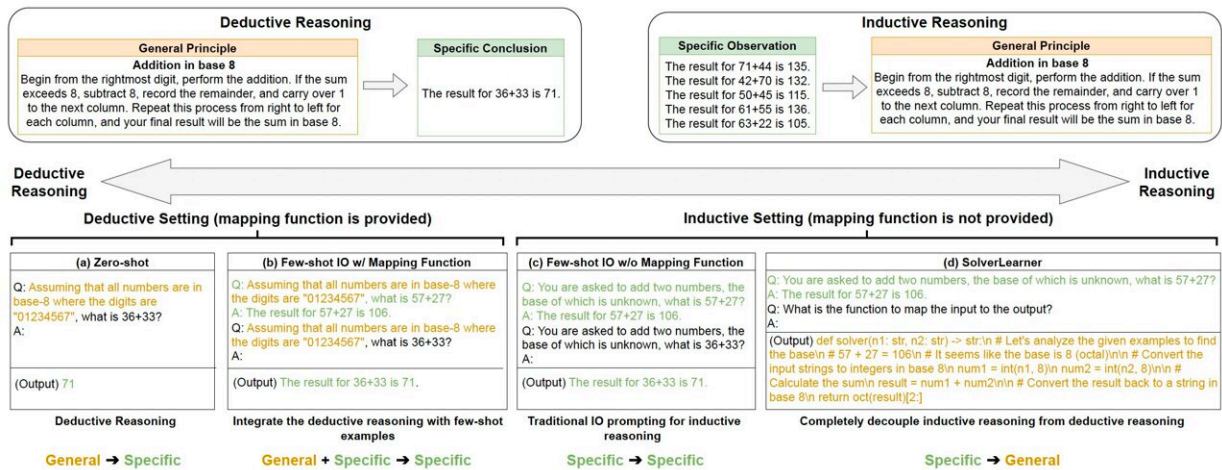


Exploring the fundamental reasoning abilities of LLMs

August 31 2024, by Ingrid Fadelli



Comparative experiments that utilize a consistent task across different contexts, each emphasizing either deductive (i.e., methods (a) and (b)) or inductive reasoning (i.e., methods (c) and (d)). Credit: Cheng et al.

Reasoning, the process through which human beings mentally process information to draw specific conclusions or solve problems, can be divided into two main categories. The first type of reasoning, known as deductive reasoning, entails starting from a general rule or premise and then using this rule to draw conclusions about specific cases.

This could mean, for instance, building on the premise that "all dogs have ears" and "Chihuahuas are dogs," to conclude that "chihuahuas have

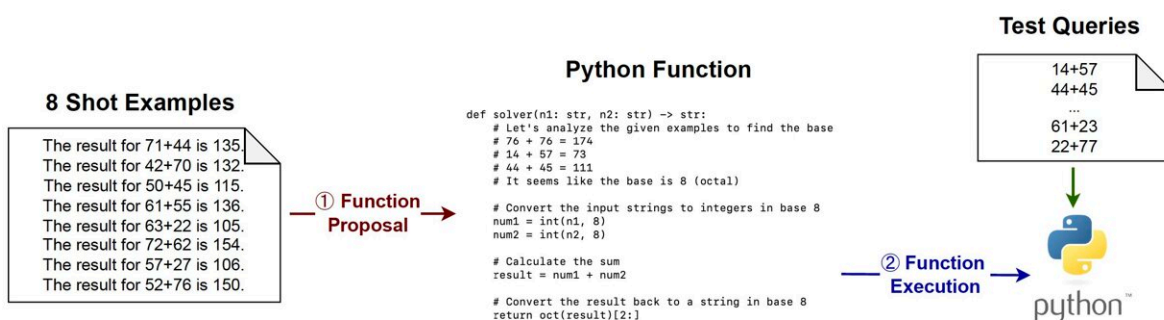
ears."

The second widely used form of reasoning is inductive reasoning, which instead consists of generalizing (i.e., formulating general rules) based on specific observations. This could mean, for instance, assuming that all swans are white because all the swans we encountered during our lifetime were white.

Numerous past research studies have investigated how humans use deductive and inductive reasoning in their everyday lives. Yet the extent to which [artificial intelligence](#) (AI) systems employ these different reasoning strategies has, so far, rarely been explored.

A research team at Amazon and University of California Los Angeles recently carried out a study exploring the fundamental reasoning abilities of large language models (LLMs), large AI systems that can process, generate and adapt texts in human languages. Their [findings](#), posted to the *arXiv* preprint server, suggest these models have strong inductive reasoning capabilities, while they often exhibit poor deductive reasoning.

The objective of the paper was to better understand gaps in LLM reasoning and identify why LLM's exhibit lower performance for "counterfactual" reasoning tasks that deviate from the norm.



Overview of the team's framework SolverLearner for inductive reasoning. SolverLearner follows a two-step process to separate the learning of input-output mapping functions from the application of these functions for inference. Specifically, functions are applied through external code interpreters, to avoid incorporating LLM-based deductive reasoning. Credit: Cheng et al.

Various past studies assessed the deductive reasoning skills of LLMs by testing their ability to follow instructions as part of basic reasoning tasks. Yet their inductive reasoning (i.e., their ability to make general predictions based on the information they processed in the past) had not been closely examined.

To clearly distinguish inductive reasoning from deductive reasoning, the researchers introduced a new model, called SolverLearner. The model uses a two-step approach to separate the process of learning rules from that of applying them to specific cases. In particular, the rules are applied through external tools, like code interpreters, to avoid relying on the LLM's deductive reasoning capability, according to an Amazon spokesperson.

Using the SolverLearner framework they developed, the team at Amazon trained LLMs to learn functions that map out input data points to their corresponding outputs, using specific examples. This in turned allowed them to investigate the extent to which the models could learn general rules based on the examples provided to them.

The researchers found that LLMs have stronger inductive reasoning capability than deductive, especially for tasks involving "counterfactual" scenarios that deviate from the norm. These findings can help people better understand when and how to use LLMs. For instance, when

designing agent systems, like chatbots, it may be better to leverage the strong inductive capabilities of LLMs.

Overall, the researchers found that LLMs performed remarkably well on inductive reasoning tasks, yet they often lacked deductive reasoning abilities. Their deductive reasoning appeared to be particularly poor in scenarios that were based on hypothetical assumptions or deviated from the norm.

The results gathered as part of this study could inspire AI developers to leverage the strong inductive reasoning capabilities of LLMs to tackle specific tasks. In addition, they could pave the way for further efforts aimed at understanding LLM reasoning processes.

According to an Amazon spokesperson, future research in this area could focus on exploring how the ability of an LLM to compress information relates to its strong inductive capabilities. This perspective may further improve the LLM's inductive reasoning capabilities.

More information: Kewei Cheng et al, Inductive or Deductive? Rethinking the Fundamental Reasoning Abilities of LLMs, *arXiv* (2024). [DOI: 10.48550/arxiv.2408.00114](https://doi.org/10.48550/arxiv.2408.00114)

© 2024 Science X Network

Citation: Exploring the fundamental reasoning abilities of LLMs (2024, August 31) retrieved 2 September 2024 from <https://techxplore.com/news/2024-08-exploring-fundamental-abilities-llms.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.