

# EXPLAIN-QUERY-TEST: SELF-EVALUATING LLMs VIA EXPLANATION AND COMPREHENSION DISCREPANCY

**Saeid Asgari Taghanaki**  
Autodesk Research

**Joao Monteiro**  
Autodesk Research

## ABSTRACT

Large language models (LLMs) have demonstrated remarkable proficiency in generating detailed and coherent explanations of complex concepts. However, the extent to which these models truly comprehend the concepts they articulate remains unclear. To assess the level of comprehension of a model relative to the content it generates, we implemented a self-evaluation pipeline where models: (i) given a topic generate an excerpt with information about the topic, (ii) given an excerpt generate question-answer pairs, and finally (iii) given a question generate an answer. We refer to this self-evaluation approach as Explain-Query-Test (EQT). Interestingly, the accuracy on generated questions resulting from running the EQT pipeline correlates moderately with the model performance as verified by typical benchmarks such as MMLU-PRO. In other words, EQT’s performance is predictive of MMLU-PRO’s, and EQT can be used to rank models without the need for any external source of evaluation data other than lists of topics of interest. Moreover, our results reveal a disparity between the models’ ability to produce detailed explanations and their performance on questions related to those explanations. This gap highlights fundamental limitations in the internal knowledge representation and reasoning abilities of current LLMs. We release the code at <https://github.com/asgsaeid/EQT>.

## 1 INTRODUCTION

Large language models (LLMs) have achieved remarkable success in natural language processing tasks, including text generation, translation, and question answering (Devlin et al., 2019; Brown et al., 2020; OpenAI, 2023). State-of-the-art models such as GPT-4 (OpenAI, 2023), O1-preview, Claude (Anthropic, 2023), Gemini (AI, 2023), and Llama (Touvron et al., 2023) are capable of producing coherent and detailed explanations about a wide array of concepts.

A critical aspect of intelligence, both human and artificial, lies in the ability to understand and apply knowledge flexibly. While LLMs demonstrate remarkable prowess in generating detailed explanations of concepts, an important question arises: Does this ability reflect true comprehension, or is it simply a sophisticated form of pattern recognition? More specifically, when an LLM explains a concept, can it answer related questions derived from that explanation without direct access to the explanation during testing?

This investigation directly relates to the self-evaluation of LLMs, where models are tasked with assessing their own generated content’s alignment with their internal understanding. Self-evaluation is crucial for understanding whether LLMs possess genuine reasoning abilities or merely exploit correlations in training data. By focusing on the relationship between explanation generation and subsequent question answering, we aim to probe the depth of their internal knowledge and the robustness of their reasoning capabilities.

Understanding this disconnect is crucial for several reasons. First, the ability to explain concepts and correctly answer related questions is fundamental for applications in education, healthcare, and decision-making systems (Bommasani et al., 2021). For instance, an LLM used in education should not only provide clear explanations to students but also demonstrate understanding by accurately answering follow-up questions. Second, if models fail at this task, it highlights limitations in their internal knowledge representation and reasoning, signaling risks for high-stakes applications where reliability and understanding are paramount. Finally, this evaluation aligns with broader efforts to

ensure that AI systems exhibit true understanding rather than merely leveraging statistical correlations in data (Bender et al., 2021).

In this study, we propose a novel self-evaluation framework, Explain-Query-Test (EQT), to assess to what extent state-of-the-art LLMs can independently answer questions derived from their own explanations, without access to those explanations during testing. EQT is performed in three steps: (i) given a topic, a model generates an excerpt with information about the topic, (ii) given an excerpt, the same model then generates question-answer pairs, and finally (iii) a model is given a question and generates an answer.

By decoupling explanation from question answering, EQT tests the models’ internal knowledge, reasoning, and consistency, requiring them to rely on deeper comprehension rather than surface-level text patterns. This allows us to rigorously measure not just whether LLMs can generate plausible explanations, but whether they can independently apply their knowledge to novel yet related tasks that revolve around the same underlying knowledge. In Section 3.2, we provide a detailed mathematical justification for the EQT framework, emphasizing its robustness and consistency metrics.

This paper makes the following contributions:

- We introduce EQT, a novel framework to assess LLMs’ comprehension by evaluating their ability to explain concepts and independently answer questions about their generated explanations.
- We propose and analyze metrics, such as answer consistency and stability scores, to evaluate LLMs’ ability to reason across paraphrased and conceptually linked questions.
- We show that EQT has the potential to be used as a proxy to measure a LLM performance without having a test set.

## 2 RELATED WORK

Our work focuses on the self-evaluation of LLMs, specifically on the alignment between their generated explanations and their ability to answer questions derived from those explanations without access to the explanations during questioning. Below, we discuss related work.

**Explanation Generation in LLMs.** LLMs have demonstrated the ability to produce detailed and coherent explanations of complex topics, often surpassing human-level articulation in specific domains. Studies such as (Wiegrefe & Marasovic, 2021) and (Bansal & Verma, 2022) have focused on the use of LLMs to generate explanations that facilitate understanding and interpretability in downstream tasks. However, these works primarily evaluate the quality of explanations in isolation, without testing whether models can leverage these explanations for reasoning.

**LLM Question-Answering Capabilities.** Question-answering (QA) tasks have long been used as benchmarks to evaluate the reasoning and comprehension abilities of LLMs (Rajpurkar et al., 2016; Brown et al., 2020). Recent work highlighted the gap between LLMs’ ability to generate correct answers and their performance when faced with nuanced or paraphrased questions (Zhao et al., 2023), or with questions around knowledge not covered during training (Monteiro et al., 2024b;a). However, little attention has been paid to QA tasks derived directly from the models’ own explanations, which is the focus of our study.

**Self-Evaluation in LLMs.** Self-evaluation, where LLMs assess or critique their own outputs, has emerged as a promising approach to improve model reliability and reasoning (Lightman & Xu, 2023; Madaan et al., 2023). For example, Lightman & Xu (2023) proposed a framework for models to self-assess their responses, while Madaan et al. (2023) introduced methods for iterative refinement of model-generated content. Sonoda & Srinivasan (2024) introduced test data which they used to evaluate variations in textual consistencies across similar content to identify failures in LLMs’ self-evaluation. Alternatively, Xia et al. (2024) proposed the use of gaps in likelihood across sequentially obtained generations given a query as means for self-evaluation, but showed that doing so requires an external model to offer a discrepancy reference. Our work complements these studies by examining whether models can demonstrate consistency between their explanations and subsequent reasoning, providing a unique lens on self-evaluation without requiring additional test data or auxiliary models.

**Evaluating Consistency and Robustness.** Consistency across paraphrased inputs has been identified as a critical metric for evaluating the robustness of LLMs Wang et al. (2023); Li & Zhang (2023). Studies have shown that even state-of-the-art models often fail to maintain stable predictions when faced with semantically equivalent but syntactically varied inputs. We extend this line of inquiry by introducing a consistency metric specific to explanation-derived QA tasks, thereby contributing to the broader effort of evaluating and improving LLM robustness.

Our work bridges the gap between explanation generation, question-answering, and self-evaluation by examining whether LLMs can effectively leverage their own explanations to reason and answer related questions. This approach complements existing research while addressing an underexplored aspect of LLM capabilities.

### 3 EXPLAIN-QUERY-TEST

Explain-Query-Test (EQT) evaluates the ability of LLMs to explain concepts and then answer questions based on those explanations. The methodology is designed to work with any set of concepts, making it a flexible framework for assessing LLM performance. We also introduce metrics to measure the performance and consistency of the models across paraphrased questions. EQT is detailed in Algorithm 1.

---

#### Algorithm 1 Explain-Query-Test

---

**Require:** Set of concepts  $C = \{c_1, c_2, \dots, c_n\}$ , Language model  $LM$ , Number of paraphrases per question  $k$ , Number of questions per concept  $q$

- 1: Initialize  $Results \leftarrow \{\}$  {Stores results for each concept}
- 2: Initialize  $Questions \leftarrow \{\}$  {Stores generated questions}
- 3: **for** each concept  $c \in C$  **do**
- 4:   Prompt  $LM$  to provide a detailed explanation  $E_c$  for concept  $c$
- 5:   Store  $E_c$
- 6:   Generate  $q$  self-contained multiple-choice questions from  $E_c$
- 7:   **for** each generated question **do**
- 8:     Create  $k$  paraphrased versions of the question
- 9:   **end for**
- 10:   Store all questions and paraphrases in  $Questions[c]$
- 11: **end for**
- 12: **for** each concept  $c \in C$  **do**
- 13:   **for** each question and its paraphrases in  $Questions[c]$  **do**
- 14:     Prompt  $LM$  to predict answers for the original and paraphrased questions
- 15:     Record predictions  $P$
- 16:   **end for**
- 17: **end for**
- 18: **return**  $Results$  {Contains Accuracy and Stability for each concept}

---

#### 3.1 EXPLANATION AND QUESTION GENERATION

For each concept, the LLM is prompted to provide a detailed explanation. From this explanation, multiple self-contained multiple-choice questions are generated. Each question may include multiple correct options. To evaluate the robustness of the model, each question is paraphrased several times, resulting in a set of original and paraphrased questions. These paraphrased questions maintain semantic meaning while varying syntactically.

As will be discussed in Section 4.2, throughout this paper, we use the MMLU-PRO dataset (Wang et al., 2024) to cross-reference results, providing a benchmark for comparison. We note however that EQT can be applied to any set of concepts where explanations and questions can be generated.

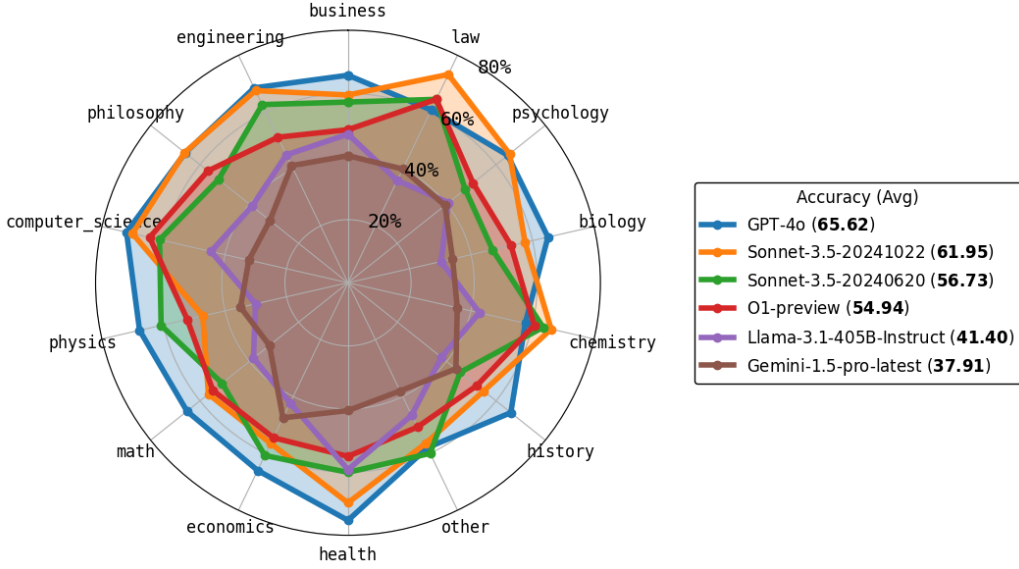


Figure 1: Comparison of EQT Accuracy across corresponding MMLU-Pro Categories.

### 3.2 A FORMAL DEFINITION OF EQT

The EQT framework evaluates the reasoning robustness of large language models (LLMs) by testing the consistency and coherence of their internal knowledge representation ( $K$ ). The framework assumes that even if intermediate artifacts, such as explanations ( $E$ ) and questions ( $Q$ ), are partially flawed, the full loop of Explain  $\rightarrow$  Query  $\rightarrow$  Test should reflect the model’s self-consistent reasoning. This section formalizes the theoretical foundation of the EQT framework.

#### 3.2.1 SELF-CONTAINED FEEDBACK LOOP

In the EQT loop, the model generates:

1. Explanations:  $E = g_\theta(K)$ , where  $g_\theta$  maps the model’s internal knowledge representation  $K$  to natural language explanations, given a topic.
2. Questions:  $Q = h_\theta(E)$ , where  $h_\theta$  transforms explanations into self-contained, logically valid questions.
3. Answers:  $A = f_\theta(Q)$ , where  $f_\theta$  represents the reasoning function that maps questions to answers.

Crucially, the EQT framework relies on the level of alignment between the model’s predicted answers ( $A$ ) and its internal knowledge representation ( $K$ ), irrespective of inaccuracies in  $E$  or  $Q$ :

$$f_\theta(h_\theta(g_\theta(K))) \approx K.$$

#### 3.2.2 FULL-LOOP ACCURACY

The framework evaluates the model’s reasoning robustness and self-consistency through full-loop accuracy, defined as:

$$\text{Acc}_{\text{loop}} = \frac{\text{Correct Answers Based on Self-Generated } Q}{\text{Total Questions}}.$$

Accurate models in the  $\text{Acc}_{\text{loop}}$  sense can reason consistently within their own generated context, even when intermediate outputs ( $E$  and  $Q$ ) are partially flawed (e.g., non-factual or non-grammatical). This metric reflects the internal knowledge consistency of the model.

### 3.2.3 SIGNIFICANCE OF HIGH AND LOW $\text{Acc}_{\text{Loop}}$

**High  $\text{Acc}_{\text{Loop}}$  Demonstrates Effective Reasoning and Self-Consistency.** *If  $\text{Acc}_{\text{Loop}}$  is high, it validates that EQT faithfully captures the model’s internal reasoning and consistency.*

Proof:

1. Consistency Across Artifacts: The EQT framework ensures that explanations ( $E$ ), questions ( $Q$ ), and answers ( $A$ ) are all derived from the model’s internal knowledge representation ( $K$ ). Formally:

$$E = g_{\theta}(K), \quad Q = h_{\theta}(E), \quad A = f_{\theta}(Q).$$

Even if  $E$  or  $Q$  contain minor flaws, their shared origin in  $K$  ensures alignment and consistency.

2. High  $\text{Acc}_{\text{Loop}}$  implies that:

$$f_{\theta}(h_{\theta}(g_{\theta}(K))) \approx K.$$

This reflects that the model’s internal knowledge representation is robust and self-consistent, as it can reproduce correct answers despite potential imperfections in intermediate artifacts.

3. Validation of Reasoning: By isolating the reasoning function  $f_{\theta}$  during the Test step, the EQT framework ensures that high  $\text{Acc}_{\text{Loop}}$  is a measure of the model’s internal reasoning rather than reliance on surface patterns or external artifacts.

**Low  $\text{Acc}_{\text{Loop}}$  Indicates Misalignment.** *If  $\text{Acc}_{\text{Loop}}$  is low, it highlights deficiencies in the EQT loop, signaling misalignment between generated artifacts ( $E$ ,  $Q$ ) and the internal knowledge representation ( $K$ ).*

Proof:

1. Low  $\text{Acc}_{\text{Loop}}$  implies that:

$$f_{\theta}(h_{\theta}(g_{\theta}(K))) \not\approx K.$$

This discrepancy can arise from:

- Inadequacies in  $g_{\theta}$ , leading to incomplete or incorrect explanations ( $E$ ).
  - Flaws in  $h_{\theta}$ , resulting in poorly constructed or ambiguous questions ( $Q$ ).
  - Weaknesses in  $f_{\theta}$ , reflecting limited reasoning or inference capabilities.
2. Implications for Model Design: Low  $\text{Acc}_{\text{Loop}}$  reveals areas where the model fails to maintain internal consistency or effectively utilize its knowledge representation. This highlights specific components ( $g_{\theta}$ ,  $h_{\theta}$ ,  $f_{\theta}$ ) that require improvement.

High  $\text{Acc}_{\text{Loop}}$  validates the EQT methodology as a robust measure of reasoning, ensuring internal knowledge consistency and alignment across self-generated artifacts. Conversely, low  $\text{Acc}_{\text{Loop}}$  acts as a diagnostic tool, pinpointing areas for improvement in the model’s reasoning and generation processes. By rigorously testing full-loop accuracy and answer consistency, the EQT framework evaluates reasoning robustness in language models. Through the isolation of the Test step and reliance on self-generated artifacts, EQT ensures that the evaluation reflects the model’s internal knowledge representation ( $K$ ) and reasoning ability ( $f_{\theta}$ ), independent of the external correctness of explanations or questions.

## 4 EXPERIMENTS

Experiments are conducted by relying on the set of topics within the widely popular MMLU-PRO dataset, which consists of 14 categories. For each such category, we extract 20 concepts. We then apply EQT for each concept. Questions derived from explanations are compared to the original MMLU-PRO dataset questions to assess the model’s ability to adapt to semantically similar yet syntactically varied prompts.

In further detail, for each category, LLMs generate explanations, and questions are derived based on those explanations. Each question is further paraphrased three times, resulting in a total of 20 questions per concept. This allows us to evaluate the stability of the models across syntactically

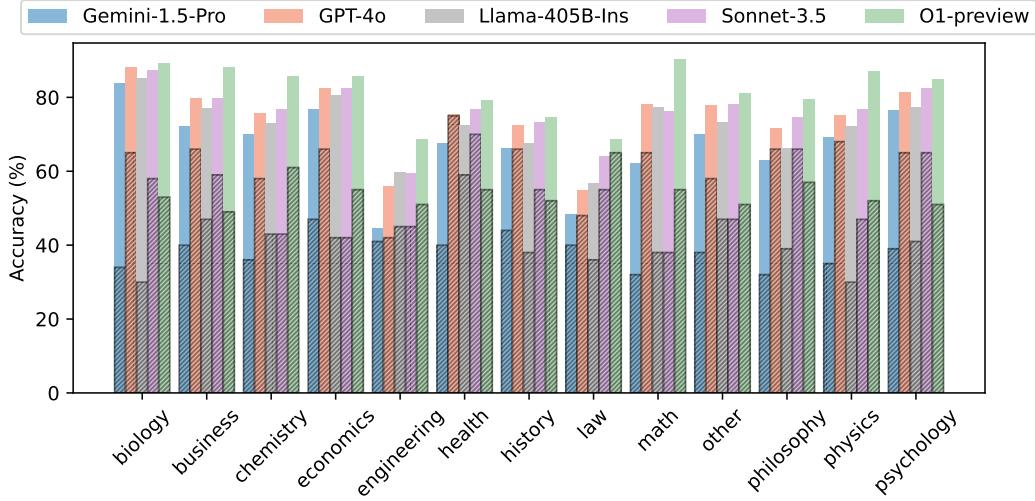


Figure 2: Accuracy (%) comparison across MMLU-PRO categories for various language models. Each model is represented by two bars: the first (solid) represents the original MMLU-PRO accuracy, and the second (hatched) indicates the adjusted accuracy due to the application of EQT since new questions are added.

diverse yet semantically equivalent prompts. The derived questions are compared to the models’ known performance on the original MMLU-PRO dataset to analyze consistency and accuracy. The categories in the dataset enable a detailed evaluation of model performance across a wide range of domains.

#### 4.1 METRICS

To evaluate the model’s ability to answer questions derived from its own explanations, we employ two metrics: question-answer accuracy, and a notion of consistency we define below.

**Answer Consistency Score (ACS):** To measure how stable the model’s predictions are across paraphrased questions, we define Answer Consistency Score (ACS). ACS evaluates to what extent the model provides consistent answers when the same question is phrased differently. Note that, in this context, *answers* refer specifically to the option letters (e.g.,  $A, B, C$ ) selected by the model, rather than the full answer text. Let:

- $A = [a_0, a_1, \dots, a_n]$ : The multiset (or list<sup>1</sup>) of predicted answers for a question and its  $n$  paraphrases, where  $a_0$  is the answer to the original question, and  $a_1, \dots, a_n$  are the answers to its paraphrases.
- $|A| = n + 1$ : The total number of answers.
- $U(A) = \{a_i \mid a_i \in A\}$ : The support of  $A$  given by the set of unique answers in  $A$ .
- $|U(A)|$ : The number of unique answers in  $A$ .

ACS is thus defined as:

$$\text{ACS} = 1 - \frac{|U(A)| - 1}{|A|}, \quad (1)$$

where higher values of ACS indicate greater consistency: *i.e.*,  $\text{ACS} = 1$  if and only if all answers are identical (maximum consistency), and  $\text{ACS} = 0$  if and only if all answers are unique (minimum consistency).

This metric highlights areas where the model’s behavior is inconsistent, providing insights into the reliability of LLMs when faced with paraphrased versions of questions.

<sup>1</sup>A multiset would be a *Counter* in pythonic jargon.

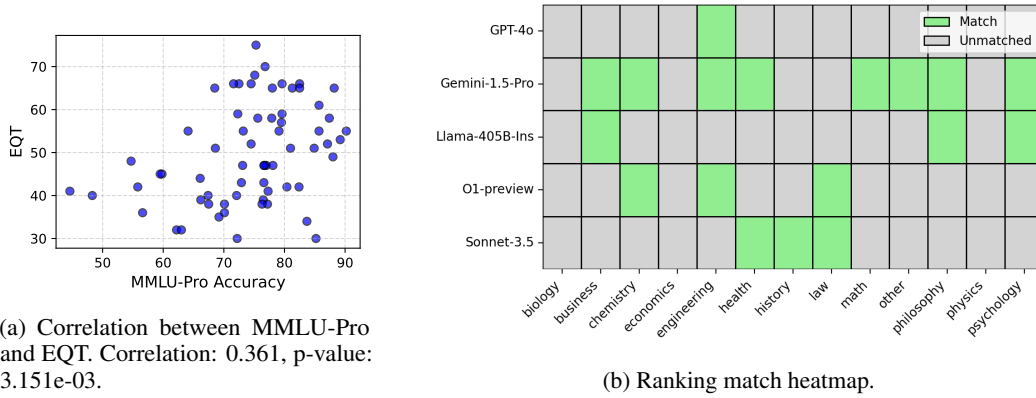


Figure 3: Analysis of MMLU-PRO and EQT results.

## 4.2 RESULTS

Figure 2 illustrates the accuracy of various language models on questions derived from the MMLU-PRO categories, as well as the corresponding performance drops compared to the original MMLU-PRO dataset. The solid bars represent the original accuracy on MMLU-PRO, while the hatched bars depict the adjusted accuracy when evaluated using EQT, which reflects the models’ concept comprehension.

Interestingly, we observe that models with higher initial accuracy on the original dataset tend to exhibit larger performance drops under EQT. For instance, categories such as *biology* and *psychology*, where models initially perform well, show significant degradation in accuracy. On the other hand, categories such as *law* and *engineering*, where models already exhibit lower baseline performance, experience smaller relative drops. This trend suggests that the drop in performance is influenced by the disparity between surface-level accuracy and the deeper understanding required to answer questions derived from explanations. Models may struggle to leverage the same high accuracy in the original dataset to maintain consistency in tasks that demand conceptual reasoning.

Overall, these results emphasize the challenges models face when answering questions derived from explanations and underline the need for improved reasoning capabilities. Figure ?? provides a spider plot visualization of the performance across different categories, offering a detailed view of the model’s strengths and weaknesses.

We further analyzed the relationship between the MMLU-PRO and EQT performances, and results in Figure 3a show a statistically significant correlation ( $r = 0.361$ ,  $p = 0.003$ ) between the two sets of results. In other words, results are moderately correlated, suggesting that EQT has predictive power of downstream reasoning performance and can serve as a proxy for assessing and ranking the performance of LLM. We highlight that the observed correlation suggests that reasoning-able LLMs can be compared without relying on external test sets, simplifying the evaluation process. By leveraging this proxy approach, we gain insights into LLM performance efficiently and consistently, reducing dependence on extensive datasets while maintaining meaningful performance comparisons.

Additionally, the ranking heatmap displayed in Figure 3b shows the consistency of model rankings across categories for both MMLU-PRO and EQT methods. Each row corresponds to a model, and each column corresponds to a category. Green cells indicate a match in model rankings between the two evaluation methods, while grey cells signify mismatches. Rankings are calculated independently for each method by sorting accuracy scores in descending order, with ties assigned the average rank. The heatmap highlights areas where EQT aligns with the established MMLU-PRO benchmark. Gemini-1.5-Pro has the most matches, while GPT-4o has the fewest, highlighting an interesting variation. Exploring these patterns further on additional benchmark datasets would be an interesting direction for future research.

Table 1: ACS stability scores (in  $[0 - 1]$ ) for each category and model. Higher scores indicate better stability.

Category	Gemini-1.5-Pro	GPT-4o	Sonnet-3.5	Llama-405B-Ins	O1-preview
business	0.85	<b>0.95</b>	0.96	0.96	0.89
law	0.85	<b>0.96</b>	0.97	0.92	0.93
psychology	0.84	0.94	<b>0.98</b>	0.92	0.89
biology	0.86	<b>0.95</b>	0.95	0.91	0.88
chemistry	0.87	0.96	<b>0.97</b>	0.91	0.91
history	0.89	<b>0.97</b>	0.96	0.92	0.92
other	0.86	0.92	<b>0.98</b>	0.95	0.90
health	0.86	0.96	<b>0.99</b>	0.92	0.93
economics	0.87	0.95	<b>0.98</b>	0.94	0.93
math	0.85	<b>0.97</b>	0.96	0.91	0.89
physics	0.86	<b>0.95</b>	0.94	0.88	0.90
computer_science	0.82	<b>0.97</b>	0.97	0.93	0.93
philosophy	0.83	0.94	<b>0.97</b>	0.95	0.89
engineering	0.84	0.96	<b>0.98</b>	0.91	0.92
Average	0.85	0.95	<b>0.97</b>	0.92	0.91

#### 4.3 STABILITY SCORES

To evaluate the consistency of the models’ predictions across paraphrased questions, we computed the ACS *stability scores* for each category and model. Table 1 summarizes these results. Higher stability scores indicate that the models provided consistent answers despite syntactical variations in the questions.

ACS scores indicate that models such as GPT-4o and Sonnet-3.5 consistently outperform others in maintaining answer consistency across paraphrased questions, especially in categories like health and engineering. Figure 1 provides a visual comparison of model performance in terms of accuracy and stability scores across the 14 categories in the MMLU-PRO. Overall, GPT-4o demonstrates superior performance compared to other methods, closely followed by the Claude-3.5-Sonnet variants. Interestingly, O1-preview, which is designed to excel in reasoning tasks, falls short of these models.

## 5 IMPLEMENTATION DETAILS AND PROMPT TEMPLATES

This section provides details on how the different steps in our methodology were implemented, along with the exact prompt templates used for explanation generation, question generation, and question paraphrasing. These prompts were carefully designed to ensure clarity, completeness, and consistency in the tasks performed by the language models (LLMs).

**Explanation Generation.** To obtain detailed explanations of concepts, we used the following prompt template:

#### Prompt for Explanation Generation

Please provide a comprehensive and detailed explanation of the concept ‘{concept}’, including its background, key principles, applications, and examples. Ensure the explanation is thorough and covers all essential aspects.

This prompt ensures that the LLM provides a complete and self-contained explanation of the given concept.

**Question Generation.** For generating questions from the explanations, the following prompt template was used:



**Prompt for Question Generation**

Create a multiple-choice question about the following concept: [CONCEPT]. Use this question type: "{question\_type}"

**Requirements:** 1. The question should have {num\_options} options (A to {last\_choice}). 2. There should one or several correct answers. 3. Base the question and all options ONLY on the information provided below. 4. Make the question entirely self-contained. Do NOT refer to any explanation, provided information, or external context. 5. Avoid phrases like "according to the text", "as described", or any similar references. 6. Ensure the question and options are clear and complete on their own.

**Format Example:** Question: [Your question here]

Options:

A) Option A text

B) Option B text

...

Correct Answers: [List of correct option letters, e.g., A, C, F]

Information about [CONCEPT]: {explanation}

This prompt ensures that the generated questions are accurate, comprehensive, and fully independent of the source material while remaining aligned with the concept’s explanation.

**Question Paraphrasing.** To generate paraphrased versions of the questions, the following prompt was used:

**Prompt for Question Paraphrasing**

Paraphrase the following question without changing its meaning. Ensure the paraphrased question is self-contained and does not reference any previous explanation or use phrases like "as mentioned earlier". ONLY generate the paraphrased question itself, and do not include any extra text such as "Here’s a paraphrased version of the question:" or similar.  
{question\_text}'

This prompt ensures that the paraphrased questions retained their original meaning while introducing syntactic variations.

**Question Answering.** Finally, for answering the questions, we used the following prompt:

**Prompt for Question Answering**

{question}

Options:

A) Option A text

B) Option B text

...

Please select all correct options (e.g., A, C, D) and provide your answer in the format: "Answer: [Your selections]".

This prompt instructs the LLM to provide answers in a clear format while respecting the requirements for multiple correct options.

**Summary of Design Choices.** Each prompt was crafted with specific guidelines to ensure:

- **Completeness:** All outputs are self-contained and do not rely on external context.
- **Clarity:** Prompts clearly define the task and expected format.
- **Robustness:** Iterative refinements were applied to address ambiguities or errors in generated content.

These carefully designed prompts form the backbone of our methodology and enable the evaluation of LLMs’ capabilities in explanation generation, question-answering, and consistency.

## 6 CONCLUSION

In this study, we investigated the extent to which large language models (LLMs) can comprehend and leverage their own explanations to answer related questions. Using the EQT approach we introduced, we evaluated models by prompting them to generate detailed explanations and then testing their ability to answer derived questions independently.

The results revealed a significant gap between the models’ ability to generate coherent explanations and their performance on questions derived from those explanations. This discrepancy highlights fundamental limitations in the internal knowledge representation and reasoning capabilities of current LLMs. Furthermore, our evaluation of answer consistency through the Answer Consistency Score (ACS) demonstrated that even state-of-the-art models struggle to maintain consistency across paraphrased questions. Interestingly, we identified EQT’s results have predictive power with respect to MMLU-PRO performance, yielding potential data-less self-evaluation framework for models able to reason.

These findings emphasize the need for further advancements in LLMs to improve their reasoning and understanding. Future work could explore techniques to align explanation generation with robust question-answering capabilities, develop training objectives that prioritize internal consistency, and design benchmarks that better reflect real-world reasoning challenges. By bridging the gap between explanation generation and question-answering, we can make significant strides toward building language models that not only articulate knowledge effectively but also demonstrate a deeper, more reliable understanding of the concepts they explain.

## REFERENCES

- Google AI. Introducing gemini: Google’s multimodal ai model. <https://ai.googleblog.com/2023/05/introducing-gemini-googles-multimodal.html>, 2023.
- Anthropic. Introducing claude: Anthropic’s ai assistant. <https://www.anthropic.com/index/claude>, 2023.
- Trisha Bansal and Prashant Verma. Explainability in ai: A survey. In *Proceedings of the 2022 International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623, 2021.
- Rishi Bommasani et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- Tom B Brown, Benjamin Mann, Nick Ryder, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *NAACL-HLT*, 2019.
- Qi Li and Wei Zhang. Evaluating para-consistency in large language models: A novel benchmark. *arXiv preprint arXiv:2305.05678*, 2023.
- Eli Lightman and Yao Xu. Self-evaluation in large language models: Bridging the gap between generation and understanding. *arXiv preprint arXiv:2304.01234*, 2023.
- Aman Madaan, Yujie Wu, and Lianmin Gao. Self-critique training: Improving large language models through iterative feedback. *arXiv preprint arXiv:2306.04500*, 2023.
- Joao Monteiro, Etienne Marcotte, Pierre-Andre Noel, Valentina Zantedeschi, David Vazquez, Nicolas Chapados, Christopher Pal, and Perouz Taslakian. Xc-cache: Cross-attending to cached context for efficient llm inference. *arXiv preprint arXiv:2404.15420*, 2024a.

Joao Monteiro, Pierre-Andre Noel, Etienne Marcotte, Sai Rajeswar, Valentina Zantedeschi, David Vazquez, Nicolas Chapados, Christopher Pal, and Perouz Taslakian. Replika: A question-answering dataset for benchmarking llms on unseen reference content. *arXiv preprint arXiv:2406.11811*, 2024b.

OpenAI. Gpt-4 technical report. <https://openai.com/research/gpt-4>, 2023.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2383–2392, 2016.

Ryosuke Sonoda and Ramya Srinivasan. A statistical analysis of llms’ self-evaluation using proverbs. *arXiv preprint arXiv:2410.16640*, 2024.

Hugo Touvron et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Xinyi Wang, Haoran Liu, and Jared Johnson. Robustness of language models to syntactic variations: An empirical study. *arXiv preprint arXiv:2302.06789*, 2023.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*, 2024.

Sarah Wiegrefe and Ana Marasovic. Teach me to explain: A review of machine learning interpretability through explanations. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 2055–2071, 2021.

Tingyu Xia, Bowen Yu, Yuan Wu, Yi Chang, and Chang Zhou. Language models can evaluate themselves via probability discrepancy. *arXiv preprint arXiv:2405.10516*, 2024.

Jingyuan Zhao, Fei Wang, and Yichen Li. Evaluating paraphrase sensitivity in large language models. In *Proceedings of the 2023 Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.

## A APPENDIX

Here is a sample explanation, questions, and paraphrases along with correct answers and predictions provided by GPT-4o for concept `graph theory` for category `math`.

### Math: Graph Theory

**Explanation:** Graph theory is a branch of mathematics that studies the properties and applications of graphs, which are mathematical structures used to model pairwise relations between objects. It has become a fundamental area of study in discrete mathematics and has applications across various fields, including computer science, biology, social sciences, and more.

**Background:** The origins of graph theory can be traced back to the 18th century with the work of Leonhard Euler. Euler’s solution to the Königsberg Bridge Problem in 1736 is often considered the first theorem of graph theory. The problem involved finding a walk through the city of Königsberg that would cross each of its seven bridges exactly once. Euler abstracted the problem into a graph, where landmasses were vertices and bridges were edges, leading to the formulation of Eulerian paths and circuits.

#### Key Principles:

- **Graphs and Their Components:**

- **Vertices (or Nodes):** Fundamental units or points in a graph.
- **Edges (or Links):** Connections between pairs of vertices.
- **Directed vs. Undirected Graphs:** Edges in directed graphs indicate one-way relationships, while undirected graphs represent two-way relationships.

- **Types of Graphs:**

- **Simple Graphs:** No loops or multiple edges between the same pair of vertices.
- **Multigraphs:** Allow multiple edges between the same pair of vertices.
- **Weighted Graphs:** Edges have weights representing costs, distances, or other metrics.
- **Bipartite Graphs:** Vertices can be divided into two disjoint sets such that every edge connects a vertex in one set to a vertex in the other.

- **Graph Properties:**

- **Degree:** Number of edges incident to a vertex.
- **Path:** Sequence of edges connecting a sequence of vertices.
- **Cycle:** Path that starts and ends at the same vertex without repeating edges or vertices.
- **Connectedness:** Graph is connected if there is a path between any pair of vertices.

**Original Question:** Explain the key characteristics of graph theory as a mathematical concept.

**Options:**

- A) Graph theory involves the study of vertices and edges to model relationships between objects.
- B) It exclusively focuses on weighted graphs where edges represent costs or distances.
- C) Graph theory originated with the solution to the Königsberg Bridge Problem by Euler.
- D) It includes the study of complete graphs, where each pair of vertices is connected by an edge.
- E) Directed graphs in graph theory have edges that indicate a two-way relationship.
- F) Graph theory is primarily used in biology and has limited applications in computer science.
- G) Concepts like cycles and paths are fundamental to understanding graph connectedness.
- H) Graph theory does not consider the use of algorithms for exploring graph structures.
- I) Multigraphs in graph theory can have multiple edges between the same pair of vertices.
- J) It is a modern mathematical field developed in the late 20th century.

**Paraphrased Questions:**

1. What are the key features of graph theory, and how is it defined as a branch of mathematics?
2. Can you describe the essential characteristics of graph theory and its fundamental principles?
3. How would you explain graph theory as a mathematical framework for understanding relationships and connections?

**Correct Answers:** A, C, D, G, I.

**Predicted Answers by GPT-4o:** A, C, D, G, I.