arXiv:2311.03033v1 [cs.LG] 6 Nov 2023

Beyond Words: A Mathematical Framework for Interpreting Large Language Models

Javier González and Aditya V. Nori

Microsoft Research

Abstract

Large language models (LLMs) are powerful AI tools that can generate and comprehend natural language text and other complex information. However, the field lacks a *mathematical framework* to systematically describe, compare and improve LLMs. We propose HEX a framework that clarifies key terms and concepts in LLM research, such as hallucinations, alignment, self-verification and chain-of-thought reasoning. The HEX framework offers a precise and consistent way to characterize LLMs, identify their strengths and weaknesses, and integrate new findings. Using HEX, we differentiate chain-of-thought reasoning from chain-of-thought prompting and establish the conditions under which they are equivalent. This distinction clarifies the basic assumptions behind chain-of-thought prompting and its implications for methods that use it, such as self-verification and prompt programming.

Our goal is to provide a formal framework for LLMs that can help both researchers and practitioners explore new possibilities for generative AI. We do not claim to have a definitive solution, but rather a tool for opening up new research avenues. We argue that our formal definitions and results are crucial for advancing the discussion on how to build generative AI systems that are safe, reliable, fair and robust, especially in domains like healthcare and software engineering.

1 Introduction

Large language models (LLMs) are powerful tools that can generate text responses to various input prompts, such as queries or commands. To interact with an LLM, a human user provides a prompt and receives a response from the model, which is based on its internal algorithms. The user then evaluates the response, and decides whether it answers the prompt satisfactorily. This process may be repeated several times until the user reaches a desired outcome (or gives up).

State-of-the-art LLMs, such as PaLM [CND+22], GPT-4 [Ope23] or LLaMA [TLI+23], can exhibit both remarkable and baffling behaviours, depending on the task. Peter Lee captures this paradox by saying that they are both "*smarter and dumber than you*" at math, statistics and logic [Sch23]. What causes this discrepancy? Research has primarily focused on testing LLMs on different tasks so the user can gain insights about their behaviour. For GPT-4, a large collection of behaviours have been studied on several tasks [BCE+23]. Benchmarking efforts have been conducted in various domains such as causal discovery [KNST23], summarization [ZLD+23] and reasoning [SPH+23]. These research directions are reasonable, given that most LLMs are black-boxes that we cannot easily inspect, interpret or explain.

In this paper, we offer a new perspective on interpreting LLMs and present a mathematical framework that formalises and generalises what is already known about LLMs. To advance LLM research, we pursue two objectives: First, we define key LLM concepts such as alignment [SJH+23], hallucinations [JLF+23] and chain-of-thought reasoning [WWS+23] using our mathematical framework. Second, we provide a reasoning tool that facilitates a common understanding of existing LLM research and form a basis for exploring new questions and challenges. Our framework, called HEX, is based on commutative diagrams that relate different levels of abstraction and computation in LLMs. The framework itself does not aim to provide all the answers but it represents a framework that can be used by researchers in LLMs to share, position and reason about their findings.

We show how HEX can capture the essence of LLMs as abstract execution machines that use natural language as an interface, and how it can reveal the assumptions and implications of various methods that use LLMs, such as self-verification and prompt programming. We also show how HEX can formalize and evaluate the output of LLMs using distance metrics and compatibility relations.

Our goal is to provide a formal framework for LLMs that can help both researchers and practitioners explore new possibilities for generative AI. We do not claim to have a definitive solution, but rather a tool for opening up new research avenues. We argue that our formal definitions and results are crucial for advancing the discussion on how to build generative AI systems that are safe, reliable, fair and robust, especially in domains like healthcare and software engineering.

The rest of the paper is organized as follows. Section 2 introduces the HEX framework and its basic definitions. Section 3 shows how the HEX framework can formalize common LLM concepts, such as hallucination, alignment, and chain-of-thought reasoning. Section 4 discusses the conditions of when chain-of-thought prompting elicits chain-of-thought reasoning, and their broader implications for LLM methods and how they are used in practice. Section 5 concludes the paper and outlines future work.

2 The mathematical framework

LLMs are abstract execution machines that use natural language as an interface. Given an input prompt that captures the user's intention, an LLM produces an output answer in the form of text that is presented to the user. The prompt aims to solve a problem, such as a mathematical operation, a summarization, or the generation of a certain type of text. The prompt always expresses an *intention* to execute some operations based on some provided information. The LLM, which operates as a black-box, uses the internal memory acquired at training time to provide an answer by sampling from some probability distribution conditioned on the input prompt[1].

This process has become extremely powerful in the latest generation of LLMs, creating a new set of terms for AI research and practice. Alignment, hallucinations, chain-of-thought reasoning, self-verification and prompt programming are some of the concepts that AI researchers and practitioners use routinely – our work allows us to formally define these concepts using a unified mathematical framework (Section 3). We introduce the definitions needed for our framework next.

Concrete state: A concrete state is a function that assigns values from a concrete domain to a (finite) set of variables. We use $\mathcal{V}$ to denote the set of variables, and $\mathcal{C}$ to denote the concrete domain. We write $\sigma : \mathcal{V} \to \mathcal{C}$ to indicate that σ is a concrete state. For example, suppose $\mathcal{V} = \{x, y\}$ and $\mathcal{C} = \mathbb{Z}$, the set of integers. Then, one possible concrete state is

[1]This is computed by choosing the next word that best fits the previous text, and repeating the process for each word.

$\sigma = \{x \mapsto 1, y \mapsto 2\}$, which means that $\sigma(x) = 1$ and $\sigma(y) = 2$. The set of all concrete states over $\mathcal{V}$ and $\mathcal{C}$ is denoted by $2^{\mathcal{V} \to \mathcal{C}}$.

Problem: Let $q : 2^{\mathcal{V} \to \mathcal{C}} \to 2^{\mathcal{V} \to \mathcal{C}}$ be a query, and let σ be a concrete state. Informally, the query q expresses what a user wishes to compute over a state σ. For example, q might select some variables, apply some operations, or filter some conditions. A problem is defined by the query-state pair (q, σ).

Prompt: A prompt is a natural language expression of a problem (q, σ) that an LLM can solve[2]. A prompt $(\overline{q}, \overline{\sigma})$ consists of two strings: a query string $\overline{q}$ that represents the query q in words, and a state string $\overline{\sigma}$ that describes the concrete state σ in words. We denote by $\mathcal{T}$ the set of all strings.

LLM: Given a prompt $(\overline{q}, \overline{\sigma})$ as input, an LLM M interprets this to produce a natural language or string output.

Abstract state: A state string $\overline{\sigma}$ has a meaning that can be captured by a vector, which is an abstract state $\hat{\sigma}$ or an embedding. An LLM works with these vectors to manipulate state strings. We denote by $\mathcal{A}$ the set of all abstract states.

Abstraction map: An abstraction map α is a function that transforms a concrete state σ into an abstract state $\hat{\sigma}$. The function α consists of two sub-functions: (i) α_c, which converts the concrete state σ to a state string $\bar{\sigma}$ (this would be part of the prompt design process) and (ii) α_a, which encodes the text state $\bar{\sigma}$ into an abstract state $\hat{\sigma}$ (this is the LLM encoder). In other words, $\alpha(\sigma) = (\alpha_a \circ \alpha_c)(\sigma)$.

Concretization map: a concretization map γ concretizes an abstract state $\hat{\sigma}$ by producing a concrete state σ. The function γ consists of two steps: (i) γ_a, which decodes the abstract state $\hat{\sigma}$ into a state string $\bar{\sigma}$ using an LLM (the LLM decoder), and (ii) γ_c, which interprets the state string $\bar{\sigma}$ as a description of the concrete output state (the human interpretation of the LLM output). In other words, $\gamma(\hat{\sigma}) = (\gamma_c \circ \gamma_a)(\hat{\sigma})$.

Figure 1 shows the HEX diagram[3] for an LLM M and a problem (q, σ), where q is a query and σ is a concrete state. The corners of the diagram are states and the vertical and horizontal edges connect different levels of abstraction. The vertical edges are based on the abstraction maps α_c and α_a, and the concretization maps γ_a and γ_c. The horizontal edges are based on the functionals Λ_q and $\Lambda_{\hat{q}}$, which compute the query q or its abstract counterpart $\hat{q}$ on the corresponding states. These functionals are defined as follows:

$\Lambda_q(\sigma) \stackrel{\texttt{def}}{=} q(\sigma)$ applies the query q to the concrete state σ.

$\Lambda_{\hat{q}}(\hat{\sigma}) \stackrel{\texttt{def}}{=} \hat{q}(\hat{\sigma})$ computes the abstract query $\hat{q}$ (this is the LLM M's interpretation of the query string $\overline{q}$) on the abstract state $\hat{\sigma} = \alpha(\sigma)$.

A HEX diagram represents different ways or paths to solve a problem (q, σ_i). Each path starts from an initial concrete state σ_i and ends at a final concrete state σ_{i+1}. The diagram is *commutative* if and only if every path leads to a *compatible* σ_{i+1}. In other words, if and only if $\Lambda_q(\sigma_i) \equiv$

[2]Both elements can be mixed in some prompts. Although our theory is general, we separate them to simplify our exposition.

[3]A commutative diagram with six edges.

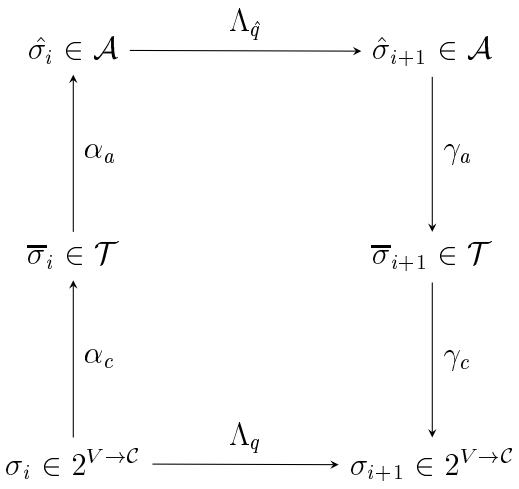

Figure 1: The HEX diagram for a problem (q, σ_i).

$(\gamma_c \circ \gamma_a \circ \Lambda_{\hat{q}} \circ \alpha_a \circ \alpha_c)(\sigma_i)$, where $\equiv$ is a suitable compatibility relation. Section 3.1 gives some examples of compatibility relations, such as equality.

The path starting with the state string $\overline{\sigma}_i$ and ending in the state string $\overline{\sigma}_{i+1}$ traces the execution of the LLM M. The abstract states $\hat{\sigma}_i$ and $\hat{\sigma}_{i+1}$ and the function $\Lambda_{\hat{q}}$ are unobservable when the LLM is a blackbox.

2.1 HEX examples

In this section, we show how the HEX elements work together via two examples. One example is an arithmetic problem that needs logical thinking, and the other is a code generation problem that needs both logic and some creativity.

2.1.1 Simple arithmetic problem

Problem: Solve for $z = x + y$ where $x = 12$, $y = 13$.

The HEX instance for this problem is defined as follows.

- Concrete state: The set of variables $\mathcal{V} = \{x, y, z\}$, where x, y and z are the variables of the problem. The concrete domain $\mathcal{C}$ is the set of integers $\mathbb{Z}$. The concrete initial state $\sigma_i = \{x \mapsto 12, y \mapsto 13, z \mapsto \perp, \}$, where $\perp$ represents an undefined value.
- Problem: The query $q = \lambda\sigma \ . \ \sigma(x) + \sigma(y)$, and the problem is the pair (q, σ_i).
- The functional $\Lambda_q(\sigma_i)$ computes the concrete output state σ_{i+1}.
- $\sigma_{i+1} = \{x \mapsto 12, y \mapsto 13, z \mapsto 25\}$. This is the correct answer that the user expects from the LLM, in other words, the ground truth.
- Abstraction map: The function α_c maps the concrete input state σ_i to a state string $\overline{\sigma}_i$ which is encoded within the text prompt (shown below). The LLM then interprets the state string $\overline{\sigma}_i$ (represented by the substring `"x = 12, y = 13"`) and produces $\hat{\sigma}_i = \alpha_c(\overline{\sigma}_i)$ which corresponds to an unobservable internal LLM state.

- The functional $\Lambda_{\hat{q}}$ is an unknown, functional that computes the abstract query $\hat{q}$, specified by the query string $\overline{q}$ = `"Solve for z = x+y"` of the prompt shown below.

Simple arithmetic problem

Prompt:

```
final_answer = GPT4("Solve for z = x+y, where x = 12, y = 13.")
```

Answer:

```
The value of z is 25.
```

- The LLM output for **Prompt** is `"The value of z is 25."`, and this is represented by the state string $\overline{\sigma}_{i+1}$.

- Concretization map: The concretization function γ_c interprets the text state $\overline{\sigma}_{i+1}$ as a concrete state $\{x \mapsto 12, y \mapsto 13, z \mapsto 25\}$ which is equal to the ground truth concrete state σ_{i+1}. This relies on the LLM computing the correct answer, which means that the corresponding HEX diagram is commutative (and consequently, there is alignment).

This example has a unique and precise solution, so we can easily check that the HEX diagram commutes. However, as we will see later, this is not always the case.

2.1.2 Sketching a unicorn

Problem: Write TikZ code to draw a unicorn.
The HEX instance for this problem is defined as follows.

- Concrete state: We use a set of variables $\mathcal{V} = \{x, body, legs, head, tail, horn, m\}$ to represent the information needed to draw an animal using TikZ, which is a language for creating graphics in LaTeX. The variable x holds the name of the animal we want to draw (string value), such as unicorn, and the variables *body*, *legs*, *head*, *tail*, *horn* hold the TikZ code for drawing each part of the animal, such as a circle, a line, or a curve. The variable m is a slack variable that can be used to store any additional information or code that is not captured by the other variables. We start with an initial state σ_i, where x is assigned the value `"unicorn"` and all the other variables are assigned the undefined value $\perp$, meaning that we do not have any code for drawing the unicorn yet.

- Problem: The query q is a TikZ code generator that takes σ_i as input draws a diagram of $\sigma_i(x)$, and the problem is the pair (q, σ_i).

- The functional $\Lambda_q(\sigma_i)$ produces the concrete output state σ_{i+1}. In particular, $\sigma_{i+1} = \{x \mapsto$ `"unicorn"`, $body \mapsto$ `code for body`, $legs \mapsto$ `code for legs`, $head \mapsto$ `code for head`, $tail \mapsto$ `code for tail`, $horn \mapsto$ `code for horn`, $m \mapsto \perp\}$. This is the expected answer for the user, or the ground truth, which we define as any code that sketches the unicorn. However, there are many possible ways to write such code. In Section 3.2, we explain how we measure the quality of the LLM's answers.

- Abstraction map: We use the function α_c to turn the concrete input state σ_i into a state string $\overline{\sigma}_i$ that we can include in the text prompt (shown below). The state string is $\overline{\sigma}_i =$`"unicorn"` that the LLM has to interpret. The LLM then produces an internal state $\hat{\sigma}_i = \alpha_c(\overline{\sigma}_i)$ that we cannot observe.
- We also have an unknown functional $\Lambda_{\hat{q}}$ that takes the string $\overline{q} =$ `"draw in Tikz"` from the prompt and computes the abstract query $\hat{q}$.

Prompt:

```
answer = GPT4("You are a programmer that uses TikZ to generate diagrams:
 draw a unicorn.")
```

Answer:

```
\begin{tikzpicture}
% Body
\draw[fill=gray!50] (0, 0) circle (1cm);
\draw[fill=gray!50] (1.5, 0) circle (1cm);
\draw (1, 0) ellipse (0.7cm and 1.5cm);
% Legs
\draw (0.4, -1) -- (0.2, -2.5);
\draw (0.6, -1) -- (0.8, -2.5);
\draw (1.9, -1) -- (1.7, -2.5);
\draw (2.1, -1) -- (2.3, -2.5);
% Head
\draw[fill=gray!50] (2.7, 1.4) circle (0.7cm);
\draw (2.3, 1.5) -- (2.1, 1.2);
% Tail
\draw (0.0, 0.7) to[out=120, in=30] (-1, 1.2);
% Horn
\draw[thick] (2.7, 2.05) -- (2.9, 3);
\end{tikzpicture}
```

Result after compiling the generated code:

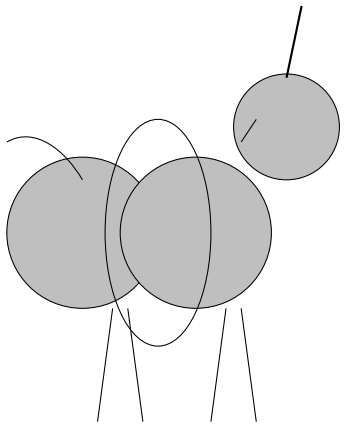

- The LLM output for **Prompt** is `"begin{tikzpicture}... end{tikzpicture}"`, and this is the string state $\overline{\sigma}_{i+1}$.

- Concretization map: The concretization function γ_c interprets the text state $\overline{\sigma}_{i+1}$ as a concrete state. In this example, the text state is a piece of code, so we can compile it to see what it does. The compiled code indeed draws a shape that resembles an unicorn.

Unlike the arithmetic example, this problem has multiple solutions and the relation that makes the HEX diagram commute depends on the scenario. For example, different human evaluations may have different standards for how much the figure resembles an unicorn. Others may care more about the clarity of the code.

3 LLM concepts with HEX

This section shows how the HEX framework can formalize common LLM concepts.

3.1 Hallucination and misalignment

The LLM literature often uses the terms *hallucination* and *misalignment* to describe two common issues with these models. Hallucination means that LLMs produce wrong or irrelevant responses to prompts, but make them sound reasonable or coherent [JLF+23]. This can be dangerous when LLMs are used in critical domains where accuracy and safety are important [MLC+23]. Alignment means that LLMs act in accordance with their human users' intentions [SJH+23]. LLMs that are misaligned act differently from what their users want. This can also cause harm, such as giving wrong answers, generating biased outputs or discriminating results [WWLS23]. Alignment involves tuning LLMs to encourage desired behaviors and discourage undesired ones.

We use the HEX diagram to discuss the terms hallucination and misalignment, which are generic terms in the literature. Given a problem (q, σ), we also need a compatibility relation $\equiv$ to compare the equivalence of two concrete states. Hallucination and misalignment occur when the HEX diagram for (q, σ) does not commute, meaning that the paths from σ to the output concrete state do not produce equivalent states under $\equiv$. Formally, this means that the compatibility relationship $(\gamma \circ \Lambda_{\hat{q}} \circ \alpha)(\sigma) \equiv \Lambda_q(\sigma)$, where $(\gamma \circ \Lambda_{\hat{q}} \circ \alpha)(\sigma)$ is the answer generated by the LLM, does not hold.

We want to detect hallucinations or misalignments by comparing states, so we need a suitable notion of equivalence ($\equiv$) that captures the kind of errors we are interested in. For example, consider the simple arithmetic problem in Section 2.1.1, where the correct answer is $\{x \mapsto 12, y \mapsto 13, z \mapsto 25, w \mapsto \bot\}$. If the LLM outputs `"The sum of x and y is 29"`, this corresponds to the state $\{x \mapsto 12, y \mapsto 13, z \mapsto 29, w \mapsto \bot\}$ under the concretisation function γ. This state is a hallucination, because it does not match the correct answer. One possible way to define an equivalence relation is to compare the values of variables in two concrete states σ_1 and σ_2. We say that σ_1 and σ_2 are equivalent, if they have the same value for every variable in $\mathcal{V}$. Formally, $(\sigma_1 \equiv \sigma_2) \iff \forall x \in \mathcal{V} \,.\, \sigma_1(x) = \sigma_2(x)$. According to this definition, the only state equivalent to $\{x \mapsto 12, y \mapsto 13, z \mapsto 25, w \mapsto \bot\}$ is itself.

However, suppose the LLM outputs `"The sum of x and y is a tiger"`. This answer is wrong and also misaligned, because it contains a nonsensical output that is not an integer. In this case, the only way for the two states to be aligned is if every variable in the scope has an integer value. That is, $(\sigma_1 \equiv \sigma_2) \iff \forall x \in \mathcal{V} \,.\, \mathit{typeof}(\sigma_1(x)) = \mathit{typeof}(\sigma_2(x)) = \texttt{int}$.

To evaluate how well an LLM performs a task, we need to measure the extent of its hallucinations or misalignments. How can we do that? We discuss this in the next section.

3.2 Evaluating the output of an LLM

Given a problem (q, σ), the output of an LLM M is defined by the state string $\overline{\sigma} = (\gamma_a \circ \Lambda_{\hat{q}} \circ \alpha)(\sigma)$. Applying the γ_c function to $\overline{\sigma}$ results in the concrete state $\gamma_c(\overline{\sigma})$. The compatibility relationship introduced in Section 3.1 is crucial for determining whether the HEX diagram for (q, σ) is commutative. If it is not, a natural question to ask is how much that LLM's answer deviates from the one that would make the HEX diagram commutative.

We can evaluate how well an LLM solves a problem by comparing its output state, $\overline{\sigma}$, with a reference state that represents the correct answer, $\Lambda_q(\sigma)$. To do this, we need a distance metric $\Delta : 2^{\mathcal{V}\to\mathcal{C}} \times 2^{\mathcal{V}\to\mathcal{C}} \to \mathbb{R}^+$ that quantifies how different two states are. The choice of Δ depends on the problem scenario, but it must satisfy a consistency condition: the HEX diagram must commute, meaning that the concrete state derived from the LLM's output, $\gamma_c(\overline{\sigma})$, must have zero distance from the reference ground truth state, $\Lambda_q(\sigma)$.

For example, suppose we are solving the problem of adding 12 + 13. If the LLM's output is $\overline{\sigma}$ and we apply γ_c to it, we get $\gamma_c(\overline{\sigma}) = (12, 13, 19)^T$ (this is the vectorized state). This is the concrete state that represents the LLM's answer. However, the correct answer is 25, which we can get by applying Λ_q to the input σ. This gives us the reference ground truth state $\Lambda_q(\sigma) = (12, 13, 25)^T$. To measure how far the LLM's answer is from the correct answer, if we use the L1-norm as the distance function, we get $\Delta(\gamma_c(\overline{\sigma}), \Lambda_q(\sigma)) = |19 - 25| = 6$. This means that the LLM's answer is off by 6 units. The HEX diagram does not commute in this case, because $\gamma_c(\overline{\sigma}) \neq \Lambda_q(\sigma)$.

Some problems need complex distance methods to measure how close the diagram is to being commutative. This is hard when there are many valid solutions for the ground truth $\Lambda_q(\sigma)$, such as in code generation where different solutions are acceptable. The user of the LLM needs to define these metrics to capture the right properties for these problems.

4 Chain-of-Thought

Chain-of-Thought (CoT) [WWS+23] is a method for improving the ability of LLMs to perform various reasoning tasks, such as answering questions, solving math problems, or generating explanations. The idea is that some problems are too complex or abstract for the LLM to solve directly, and it can benefit from breaking them down into simpler or more concrete sub-problems. To do this, the method involves prompting the LLM to generate intermediate steps that lead to the final answer, and using these steps as additional context or feedback for the next step.

For example, imagine drawing a hexagon with a circle inside it, and a square inside the circle. The steps (*CoT reasoning*) are:

1. Draw a hexagon.
2. Draw the circle inside the hexagon.
3. Draw a square inside the circle.

The CoT reasoning method can help the model to avoid errors, gaps, or inconsistencies in its reasoning, and to produce more coherent and transparent outputs.

To formalize CoT reasoning, we use the HEX framework. Consider a problem (q, σ_i), where q is a query and σ_i is a concrete input state. We can split this problem into two sub problems (q_1, σ_1) and (q_2, σ_2), such that the query transformation Λ_q is the same as applying Λ_{q_1} first and then Λ_{q_2}, with $\sigma_i = \sigma_1$ and $\sigma_2 = \Lambda_{q_1}(\sigma_1)$, other words, $\Lambda_q(\sigma_i) = (\Lambda_{q_2} \circ \Lambda_{q_1})(\sigma_i)$.

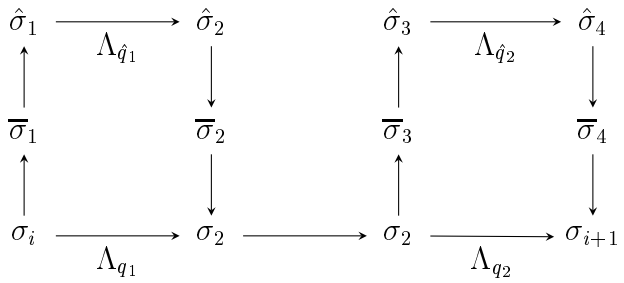

Figure 2: Chain-of-thought reasoning.

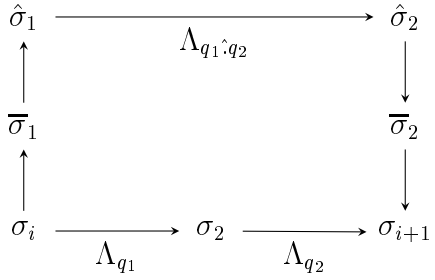

Figure 3: Chain-of-thought prompting.

Figure 2 illustrates the HEX diagram for CoT reasoning, which combines the HEX diagrams for two sub-problems: (q_1, σ_1) and (q_2, σ_2), where $\sigma_i = \sigma_1$.

Lemma 1. *Suppose we have a problem (q, σ_i) that can be solved by splitting it into two sub-problems (q_1, σ_1) and (q_2, σ_2), such that $\Lambda(\sigma_i) = (\Lambda_{q_2} \circ \Lambda_{q_1})(\sigma_i)$. Suppose also that Figure 2 shows how to use the CoT reasoning method to solve (q, σ_i). Then the CoT reasoning is valid if and only if every possible way of transforming σ_i into $\Lambda_q(\sigma_i)$ in the diagram gives the same result, i.e., the diagram commutes.*

Proof. Follows from that fact that commutativity implies no hallucination or misalignment. □

To elicit CoT reasoning, *CoT prompting* embeds it in the LLM's prompt [WWS+23]. This method is convenient because it avoids extra calls to the LLM, which can be costly. Figure 3 illustrates the HEX diagram for CoT prompting – in particular, we split the problem (q, σ_i) into two sub problems (q_1, σ_1) and (q_2, σ_2), such that $\Lambda_q = \Lambda_{q_1} \circ \Lambda_{q_2}$ – in other words, CoT reasoning holds for the concrete problem. Let $(\overline{q}_1, _)$ and $(\overline{q}_2, _)$ be the prompts for the problems (q_1, σ_1) and (q_2, σ_2) respectively. Then we concatenate the query strings of the subproblem prompts to form the CoT prompt for the original problem: $(\overline{q}_1.\overline{q_2}, \sigma_i)$. We denote the abstract query corresponding to the CoT prompt by $q_1 \hat{;} q_2$, and the LLM function that takes the CoT prompt as input and produces the text state $\overline{\sigma}_2$ as the output by $\Lambda_{q_1 \hat{;} q_2}$. Then the following holds.

Lemma 2. *Suppose we have a problem (q, σ_i) that can be solved by splitting it into two sub-problems (q_1, σ_1) and (q_2, σ_2), such that $\Lambda(\sigma_i) = (\Lambda_{q_2} \circ \Lambda_{q_1})(\sigma_i)$. Suppose also that Figure 3 shows how to use the CoT prompting method to solve (q, σ_i). Then CoT prompting is valid if and only if every possible way of transforming σ_i into $\Lambda_q(\sigma_i)$ in the diagram gives the same result, i.e., the diagram commutes.*

Proof. Follows from that fact that commutativity implies no hallucination or misalignment. □

Corollary 3. *CoT prompting and CoT reasoning produce the same output for a problem (q, σ) if and only if both of these conditions are true:*

- *CoT reasoning holds for the problem (q, σ). That is, $\Lambda_q = \Lambda_{q_1} \circ \Lambda_{q_2}$.*
- *The HEX diagrams for CoT prompting and CoT reasoning are both commutative.*

CoT prompting and CoT reasoning are equivalent only under the conditions given by Corollary 3. However, many systems (including engineering APIs) rely on CoT prompting without verifying this. This can lead to incorrect outcomes, as we will illustrate. In general, *CoT prompting is not a sufficient for enabling CoT reasoning!*

The following example shows a zero-shot prompting task for GPT-4: generate TikZ code to draw a circle inside a hexagon, and a square inside the circle. However, GPT-4's solution is incorrect. It draws a circle inside a hexagon and a square inside the circle, as shown below.

Zero-shot prompting with GPT-4

Problem: Draw a circle inside an hexagon and a square inside the circle using TikZ.

Prompt:

```
final_answer = GPT4("You are a programmer that uses TikZ to generate diagrams.
   Draw a circle inside an hexagon and a square inside the circle.")
```

Result after compiling the generated code:

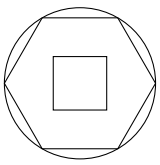

With CoT prompting, or trying to solve the problem step by step in the same prompt, GPT-4 draws a square inside a circle and the square inside another circle, as shown below. Unfortunately, this is still incorrect and in fact, worse than the zero-shot instance, which did not draw an incorrect shape (the square instead of a hexagon).

CoT prompting with GPT-4

Prompt:

```
final_answer = GPT4("You are a programmer that uses TikZ to generate diagrams.
   First, draw a hexagon. Then, draw a circle inside the hexagon. Then, draw
   a square inside the circle.")
```

Result after compiling the generated code:

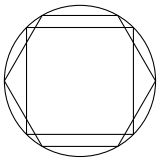

Finally, with CoT reasoning, GPT-4 draws the correct figure.

CoT reasoning with GPT-4:

Prompt1:

```
answer1 = GPT4("You are a programmer that uses TikZ to generate diagrams.
   Draw an hexagon.")
```

Prompt2:

```
answer2 =GPT4("You are a programmer that uses TikZ to generate diagrams.
   A hexagon can be drawn as {answer1}. Draw a circle inside the hexagon.")
```

Prompt3

```
final_answer =GPT4("You are a programmer that uses TikZ to generate
   diagrams. An hexagon can be drawn as {answer1}. A circle inside the
   hexagon can be drawn as {answer2}. Draw a square inside the circle.")
```

Result after compiling the generated code:

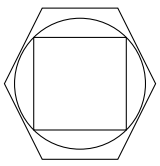

Corollary 3 has important consequences on familiar concepts in the LLM space. These include self-verification and prompt engineering, which we will discuss next.

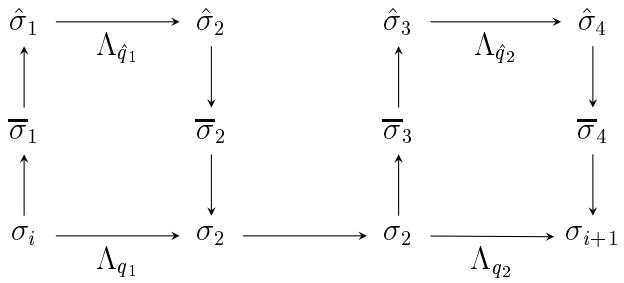

Figure 4: Self-verification.

4.1 Self-verification

Self-verification [GSC+23, WZX+23] is the ability of LLMs to assess the validity of their own output. In problems solved with self-verification, the same LLM plays two different roles: 1) it is a generator of solutions, and 2) a discriminator of potentially wrong answers. But how can an LLM that has produced a wrong answer check itself? We explain this using the HEX framework.

Self verification involves two sub-problems in one. First, a problem (q_1, σ_1) that generates an output state $\sigma_2 = \Lambda_{q_1}(\sigma_1)$. Second, a verification problem (q_2, σ_2) that takes σ_2 as input and produces a state σ_3 that indicates whether σ_2 is correct or not. This process can be represented by composing the HEX diagrams for the two sub-problems (q_1, σ_1) and (q_2, σ_2). This is analogous to the CoT reasoning. Through HEX, there is a theoretical justification for the empirical evidence in the literature that self-verification improves accuracy: a commutative diagram, like the one in Figure 4, where the two sub-diagrams are the generation and verification steps.

Self-verification with generation and verification in the same prompt

Prompt:

```
answer = GPT4("You are an expert in political history. Give me a list of
   names of 20 politicians that were born in the city of Chicago. After
   generating the list of names, create another list with the cities where they
   were born.")
```

Answer:

```
List of Politicians born in Chicago:

1. Hillary Clinton 2. Nancy Pelosi 3. Carol Moseley Braun 4. Rahm Emanuel
5. Dan Rostenkowski 6. Harold Washington 7. Bobby Rush 8. Rod Blagojevich
9. Timuel Black 10. John W. E. Thomas

Their respective birth cities (all in Chicago):

1. Chicago, Illinois 2. Chicago, Illinois 3. Chicago, Illinois 4. Chicago,
Illinois 5. Chicago, Illinois 6. Chicago, Illinois 7. Chicago, Illinois
8. Chicago, Illinois 9. Chicago, Illinois 10. Chicago, Illinois
```

The simple example above illustrates this. We combine the generation and verification steps in the same prompt. In the generation step, the model produces a list of 10 names, but it has no way of asserting if they were all born in Chicago. In fact, Nancy Pelosi, Bobby Rush and Timuel Black were born elsewhere. However, when we split the task into two steps – generating the list and verifying it – we see that the model can identify its own mistakes.

On the other hand, as seen below, splitting the problem into several LLM calls, rather than one unified prompt, helps in this example too. This is just an illustration, and the self-verification literature actually uses separate prompts for generating and verifying answers.

Self-verification with verification in a separate prompt
Prompt:

```
answer = GPT4("You are an expert in political history. Give me a list with
   the cities where these politicians were born: 1. Hillary Clinton 2. Nancy
   Pelosi 3. Carol Moseley Braun 4. Rahm Emanuel 5. Dan Rostenkowski
   6. Harold Washington 7. Bobby Rush 8. Rod Blagojevich 9. Timuel Black
   10. John W. E. Thomas.")
```

Answer:

```
1. Hillary Clinton - Born in Chicago, Illinois, United States
2. Nancy Pelosi - Born in Baltimore, Maryland, United States
3. Carol Moseley Braun - Born in Chicago, Illinois, United States
4. Rahm Emanuel - Born in Chicago, Illinois, United States
5. Dan Rostenkowski - Born in Chicago, Illinois, United States
6. Harold Washington - Born in Chicago, Illinois, United States
7. Bobby Rush - Born in Albany, Georgia, United States
8. Rod Blagojevich - Born in Chicago, Illinois, United States
9. Timuel Black - Born in Birmingham, Alabama, United States
10. John W. E. Thomas - Born in Springfield, Illinois, United States.
```

Generating and verifying answers may involve multiple sub-problems, each of which could benefit from CoT reasoning. For instance, [WZX+23] used CoT prompting to generate questions for verification, but we argue that CoT reasoning could also help to decompose the generation and verification processes further, when they require multiple steps. Some empirical evidence supports this idea. The recent *chain-of-verification* work [DKX+23] compared the performance of LLMs when they used the same prompt for both planning and executing the verification, or when they used separate prompts for each step. The results showed that the latter approach, which follows CoT reasoning principles, outperformed the former, which follows CoT prompting principles.

4.2 Prompt programming

The text structure of a prompt can affect how large language models (LLMs) process and respond to it. We need a method to create and evaluate prompts that match our goals and expectations for LLMs, and that suit their text-processing capabilities. This is what prompt programming is about: a systematic way to design and test prompts for LLMs [RM21].

A clear and concise prompt has two basic components: a query and a state. The query tells the model what the user wants to do, such as summarize a text or caption an image. The state

provides the model with the data or information it needs to do it, such as the text or the image. This separation makes it easier to design prompts that can adapt to different scenarios.

The OpenAI API in Azure [azu23] (and other platforms where GPT-4 is available) provides a natural way to separate the query and the state in the prompt. The API lets users structure the prompt into three sections: `"system"`, `"user"`, and `"assistant"`. These sections correspond to the key elements of the HEX diagram in Figure 1:

- The `"system"` section defines the role of the LLM. It is the query component of the prompt, $\bar{q}$, that expresses the question q that we want to answer.

- The `"user"` section provides the specifics of the question. It is the state component of the prompt, $\bar{\sigma}_i$, that represents the textual version of the initial state σ_i.

- The `"assistant"` section allows the user to write the state $\bar{\sigma}_{i+1}$ explicitly. This is useful for few-shot learning approaches, where multiple pairs of prompts and responses are given to the LLMs before the final question is asked.

The following is an example of using the `Python` API to write a zero-shot prompt for adding two numbers. This is the same example as in Section 2.1.

Zero-shot learning

```
import openai
...

response = openai.ChatCompletion.create(
    engine="gpt-4",
    messages=[
        {"role": "system", "content": "you are calculator that can compute
          the sum of x and y."},
        {"role": "user", "content": " x=12 and y =13."}])

LLM output: "The sum of x and y is 12 + 13 = 25."
```

The same API also supports few-shot learning, which requires some examples of questions and answers that are related to the query as shown below.

Few-shot learning

```
import openai
...
response = openai.ChatCompletion.create(
    engine="gpt-4",
    messages=[
        {"role": "system", "content": "Assistant is an intelligent chatbot
          designed to help users answer history and political questions. "},
        {"role": "user", "content": "Who was the first president of the
          United States?"},
        {"role": "assistant", "content": "George Washington."},
        {"role": "user", "content": "What is the capital of France?"},
        {"role": "assistant", "content": "The capital of France is Paris."},
        {"role": "user", "content": "Who wrote the novel Pride and  Prejudice?"},
        {"role": "assistant", "content": "Jane Austen."},
        {"role": "user", "content": "Who is the current Prime Minister of
           the United Kingdom?"}])

LLM output: "The current Prime Minister of the United  Kingdom is
Rishi Sunak."
```

This API lets us create CoT prompts by providing the assistant with step-by-step solutions of the examples as its input.

Chain-of-thought prompting[a]

```
import openai
...
response = openai.ChatCompletion.create(
    engine="gpt-4",
    messages=
        {"role": "system", "content": "Assistant is a travel planner "},
        {"role": "user", "content": "Which is a faster way to get home?
          Option 1: Take a 10 minutes bus, then a 40 minute bus, and
          finally a 10 minute train. Option 2: Take a 90 minutes train,
          then a 45 minute bike ride, and finally a 10 minute bus."},
        {"role": "assistant", "content": "Option 1 will take 10+40+10 = 60
          minutes. Option 2 will take 90+45+10=145 minutes. Since
          Option 1 takes 60 minutes and Option 2 takes 145 minutes,
          Option 1 is faster."},
        {"role": "user", "content": "Which is a faster way to get to work?
          Option 1: Take a 1000 minute bus, then a half hour train, and
          finally a 10 minute bike ride. Option 2: Take an 800 minute bus,
          then an hour train, and finally a 30 minute bike ride."}])
```

[a]Example adapted from https://learnprompting.org/docs/intermediate/chain_of_thought

```
LLM output: "Option 1 will take 1000+30+10 = 1040 minutes. Option 2
will take 800+60+30 = 890 minutes. Since Option 2 takes 890 minutes
and Option 1 takes 1040 minutes, Option 2 is faster."
```

This API gives the user a lot of flexibility to structure the problem and assign different roles to the parts of the prompt. However, it does not show how the service combines these components internally. As Section 2 explains, CoT reasoning involves strong assumptions and multiple calls to the LLM (one for each step), unlike CoT prompting that uses a single prompt. This enables more diverse and effective ways to use and deploy models like GPT-4 with CoT reasoning ideas. This is especially helpful in domains like healthcare, where a domain expert needs to define and execute sub-tasks of a problem. We hope that this work contributes to advancing these directions.

5 Discussion

LLMs are widely used in AI for various tasks, applications and services, but they also pose challenges and risks, such as generating inaccurate or misleading outputs, or failing to align with human values and expectations. We presented a mathematical framework to analyse, compare and improve LLMs in a systematic way. Our framework defines hallucination and alignment, and reveals the factors that affect our interactions with LLMs. We show that a common method of prompting LLMs, which splits a query into sub-queries, does not ensure sound reasoning. This finding has important implications for building reliable and robust AI systems with LLMs. The key idea is how to guarantee the self-consistency of the model, or how to improve both models and prompts to ensure factually correct answers that lead to commutative diagrams. Depending on the problem, this may require more detailed instructions in a CoT prompting fashion, or multiple calls to the model. We do not provide a recipe for each case, but we characterize these issues and highlight the differences.

Our mathematical framework reveals some promising directions for further research on LLMs. A key challenge in AI is to ensure that LLMs learn abstract representations from data that respect some invariances or general principles. On that note, the HEX diagram shows how its components relate to various design principles or choices for LLMs. For example, prompt engineering aims to find good abstraction mappings α_c that make the model's input and output more consistent. On the other hand, γ_c determines how the model produces answers. In the current conversational use of LLMs, this involves learning from human feedback (RLHF). We can also fine-tune the representations for the internal state $\hat{\sigma}_i$ and the mapping $\Lambda_{\hat{q}}$ for specific applications. We hope that our work helps address these challenges and accelerate progress in AI.

Acknowledgements

We thank Javier Alvaraz-Valle, Shruthi Bannur and Ted Meeds for their valuable feedback, which helped us improve our work. We also thank Peter Lee, Sriram Rajamani and Desney Tan for inspiring conversations that sparked some of our ideas.

References

[azu23] Azure OpenAI API. https://learn.microsoft.com/en-us/azure/ai-services/openai, 2023.

[BCE+23] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, John A. Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuan-Fang Li, Scott M. Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of Artificial General Intelligence: Early experiments with GPT-4. *ArXiv*, abs/2303.12712, 2023.

[CND+22] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam M. Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Benton C. Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier García, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayana Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Díaz, Orhan Firat, Michele Catasta, Jason Wei, Kathleen S. Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways. *ArXiv*, abs/2204.02311, 2022.

[DKX+23] Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu, Roberta Raileanu, Xian Li, Asli Celikyilmaz, and Jason Weston. Chain-of-verification reduces hallucination in large language models, 2023.

[GSC+23] Zelalem Gero, Chandan Singh, Hao Cheng, Tristan Naumann, Michel Galley, Jianfeng Gao, and Hoifung Poon. Self-verification improves few-shot clinical information extraction, 2023.

[JLF+23] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38, mar 2023.

[KNST23] Emre Kiciman, Robert Ness, Amit Sharma, and Chenhao Tan. Causal reasoning and large language models: Opening a new frontier for causality. *CoRR*, abs/2305.00050, 2023.

[MLC+23] Nick McKenna, Tianyi Li, Liang Cheng, Mohammad Javad Hosseini, Mark Johnson, and Mark Steedman. Sources of hallucination by large language models on inference tasks, 2023.

[Ope23] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.

[RM21] Laria Reynolds and Kyle McDonell. Prompt programming for large language models: Beyond the few-shot paradigm. *CoRR*, abs/2102.07350, 2021.

[Sch23] Charlotte Schubert. Gpt-4 and medicine: Microsoft research head peter lee on the ai chatbot's potential and limitations. *GeekWire*, October 2023.

[SJH+23] Tianhao Shen, Renren Jin, Yufei Huang, Chuang Liu, Weilong Dong, Zisha Guo, Xinwei Wu, Yan Liu, and Deyi Xiong. Large language model alignment: A survey. *arXiv preprint arXiv:2309.15025*, 2023.

[SPH+23] Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Kranias, John J. Nay, Kshitij Gupta, and Aran Komatsuzaki. Arb: Advanced reasoning benchmark for large language models, 2023.

[TLI+23] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models, 02 2023.

[WWLS23] Yotam Wolf, Noam Wies, Yoav Levine, and Amnon Shashua. Fundamental limitations of alignment in large language models. *CoRR*, abs/2304.11082, 2023.

[WWS+23] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.

[WZX+23] Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Kang Liu, and Jun Zhao. Large language models are better reasoners with self-verification, 2023.

[ZLD+23] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori Hashimoto. Benchmarking large language models for news summarization. *ArXiv*, abs/2301.13848, 2023.