What is LLM Reasoning?

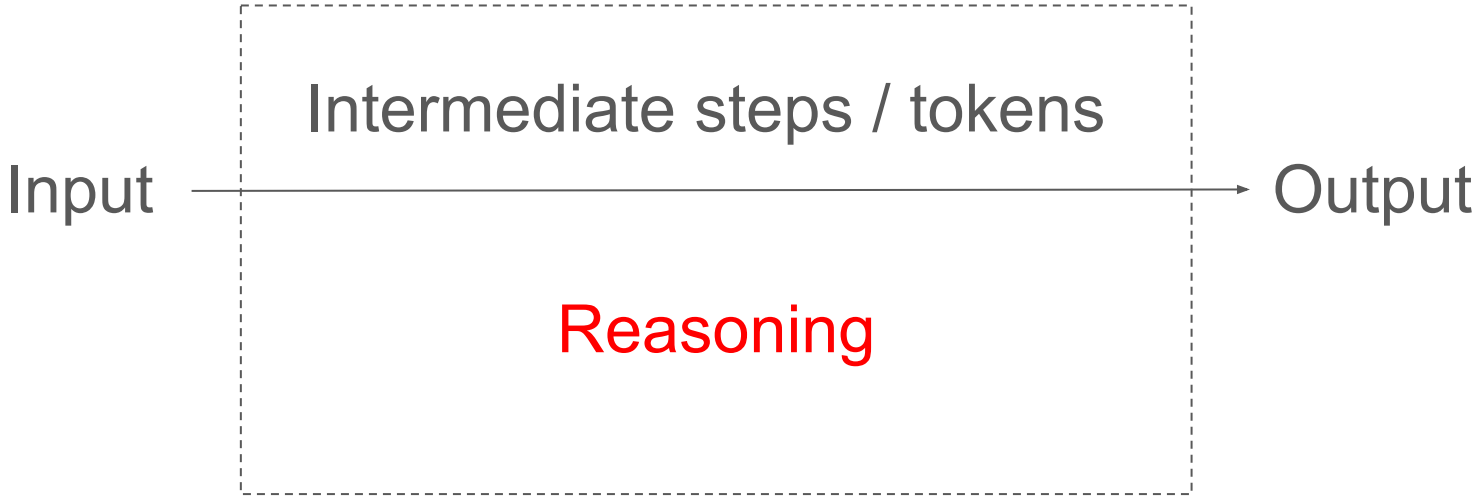

Ling et al. Program Induction by Rationale Generation: Learning to Solve and Explain Algebraic Word Problems. ACL 2017

Chen et al. Compositional generalization via neural-symbolic stack machines. NeurIPS 2020.

What is the output when concatenating the last letter of each word in “artificial intelligence”?

No reasoning

The answer is “le”.

Reasoning

The last letter of “artificial” is “l”. The last letter of “intelligence” is “e”. Concatenating “l” and “e” leads to “le”. So the answer is “le”.

Parisotto, Emilio, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Denny Zhou, and Pushmeet Kohli. Neuro-symbolic program synthesis. arXiv preprint arXiv:1611.01855 (2016).

Common Belief

Pretrained LLMs cannot reason without further prompting engineering or finetuning

WRONG

Pretrained LLMs are ready to reason

All we need is decoding

Xuezhi Wang and Denny Zhou. Chain-of-Thought Reasoning Without Prompting. NeurIPS 2024.

Select responses with high confidence on answers!

I have 3 apples. My dad has 2 more apples than me. How many apples do we have in total?

5 apples. (Greedy Decoding)

I have 3 apples, my dad has 2 more apples than me, so he has 5 apples. 3+5=**8**.

We have 8 apples in total.

You have 3 apples, your dad has 2 more apples than you, so he has 5 apples. 3+5=**8**.

The answer is 5.

Way higher confidence on reasoning-based answers!

Chain-of-Thought Decoding

1. Go beyond greedy decoding by checking more generation candidates
2. Choose candidates which have the highest confidence on the final answer

Xuezhi Wang and Denny Zhou. Chain-of-Thought Reasoning Without Prompting. NeurIPS 2024.