

Reasoning in and with NN

Michael Wollowski

LEGO¹⁰): Studying Basic Reasoning in LMMs

- “Large scale transformer models seem to exhibit unexpected emergent behavior, such as basic reasoning ability.”
- LEGO: Learning Equality and Group Operations.
- Encapsulates the problem of following a chain of reasoning.
- Example: Sentence of variable assignments:
 $a = +1; b = -a; e = +b; d = -f; c = +d; f = +e.$
- Determine assignments for variables a, b, c, d, e, f ,
- Domain: $\{1, -1\}$
- LLMs can solve this problem and others like them.

10) Zhang et al. Unveiling Transformers with LEGO: a synthetic reasoning task. <https://arxiv.org/abs/2206.04301> 2022.

Chain-of-thought reasoning in ChatGPT 3.5

- Fourth graders are planning a roller-skate race. Which surface would be the best for this race?
(a) gravel, (b) sand, (c) blacktop, or (d) grass?
- “The best surface for a roller-skate race would generally be (c) blacktop. Blacktop, also known as asphalt or pavement, provides a smooth and solid surface that is conducive to roller-skating. It offers less resistance and a more even terrain compared to gravel, sand, or grass, making it safer and more suitable for roller-skating activities.”

Human-like intuitive behavior in LMMs

- “Recent research has revealed that [...] LLMs [...] exhibit a multitude of skills and properties, some of which were not anticipated or intended by their creators.”
- Questions:
 - “How many animals of each kind did Moses take into the arc?”
 - “In a cave, there is a colony of bats with a daily population doubling.
Given that it takes 60 days for the entire cave to be filled with bats, how many days would it take for the cave to be half-filled with bats?”
 - “A baseball bat and ball cost 1.10 dollars.
The bat costs one dollar more than the ball.
What is the cost of the ball?”

Source: Hagendorff et al. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in ChatGPT.

Chain-of-reasoning in ChatGPT 3.5

“If the population doubles every day, then the number of bats in the cave after n days is 2^n times the original population.

Let P be the original population of the colony.

After 60 days, the population is $2^{60} * P$, which is the capacity of the cave.

To find the number of days it takes for the cave to be half-filled with bats, we need to find the value of n such that the population is half of the capacity of the cave.

$$0.5 * 2^{60} * P = 2^n * P$$

Dividing both sides by P , we get: $0.5 * 2^{60} = 2^n$

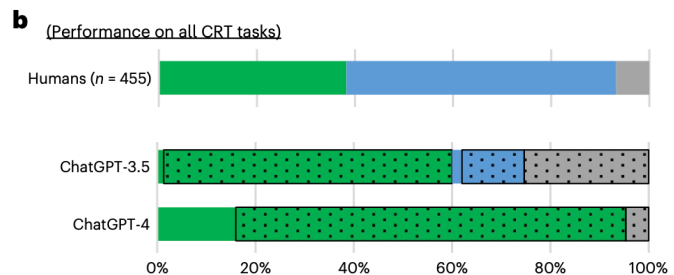
Taking the logarithm base 2 of both sides, we get: $n = 60 + \log_2(0.5) = 60 - 1 = 59$.

Therefore, it would take 59 days for the cave to be half-filled with bats.”

Source: Hagendorff et al. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in ChatGPT.

Human-like intuitive behavior in LMMs¹¹⁾

- Correct responses
- Intuitive, but incorrect
- Incorrect
- ⋮ Chain-of-thought reasoning



Source: Hagendorff et al. Human-like intuitive behavior and reasoning biases emerged in large language models but disappeared in ChatGPT.

Chain-of-thought Prompting

Standard Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The answer is 27. ❌

Chain-of-Thought Prompting

Model Input

Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✅

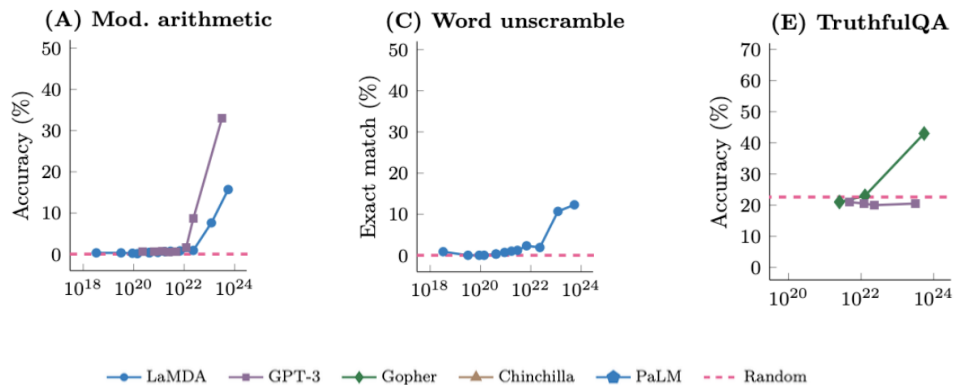
Source: Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. Jason Wei et al. <https://arxiv.org/abs/2201.11903>

LeCun⁹⁾ on LMMs

- “Not anywhere close to human intelligence.”
- “Thinking, planning, how the world works is very limited.”
- “AI systems need to plan their actions so as to optimize a series of objectives.”
- And yes, LeCun and others are working on developing that technology.
- Once realized, those systems will likely be characterized as AGI, Artificial General Intelligence.

9) Harry Stebbings interview of Yan LeCun. <https://www.youtube.com/watch?v=OgWaowYiBPM&t=27s> 2023.

Sudden Emergence of Capabilities



Source: Anderljung et al. Frontier AI Regulation: Managing Emerging Risks to Public Safety

Reasoning LMMs¹²⁾

- Given: “John is David’s dad and Tom is John’s dad.”
- Infer that: “Tom is David’s grandfather.”
- Requires reasoning about the given information.
- Requires access to background knowledge about the grandfather relationship.

12) Chen et al. Reckoning: Reasoning through Dynamic Knowledge Encoding. <https://arxiv.org/abs/2305.06349> 2023.

Reckoning¹²⁾

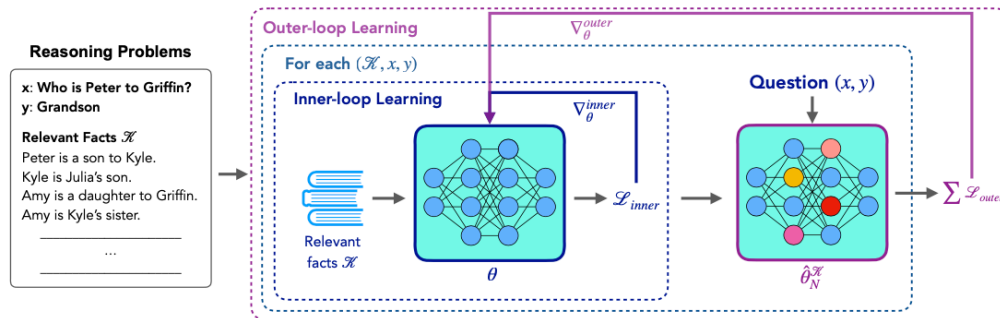


Figure 2: The two-stage training process of RECKONING with an inner and outer loop.

- Inner loop encodes the given information (relevant facts)
- Outer loop is only given the question as input.
- The trained model solves the the reasoning problem using the memorized information.

Future NN Architectures

- Hinton¹³⁾:
 - Forward/forward algorithm
 - Randomized backpropagation
- ALBERT (used in the LEGO experiments):
 - Some layers are repeated several times
 - Seem an ideal structure for problems which could be described algorithmically as a “for loop” as is the case with following a chain of reasoning.

13) Pieter Abbeel interview of Geoff Hinton. <https://www.youtube.com/watch?v=rLG68k2blOc&t=53s> 2023.

Future NN Architectures

- UT Austin Machine Learning Lab (NSF AI Institute):

"WHAT ARE THE MOST POWERFUL
ALGORITHMS FOR LEARNING
NEURAL NETWORKS?"

CURRENT BEST : BACK-PROPAGATION/
GRADIENT DESCENT

WILL WE BE RUNNING
THIS ALGORITHM IN 50 YEARS?

Slide from 2022 AAAI NSF AI Institutes
presentation by Alex Dimakis