe, the diameter is 2 cm and the flow speed is 18 m/s . What is the flow speed at another point in this pipe, where the diameter is 3 cm ? Which of the following is the right choice? Explain your answer. (A) 4 m/s (B) 6 m/s (C) 8 m/s (D) 12 m/s

MODEL ANSWER Let A_1 be the cross-sectional area of the pipe at the first point and A_2 be the cross-sectional area of the pipe at the second point. Let v_1 be the flow speed at the first point and v_2 be the flow speed at the second point. Since the flow rate is constant, the product of the cross-sectional area and the flow speed is constant. Hence,

$$A_1 v_1 = A_2 v_2.$$

The cross-sectional area of the pipe at the first point is $A_1 = \pi r_1^2 = \pi \left(\frac{2}{2}\right)^2 = \pi$, where r_1 is the radius of the pipe at the first point. Similarly, $A_2 = \pi r_2^2 = \pi \left(\frac{3}{2}\right)^2 = \frac{9\pi}{4}$. Therefore,

$$\pi \cdot 18 = \frac{9\pi}{4} v_2 \Rightarrow v_2 = \boxed{8\text{ m/s}}.$$

Final Answer: The final answer is (C).

REFERENCE ANSWER C

PROBLEM SOURCE MMLU High School Physics

EVALUATION RESULT Correct

QUESTION A stone is thrown horizontally with an initial speed of 30 m/s from a bridge. Find the stone's total speed when it enters the water 4 seconds later, assuming that air resistance is negligible. Which of the following is the right choice? Explain your answer. (A) 30 m/s (B) 40 m/s (C) 50 m/s (D) 60 m/s

MODEL ANSWER Since the stone is thrown horizontally, its horizontal speed is 30 m/s . After 4 seconds, its vertical speed is $v_y = gt = 9.8 \text{ m/s}^2 \cdot 4 \text{ s} = 39.2 \text{ m/s}$. Thus the total speed of the stone is

$$\begin{aligned} v &= \sqrt{v_x^2 + v_y^2} \\ &= \sqrt{30^2 + 39.2^2} \approx \boxed{50 \text{ m/s}}. \end{aligned}$$

Final Answer: The final answer is (C).

REFERENCE ANSWER C

PROBLEM SOURCE MMLU high school physics

EVALUATION RESULT Correct

QUESTION There are 42 people in pilates class. Today, 36 of them brought yoga mats, 26 of them brought eco-friendly water bottles, and 20 of them remembered to bring both yoga mats and water bottles. How many people had to suffer through class having neither a yoga mat nor a water bottle?

MODEL ANSWER Let A be the set of people who brought yoga mats, and let B be the set of people who brought eco-friendly water bottles. We know that $|A| = 36$, $|B| = 26$, and $|A \cap B| = 20$. We want to know how many people were in $A^c \cap B^c$. To do this, we will use the following formula:

$$\begin{aligned} |A^c \cap B^c| &= |(A \cup B)^c| \\ &= |U| - |A \cup B| \\ &= |U| - |A| - |B| + |A \cap B| \\ &= 42 - 36 - 26 + 20 \\ &= \boxed{0} \end{aligned}$$

Final Answer: The final answer is 0.