Perception of Knowledge Boundary for Large Language Models through Semi-open-ended Question Answering

Zhihua Wen, Zhiliang Tian[*]**, Zexin Jian, Zhen Huang,**
Pei Ke, **Yifu Gao**, **Minlie Huang**, **Dongsheng Li**[*]
College of Computer, National University of Defense Technology, Hunan, China
Tsinghua University, Beijing, China
{zhwen, tianzhiliang, gaoyifu, jianzexin21,
huangzhen, dsli}@nudt.edu.cn, aihuang@tsinghua.edu.cn, kepei1106@outlook.com

Abstract

Large Language Models (LLMs) are widely used for knowledge-seeking purposes yet suffer from hallucinations. The knowledge boundary of an LLM limits its factual understanding, beyond which it may begin to hallucinate. Investigating the perception of LLMs' *knowledge boundary* is crucial for detecting hallucinations and LLMs' reliable generation. Current studies perceive LLMs' knowledge boundary on questions with a concrete answer (close-ended questions) while paying limited attention to *semi-open-ended questions* that correspond to many potential answers. Some researchers achieve it by judging whether the question is answerable or not. However, this paradigm is not so suitable for semi-open-ended questions, which are usually "partially answerable questions" containing both answerable answers and ambiguous (unanswerable) answers. Ambiguous answers are essential for knowledge-seeking, but they may go beyond the knowledge boundary of LLMs. In this paper, we perceive the LLMs' knowledge boundary with semi-open-ended questions by discovering more ambiguous answers. First, we apply an LLM-based approach to construct semi-open-ended questions and obtain answers from a target LLM. Unfortunately, the output probabilities of mainstream black-box LLMs are inaccessible to sample for low-probability ambiguous answers. Therefore, we apply an open-sourced auxiliary model to explore ambiguous answers for the target LLM. We calculate the nearest semantic representation for existing answers to estimate their probabilities, with which we reduce the generation probability of high-probability existing answers to achieve a more effective generation. Finally, we compare the results from the RAG-based evaluation and LLM self-evaluation to categorize four types of ambiguous answers that are beyond the knowledge boundary of the target LLM. Following our method, we construct a dataset to perceive the knowledge boundary for GPT-4. We find that GPT-4 performs poorly on semi-open-ended questions and is often unaware of its knowledge boundary. Besides, our auxiliary model, LLaMA-2-13B, is effective in discovering many ambiguous answers, including correct answers neglected by GPT-4 and delusive wrong answers GPT-4 struggles to identify.

---

[*] Corresponding Authors.

1 Introduction

Large language models (LLMs) have revolutionized our interactions with AI, enabling users to acquire knowledge by posing questions in natural language [12, 6]. However, LLMs are prone to hallucination and generate non-factual responses, hindering the development of trustworthy AI.

One main cause of LLM hallucination is its unfamiliarity with the long-tail knowledge that appears less frequently than common-sense knowledge in the training data. To alleviate this issue, many researchers collect more domain-specific training data [22] or incorporate external information [31, 24] via retrieval-augmented generation (RAG) during inference. Another line of work investigates the perception of `knowledge boundaries` for LLMs, which indicates the extent of knowledge that the LLM can grasp well, beyond which it may begin to hallucinate [17]. Studying the perception of knowledge boundaries for LLMs helps alleviate hallucinations in many ways. For example, 1) It helps detect the hallucinations of a target LLM and the extent and scope of its factual knowledge [15, 39]. 2) It helps align LLMs for more honest generation [37, 36].

Existing studies on the perception of knowledge boundaries are primarily in the form of Question-Answering (QA). Their methods mainly aim to judge whether a question is answerable or unanswerable, and regard their border as the knowledge boundary. An answerable question refers to when the LLM is capable of generating a response matching the ground truth, and conversely, an unanswerable question means unable to answer correctly. These studies can be categorized into two groups. Prompt-based perception employs prompt engineering [39, 13] to assess whether the LLM can answer the question via LLM self-evaluation. They question whether the LLM knows the answer [27, 8, 40] or needs external knowledge to answer the question [31]. As LLMs tend to be overconfident [24, 27, 16, 43], more researchers explore representation-based perception. These studies optimize different representations for answers with different answerability [13, 10, 31, 43] or extract representations from a fixed encoder to train a classifier [20].

However, directly discriminating questions into answerable and unanswerable ones may not apply to some partially answerable questions. In many scenarios, the questions are relatively open-ended (i.e. having a list of correct answers) that may include (1) a subset of easy answerable answers, and (2) a subset of hard and unpopular answers, which may be unanswerable. These questions (referred to as "semi-open-ended questions") are particularly challenging and knowledge-extensive. Investigating the ambiguous answers to these semi-open-ended questions in various fields benefits knowledge-seeking. Ambiguous answers often go beyond the knowledge boundaries of LLMs and could lead to misinformation (see Sec. 4.6). Therefore, we argue that investigating these questions with their ambiguous answers can augment the perception of the knowledge boundaries for LLMs.

In this paper, we propose to perceive the knowledge boundary for a target LLM with semi-open-ended questions by discovering pieces of unfamiliar knowledge where the LLM learns badly. Particularly, We first construct a dataset with semi-open-ended questions on the open domain and query the target LLM for their corresponding answers. We define the low-probability correct answers and delusive incorrect answers are the ambiguous answers corresponding to the LLM's unfamiliar knowledge.

A challenge is that obtaining LLMs' low-probability answers needs accessing LLMs' output probabilities (or violently sampling LLMs' outputs many times to approximate the probabilities), which is inaccessible (or expensive) for mainstream black-box LLMs, i.e. GPT-4. Therefore, we approximate the generation probabilities of the target LLM with an open-sourced auxiliary model. We use the Pseudo-inverse of model embedding to estimate the nearest semantic representation for the existing answers. Consequently, we obtain the probability distribution of existing answers and repetitively filter the existing answers (and their semantic-related counterparts) to obtain answers with low-probabilities. Finally, we recognize answers beyond the knowledge boundary of the target LLM by comparing its self-evaluation results against the ground truth answers obtained from RAG-based evaluation.

Empirically, we use our method to construct a dataset of approximately 1k samples and evaluate GPT-4's performance. We find that GPT-4 makes mistakes in 82.90% of questions and 40.15% of its ambiguous answers generated are unqualified. Besides, GPT-4 also makes inaccurate self-evaluation 28.77% of the time, indicating that these are beyond the knowledge boundary of GPT-4. Moreover, we find nearly 50% of the candidate answers discovered by our auxiliary model, LLaMA-2-13B, are also beyond the knowledge boundary of GPT-4, including both factual answers that GPT-4 fails to produce and delusive wrong answers GPT-4 evaluates incorrectly.

Our contributions are as threefold: (1) We are the first to investigate the importance of semi-open-ended questions to the perception of knowledge boundaries for LLMs. (2) We propose an ambiguous answer discovery strategy that discovers many ambiguous answers with pieces of knowledge that are beyond the LLM's knowledge boundary. (3) Experimental results show the poor performance of an advanced LLM, GPT-4, on semi-open-ended questions and the effectiveness of our ambiguous answer discovery method in finding more pieces of knowledge which the LLMs are unfamiliar with.

2 Related Work

2.1 Perception of Knowledge Boundaries for LLMs

Existing studies on the perception of knowledge boundaries for LLMs can be categorized into prompt-based perception and pattern-based perception. Prompt-based perception perceives the knowledge boundary by querying the target LLM. Many researchers instruct the LLM before and after response generation, asking whether it can correctly answer the questions [27, 31, 8, 40] and if the generated answers are accurate [27, 37]. In addition, Yin et al. (2024) seek the optimal prompt for benchmarking LLM knowledge boundaries. Amayuelas et al. (2023) study the LLMs' ability to understand their knowledge and measure their uncertainty. Kadavath et al. (2022) also instruct LLMs to generate their confidence score for their responses. As studies find that LLMs tend to be overconfident [24, 27, 16, 43], many researchers explore representation-based perception. Researchers identify unknown questions [13] or evaluate correct and incorrect answers [10] by implicitly learning their different representations. Chen et al. (2023) train LLMs to identify incorrect answers via parameter-efficient tuning. Besides, Wang et al. (2023) extract representations of answerable and unanswerable questions to train a classifier to predict whether a question is answerable and assume questions with similar representations share the same answerability. Si et al. (2023) take token probability as the answer's confidence score during generation. Zhao et al. (2023) detect unanswerable questions by paraphrasing questions and checking the divergence of their answer distribution. The above studies primarily perceive knowledge boundaries for LLMs by distinguishing between answerable and unanswerable questions. This type of binary division does not apply to questions with both common easy answers and unpopular hard answers. Our study is the first to investigate the perception of knowledge boundaries on semi-open-ended questions.

2.2 Questions Answering for LLMs

Existing studies on Question Answering (QA) can be categorized into open-ended QA and close-ended QA based on the type of questions. Close-ended questions correspond to a limited number of correct answers, usually in the form of yes or no, true or false, or multiple-choice options, constraining the answers to a predetermined answer set [26, 14, 34, 28]. In addition, Researchers also study open-ended questions that allow the respondent to provide a more detailed and subjective response such as personal opinions and explanations [19, 34, 4, 3].

Researchers study the performance of LLMs on QA tasks mainly through various prompting strategies. Wei et al. (2023) explore "Chain-of-Thought" prompting (CoT), a simple and broadly applicable method for enhancing question answering ability of LLMs. Yao et al. (2023) and Besta et al. (2024) introduce similar frameworks for more complex QA tasks, namely, "Tree of Thought" and "Graph of Thoughts" prompting. As studies show that relying solely on an LLM's internal knowledge may lead to hallucinations [23], many researchers have also improved model performance in QA by incorporating external information (RAG systems [41] and knowledge graph [11, 1]). More recently, researchers have studied adaptive retrieval to avoid misinformation in the retrieved documents [35]. Ni et al. (2024) estimate the answerability of the given question and determines whether to retrieve [10, 24]. Xu et al. (2024) learn to identify the knowledge boundaries of LLMs and refuse to answer certain questions to avoid risks [13, 36].

3 Perception of Knowledge Boundary for LLM via Semi-open-ended QA

3.1 Overview

Our framework consists of three parts (see Fig. 1). We first exploit the instruction-following ability of a strong LLM to create a dataset consisting of semi-open-ended questions on various domains with

LLM's answers. To discover more pieces of unfamiliar knowledge for the target LLM, we apply an open-sourced auxiliary model to incur more ambiguous answers by encouraging more distinctive generations. Finally, we evaluate whether the ambiguous answers to each question are beyond the knowledge boundary of the target LLM by comparing the self-evaluation results against RAG-based evaluation.

3.2 Semi-open-ended QA Dataset

3.2.1 Dataset Construction

To study the performance of semi-open-ended questions across various domains, we employ an LLM-based 2-step approach to obtain semi-open-ended questions and collect answers.

- **Domain selection.** We first prompt the LLM to generate a list of domains, encompassing world knowledge, which includes areas such as biology, geology, music, etc.

- **Question generation.** We prompt the LLM multiple times under each domain to generate a set of semi-open-ended questions Q. To ensure the quality of the generation quality, we provide human-written sample questions as demonstrations and specify the following requirements for the generation of candidate questions: 1. The question should correspond to multiple correct answers, making it challenging to answer. The question should also be relatively easy for non-expert users to understand. 2. The judgment of question answers should be based on objective standards in the real world instead of the subjective standards of the evaluator. 3. The truthfulness of an answer to a question should not change constantly over time. 4. The questions share the same template: *Tell me a list of ....* We use the same vanilla prompt to eliminate the influence of different question styles.

- **Answer collection.** For each question q in Q, we query the LLM I times and collect all responses $\mathsf{A} = \{a_0, a_1, ..., a_{I-1}\}$. In the i-th interaction, we inform the LLM of all previously generated responses $\mathsf{A}[:i]$ and obtain a_{i-1} by querying the LLM with question q', which repeats the same criteria specified in q and highlights the need for more answers.[2]. Finally, we extract all answer entities in A and construct an answer list A.

3.2.2 Dataset Descriptions

We create a dataset to investigate the performance of mainstream LLM, GPT-4, on semi-open-ended questions. In dataset construction, we set I to 3 to exploit the knowledge of GPT-4 through multi-round conversations. Like humans, when faced with such questions (e.g., *What are the animals unique to Australia?*), LLMs tend first to give answers in which they hold high confidence (like the *red angaroo*). The latter answers are less certain and may have more mistakes (like *echidna*). We define the initial 75% of answer entities from GPT-4 generations as common-sense answers, while the remaining 25% as ambiguous answers. Our dataset comprises 953 questions covering 32 domains, with well-distributed data within each domain. On average, GPT-4 yields 52 answers for each question, including an average of 13 ambiguous answers. See the data samples in App. B. See the full prompts and demonstrations in App. C.

3.3 Ambiguous Knowledge Discovery

We apply an open-sourced auxiliary LLM to effectively discover more ambiguous answers that may be beyond the knowledge boundaries for black-box LLMs. Our intuition is that low-probability ambiguous answers reflect LLM's unfamiliarity with certain pieces of knowledge. However, it is challenging to collect low-probability ambiguous answers as 1) the generation probability of black-box LLMs (e.g. GPT-4) are inaccessible and their hyper-parameters (e.g. temperature) cannot target ambiguous answer-related tokens (see Sec. 4.2). 2) violently prompting the LLM with question q many times to approximate the generation probability is inefficient as it prioritizes high-probability answers that may already be present in existing answers. Answers that are semantically similar to high-probability answers also tend to have a high probability during generation. Hence, we propose to prevent the generation of high-probability answers with their semantic-related counterparts on an

---

[2]For example, if q is: *Tell me a list of animals unique to Australia.*, then q' is *Tell me more animals unique to Australia.*
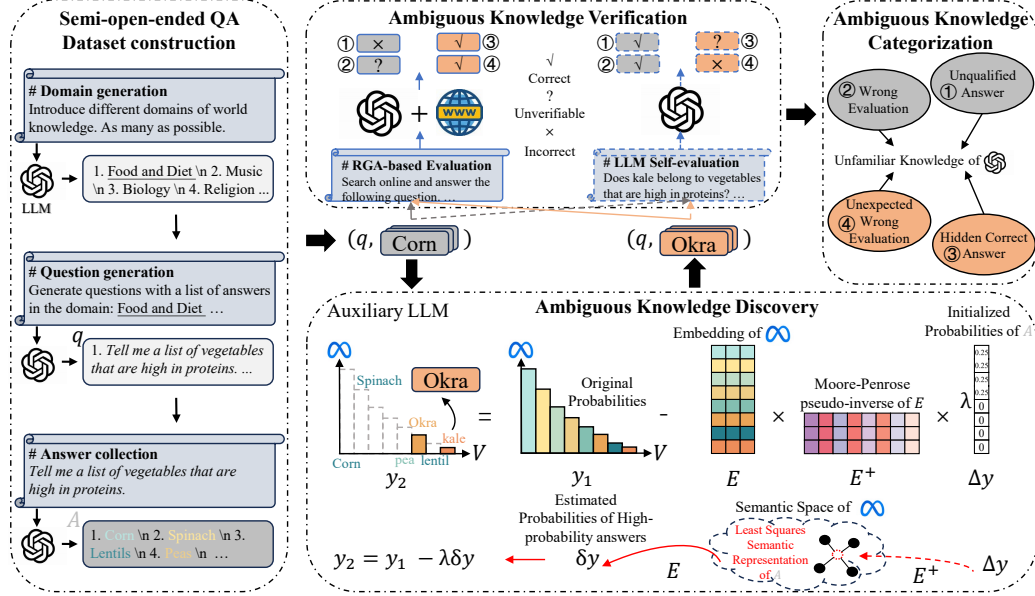
Figure 1: The overview of our framework. For the question q in the constructed dataset, The open-sourced auxiliary model prevent the high-probability answers based on existing answer entities A from the black-box LLM (i.e. GPT-4) and generate 4 categories of ambiguous knowledge that are unfamiliar knowledge for the target model.

open-sourced auxiliary LLM to incur more low-probability ambiguous answers for the perception of the knowledge boundary for the black-box LLM.

We denote that the open-sourced auxiliary model's final layer is $E \in \mathbb{R}^{|V| \times d}$, where d and V is the dimension size and V is the output vocabulary. In each generation step, LLM encodes the semantic representation of the input context as $x \in \mathbb{R}^{d \times 1}$ and uses it to calculate the next token generation probability as $y_1 = Ex$. Specifically, we take the following 3 steps to estimate and decrease the probability of high-probability answers.

1. Probability initialization for existing answers. We choose the first token in an answer entity $a \in A$ as "anchor token", which is suppose to represent the primary information about a since the first token is indispensable to its generation[9]. For a given question q, we define a vector $\Delta y \in \mathbb{R}^{V \times 1}$, indicating the existence of unique anchor tokens in all answer entities to the question. In Δy, we assign the value $\frac{1}{n}$ to the anchor tokens' position and assign 0 to other positions, where n is the number of all anchor tokens. As $\sum_i \Delta y_i = 1$, we deem Δy as the initialized probability distribution of all existing answers for the given question.

2. Semantic estimation of high-probability answers. We estimate the semantic representation δx for high-probability answer entities from initialized existing answer probability Δy. We calculate $\delta x = E^+ \Delta y$, where $E^+ E = I$. Here, E^+ is the left Moore-Penrose pseudo-inverse, and I is the identity matrix. Because the pseudo-inverse of a non-square matrix is often used to find the least squares solution of equations, In this way, we obtain δx as the least squares semantic representation that is the nearest to approximate the real semantic representation (which is unknown) of anchor tokens.

3. Probability Reduction. We calculate the probability of high-probability answers as $\delta y = E\delta x$ and ultimately obtain an adjusted generation probability y_2:

$$y_2 = y_1 - \lambda \delta y = y_1 - \lambda E \delta x = y_1 - \lambda E E_l^+ \Delta y,$$

where λ is a scaler that controls the extension of the reduction. Since anchor tokens and their semantic tokens almost share the same semantics[3], interpreting the estimated semantic

---

[3]Near-duplicate tokens roughly share the same semantic meaning but are different under tokenization due to typos, capitalization, or whitespace marking. For example, *Pea* and *peas* are two near-duplicate answers.

representation of anchor tokens propagates the semantic information to their semantic-related ones. Thereby, we obtain an estimated probability distribution of high-probability tokens δy and use it to prevent the generation of high-probability answers.

During generation, the auxiliary LLM samples words on the adjusted probability distribution y_2 instead of the original distribution y_1. In this way, we only reduce the probability of answers that are semantically related to existing high-probability answers. while preserving the remaining probability distribution almost intact. The black-box model has already exhausted almost all high-probability common answers in A, forcing the auxiliary model to generate new ambiguous answers with low probabilities. These ambiguous answers are either low-probability correct answers neglected by the black-box model or delusive incorrect answers.

3.4 Ambiguous Knowledge Verification

We compare the results of LLM self-evaluation against the ground truth from RAG-based evaluation to verify the truthfulness of ambiguous answers, thereby identifying answers beyond the knowledge boundary for the target LLM. We conduct self-evaluation on the target LLM with well-designed instructions. We craft a prompt template including 1) an incentive statement that encourages better performance: *I'll pay you $100 for a factually correct answer* [7]; 2) an instruction *think step by step* which prompts the LLM to analyze before reaching to a conclusion [32, 21]. 3) multiple human-crafted examples as in-context learning demonstrations.

We believe an answer a in each test case $(q, \mathcal{A})$ to be factually correct to q if verified on public, trustworthy sources. As questions in our dataset correspond to numerous low-probability hard answers, it is cost-prohibitive to annotate the truthfulness of each answer with expert knowledge. Inspired by Web-GPT [23], we adopt a cost-efficient approach to mimic human behavior when faced with unfamiliar knowledge. We instruct an Internet-connected LLM with the same prompt for self-evaluation and require it to search online for related information before making judgments.

For both evaluations, we evaluate each answer a to its corresponding question as 1) *incorrect*, which contradicts reliable sources; 2) *correct*, which is supported by reliable sources, and 3) *unverifiable*, which is for cases that cannot be verified based on available information. Finally, for each answer a, we compare the differences between LLM self-evaluation and RAG-based evaluation to categorize different types of ambiguous knowledge.

3.5 Ambiguous Knowledge Categorization

We classify the ambiguous answers into four categories based on the above two evaluation results and posit that they are beyond the knowledge boundary of the target LLM. For a question, we categorize the following types of answers to be the LLM's ambiguous answers:

- Unqualified answers: answers from the target LLM's response that are identified as incorrect or unverifiable according to the ground truth.
- Inaccurate evaluations: answers from the target LLM's response whose self-evaluation results contradict their ground truth.
- Hidden correct answers: answers that are neglected by the target LLM, yet supplemented by the auxiliary model, which are correct according to ground truth.
- Unexpected wrong evaluations: answers that are neglected by the target LLM but generated by the auxiliary model whose self-evaluation results misalign with their ground truth.

The above categorization helps us understand different types of misunderstanding of the target LLM regarding specific pieces of unfamiliar knowledge.

4 Experiments

4.1 Settings

In our experiment, we investigate the knowledge boundary of GPT-4 on our constructed dataset. We use two LLaMA-2-13b [30], as our auxiliary models in Sec. 3.3. Our method sets λ in Sec. 3.3 to 80. See more implementation details in App. A.

We compare our method with several baselines. Following prompt-based approaches [31],*Prompt* instructs the auxiliary model to generate more answers via prompt-engineering. Inspired by Zhang et al. (2023), *MASK* belongs to the representation-based perception that uses an average initialization for the probability of tokens from existing answers to represent the likelihood of high-probability answers and reduce their generation probabilities in the auxiliary model.

For the evaluation metrics, we use widely adopted Exact Match [27] (EM) and F1 scores [27, 31] to measure the performance of in discovering ambiguous answers. Different from previous research, our semi-open-ended questions correspond to a large number of correct answers, making it hard to build a comprehensive answer set for evaluation[4]. Instead, we select ambiguous answers identified by the RAG-evaluation and confirmed by our human annotators as their ground truth and compare them with the full response to calculate the **EM** and **F1**. In this way, EM is an entity-level metric that measures the percentage of ambiguous answers within the response. F1 is a word-level metric that quantifies the word overlap between the ambiguous answers and the ground truth. We adopt Bleu [25] to measure word-level overlap between responses from the GPT-4 response and the auxiliary model. We also use answer overlap rate (AOR) to evaluate the efficiency of generating distinctive ambiguous answers. AOR is an entity-level metric that calculates the proportion of words in a list of answer entities that duplicate the reference response. See more evaluation details in App. A.

4.2   Overall Performance

Table 1: The performance of different auxiliary models with various strategies in discovering more ambiguous answers.

Method	Auxiliary Model Size	EM ↑	F1 ↑	AOR ↓	Bleu1↓	Bleu2↓	Bleu3↓	Bleu4↓
Prompt	LLaMA-2-13B	0.300	0.461	0.490	0.252	0.118	0.052	0.023
MASK	LLaMA-2-7B	0.458	0.570	0.342	0.185	0.075	0.021	0.004
	LLaMA-2-13B	0.470	0.587	0.344	0.189	0.075	0.021	0.004
Ours	LLaMA-2-13B	**0.481**	**0.587**	**0.326**	**0.181**	**0.071**	**0.018**	**0.004**

We analyze the effectiveness of our auxiliary model in exploring the knowledge boundary of GPT-4 by comparing it with multiple baselines. We randomly sample 200 questions from our dataset and discover ambiguous answers with an auxiliary model. Then, we verify the truthfulness of these answers using RAG-based evaluation with human annotation following Sec 3.4 and measure their performance with our evaluation metrics. Tab 1 shows the results of ambiguous answers on different evaluation metrics when using different auxiliary models and strategies. *Prompt* directly prompts the auxiliary model to generate answers, achieving the worst performance on all metrics. This suggests that directly prompting the auxiliary model may generate many repetitive answers (results in high AOR and Bleu scores) and, therefore inefficient in discovering new ambiguous answers (results in low EM and F1). *MASK* reduces the generation probabilities of anchor tokens during generation. When employing *MASK* on the same auxiliary model, its EM and F1 increase to 0.47 and 0.587 respectively while achieving a lower AOR (0.344). It indicates that reducing the anchor words' generation probability is effective in achieving a more diverse generation. Replacing the auxiliary model with LLaMA-7B results in a slightly lower EM of 0.458. This marginal decrease implies that while a larger model can offer a broader knowledge base, the strategy of reducing anchor word probabilities is more influential in generating distinctive answers. Our strategy estimates and reduces the generation probability of near-duplicate answers. This approach achieves the best performance on all metrics. It underscores the effectiveness of our strategy in generating a diverse set of ambiguous answers that are less likely to duplicate existing ones, thus exploring the knowledge boundary for the target LLM. See our case study in Sec.4.6.

Table 2: Ablation study on the key components in our method. We use the same metrics as in Sec. 4.2, apart from those that require manual annotation. AOR and Bleu measure the entity- and word-level overlap respectively between answers from GPT-4 and different model variants. − Auxiliary Model prompt GPT-4 with existing answers as examples for more ambiguous answers. − Inverse Matrix keeps the near-duplicate tokens during generation. Besides, we adjust the probability influence scaler (λ) to verify its impact on the generation results.

Variants	AOR↓	Bleu1↓	Bleu2↓	Bleu3↓	Bleu4↓
− *Auxiliary Model*	0.535	0.267	0.106	0.037	0.010
− *Inverse Matrix*	0.344	0.189	0.075	0.021	0.004
Ours(λ=60)	0.419	0.224	0.098	0.038	0.013
Ours(λ=70)	0.352	0.192	0.076	0.021	0.005
Ours	**0.326**	**0.181**	**0.071**	**0.018**	**0.004**

4.3 Ablation Study

We conduct an ablation study on our proposed method to verify the importance of each component in eliciting more distinctive ambiguous answers (as shown in Tab 2). − *Auxiliary Model* abandons the auxiliary model and use existing answers as in-context learning examples to prompt GPT-4 for more answers. It achieves the highest on all metrics, indicating that prompting the black-box model violently for more answers is inefficient as it results in many repetitive answers. We also try to encourage a more diverse generation by increasing the generation temperature of GPT-4. However, we find that GPT-4 starts generating scrambled texts after just a few words and the perplexity of these texts exceeds 1000, while normally GPT-4 generates a low perplexity of around 10. This indicates that increasing the sampling temperature results in the generation of scrambled texts. Although adjusting the generation temperature can change the generation probabilities, it does not alter the original probability relationships, nor can it specifically target tokens related to ambiguous answers. − *Inverse Matrix* only reduces the probability of existing answers without considering their near-duplicate answers. It performs better than − *Auxiliary Model* while underperforms Ours on all metrics. It shows that estimating and reducing the probability of near-duplicate tokens augment the auxiliary model for higher generation diversity. We increase the intervention on the generation probability by lowering λ, the probability influence scaler. From row 3 to row 5 in Tab. 2, λ increases from 60 to 80, resulting in a decrease on all metrics. This suggests that the extent to which we intervene in the generation probability is positively correlated with the diversity of ambiguous answers produced by the auxiliary model.

4.4 Results of Perceiving the Knowledge Boundary for GPT-4

Table 3: Percentages of different categories of answers comparing the LLM self-evaluation and ground truth labels. Following the categorization in Sec. 3.5, we calculate that the percentage of unqualified answers is 40.15% by adding up the underlined results that are incorrect or unverifiable according to the ground truth. We also obtain the percentage of inaccurate evaluations as 28.47%, by adding up the results highlighted in red where self-evaluation is inconsistent with the ground truth.

Ground Truth\ Self-evaluation	Incorrect	Correct	Unverifiable
Incorrect	<u>20.98</u>	<u>8.37</u>	<u>2.25</u>
Correct	9.30	47.77	2.78
Unverifiable	<u>3.30</u>	<u>3.73</u>	<u>1.52</u>

By analyzing the ambiguous answers in our dataset, we arrive at the following findings: (1) **GPT-4 performs poorly on the semi-open-ended questions and generates many unqualified answers.** We calculate the percentage of questions where at least one of the GPT-4's answers is unverifiable

---

[4]There are many low-frequency answers for these questions. For example, there are maybe hundreds of accurate answers to the question:*What animals are native to Australia?*

or incorrect according to the ground truth. We find that GPT-4 generates incorrect or unverifiable answers in 82.90% of questions. By adding up the proportion of incorrect answers (row 1 in Tab.3) and unverifiable answers (row 3 in Tab.3), we identify that 40.15% of ambiguous answers belong to unqualified answers. (2) **GPT-4 makes many inaccurate evaluations regarding the truthfulness of ambiguous answers, indicating that the LLM lacks understanding of the relevant knowledge.** We identify answers whose ground truth misaligns with self-evaluations in Tab.3 and find that 28.47% (by adding up the results in red color from Tab.3) of ambiguous answers belong to inaccurate evaluations for GPT-4. It indicates that GPT-4 is unfamiliar with these pieces of knowledge, and retrieval is helpful for LLM to draw the correct conclusion. (3) **GPT-4 has limited ability to recognize its knowledge boundary, while in most cases it continues to produce unqualified answers.** We search for keywords in all responses that reflect GPT-4's admission of its knowledge boundary (e.g. *I apologize* and *I'm afraid*) and calculate the proportion of corresponding questions in the dataset. We find that in about 7% of questions, GPT-4 admits that it has listed all the answers and refuses to provide more answers (it generates a response like *I apologize for any confusion, but to the best of my knowledge, the list I provided includes all the correct answers.*). However, it fails to recognize its knowledge boundary in the rest questions and continues to generate unqualified answers. Our findings indicate that advanced LLM (i.e. GPT-4) is easy to hallucinate on semi-open-ended questions, indicating the importance of detecting the LLM knowledge boundary via these questions.

4.5 Results of the Auxiliary Model on Perceiving the Knowledge Boundary for GPT4

Table 4: Percentages of different categories of ambiguous answers comparing the GPT-4 self-evaluation results and their ground truth. Following the categorization in Sec.3.5, we calculate that the percentage of hidden correct answers is 75.12% by adding up the starred (*) results that are correct according to the ground truth. We also obtain the percentage of unexpected wrong evaluations as 62.43%, by adding up the results highlighted in orange where GPT-4-evaluation is inconsistent with the ground truth.

Ground Truth\ GPT-4-evaluation	Incorrect	Correct	Unverifiable
Incorrect	0.04	9.53	11.31
Correct	23.97*	37.54*	13.61*
Unverifiable	3.18	0.83	0.00

Tab.4 shows the fine-grained results of GPT-4 self-evaluation and the ground truth of ambiguous answers discovered by our auxiliary model. By analyzing the results, we conclude some interesting findings: (1) **LLaMA-2-13B effectively supplements GPT-4 by identifying hidden correct answers.** We add up the starred results in Tab.4 and obtain the proportion of LLM neglected hidden correct answers (75.12%). Notably, 23.97% and 13.61% of correct ambiguous answers are both neglected by GPT-4 and deemed to be incorrect or unverifiable under GPT-4 self-evaluation, showcasing the GPT-4's unfamiliarity with the corresponding knowledge. (2) **LLaMA-2-13B is easy to incur unexpected wrong evaluations during GPT-4 self-evaluation.** We add up the results highlighted in orange from Tab.4 and find that 62.43% of the GPT-4 self-evaluations are inconsistent with the ground truth. It implies that GPT-4's self-evaluation mechanism may not be fully aligned with the actual correctness of the ambiguous answers, especially when they are supplemented by the auxiliary model. (3) **LLaMA-2-13B is also able to discover situations where GPT-4 admits for its knowledge boundary.** We add up the results in column 3 where GPT-4 admits that it cannot make a judgment (unanswerable) during self-evaluation and obtain 24.92% of aligned ambiguous answers. It means that GPT-4 aligns well with these ambiguous answers because it neglected these answers during generation and deems them as unknown knowledge during self-evaluation.

4.6 Case Study

We showcase the importance of ambiguous answers in perceiving the knowledge boundary for LLMs. Tab. 8 shows an example with two ambiguous answers for the same question. First, we sample a semi-open-ended question with its GPT-4 responses and ambiguous answers augmented by the auxiliary model (row 2 and 3 in table Tab.8). Then, we manually construct two related questions,

Table 5: Examples of two ambiguous answers with their related questions. The Raspberry is an unqualified answer generated by GPT-4, which yields a different answer in another question. Plantain is an answer that is neglected by GPT-4, yet supplemented by the auxiliary model, which are correct according to ground truth. However, GPT-4 believes it to be wrong and generates wrong information for another question. Texts in yellowgreen are truthful information, while texts in red are non-factual.

Semi-open-ended Question	Tell me a list of foods that are rich in Vitamin A but low in fat.	
GPT-4 Response for Semi-open-ended Question	1. Carrots \n 2. Spinach ... 46. Raspberries \n 47. Red Leaf Lettuce ...	
Auxiliary Model Response	1. Bell peppers \n 2. Liver ... 14. Meat such as beef liver \n 15. Plantains ...	
Ambiguous Answer	Raspberries	Plantains
Answer Type	Unqualified Answer	Hidden Correct Answer
Related Question	What vitamins are rich in raspberries?	Is it a good choice to eat Plantains for many Vitamin A?
GPT-4 Response	Raspberries are rich in vitamins C, K, E, and B-complex. They also contain small amounts of vitamin A.	Plantains do contain vitamin A, but not in very high amounts. ...

each involving different types of ambiguous answers, and request new responses from GPT-4. For the case displayed on the left side of Tab.8, GPT-4 falsely deems the raspberry as food that is rich in vitamin A but low in fat, yet it answers correctly in another related question. It indicates that GPT-4 is inconsistent in answering different questions involving the same ambiguous answer. For the case on the right side of Tab.8, given the same semi-open-ended question, the auxiliary model discovers another ambiguous answer, Plantain, which is correct to the question. Interestingly, GPT-4 generates misinformation regarding this answer entity. It shows that GPT-4 falsely believes plantain is incorrect. It also indicates that the auxiliary model helps discover ambiguous answers that elicit misinformation in the target LLM. It strengthens the importance of perceiving knowledge boundaries for LLMs by discovering ambiguous answers to semi-open-ended questions. See App. D for more examples.

5 Conclusion

We investigate the perception of knowledge boundary for LLMs with semi-open-ended questions, an important yet underexplored type of question corresponding to a large number of accurate answers. We introduce an LLM-based approach to construct semi-open-ended questions and collect LLM answers from the target LLM. Then, we discover more pieces of unfamiliar knowledge for the target LLM by eliciting ambiguous answers from an auxiliary model that the LLM neglects. To achieve a more effective generation, we estimate and reduce the generation probability of existing answers with their near-duplicate counterparts. With our methods, we construct a dataset to evaluate the performance of GPT-4 and discover many ambiguous answers with our auxiliary model, LLaMA-2-13B. Our findings reveal that GPT-4 produces many unqualified answers and suffers from inaccurate evaluations. Besides, we verify that LLaMA-2-13B is effective in discovering more unpopular correct answers and delusive wrong answers neglected by GPT-4. Our findings underscore the importance of semi-ended questions and the effectiveness of our method in assisting in perceiving knowledge boundaries for LLMs.

References

[1] Garima Agrawal, Tharindu Kumarage, Zeyad Alghamdi, and Huan Liu. Can knowledge graphs reduce hallucinations in llms? : A survey, 2024.

[2] Alfonso Amayuelas, Liangming Pan, Wenhu Chen, and William Wang. Knowledge of knowledge: Exploring known-unknowns uncertainty with large language models. *arXiv preprint arXiv:2305.13712*, 2023.

[3] Yejin Bang, Nayeon Lee, Tiezheng Yu, Leila Khalatbari, Yan Xu, Samuel Cahyawijaya, Dan Su, Bryan Wilie, Romain Barraud, Elham J. Barezi, Andrea Madotto, Hayden Kee, and Pascale Fung. Towards answering open-ended ethical quandary questions, 2023.

[4] Man Luo;Shailaja Keyur Sampat;Riley Tallman;Yankai Zeng;Manuha Vancha;Akarshan Sajja;Chitta Baral. 'just because you are right, doesn't mean i am wrong': Overcoming a bottleneck in the development and evaluation of open-ended visual question answering (vqa) tasks. In *Conference of the European Chapter of the Association for Computational Linguistics*, 2021.

[5] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michał Podstawski, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefler. Graph of Thoughts: Solving Elaborate Problems with Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17682–17690, Mar 2024. doi: 10.1609/aaai.v38i16.29720. URL https://ojs.aaai.org/index.php/AAAI/article/view/29720.

[6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf.

[7] Sondos Mahmoud Bsharat, Aidar Myrzakhan, and Zhiqiang Shen. Principled instructions are all you need for questioning llama-1/2, gpt-3.5/4, 2024.

[8] Lang Cao. Learn to refuse: Making large language models more controllable and reliable through knowledge scope limitation and refusal mechanism. *arXiv preprint arXiv:2311.01041*, 2023.

[9] Haw-Shiuan Chang, Ruei-Yao Sun, Kathryn Ricci, and Andrew McCallum. Multi-CLS BERT: An efficient alternative to traditional ensembling. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 821–854, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.48. URL https://aclanthology.org/2023.acl-long.48.

[10] Jiefeng Chen, Jinsung Yoon, Sayna Ebrahimi, Sercan Arik, Tomas Pfister, and Somesh Jha. Adaptation with self-evaluation to improve selective prediction in LLMs. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5190–5213, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.345. URL https://aclanthology.org/2023.findings-emnlp.345.

[11] Zhuo Chen;Jiaoyan Chen;Yuxia Geng;Jeff Z. Pan;Zonggang Yuan & Huajun Chen. Zero-shot visual question answering using knowledge graph. In *The Semantic Web – ISWC 2021*, 2021.

[12] 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 Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben 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 Garcia, 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 Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022.

[13] Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng, and Tat-Seng Chua. Gotcha! don't trick me with unanswerable questions! self-aligning large language models for responding to unknown questions. *arXiv preprint arXiv:2402.15062*, 2024.

[14] Mor Geva, Yoav Goldberg, and Jonathan Berant. Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets. In Kentaro Inui,

Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1161–1166, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1107. URL `https://aclanthology.org/D19-1107`.

[15] Xuming Hu, Junzhe Chen, Xiaochuan Li, Yufei Guo, Lijie Wen, Philip S. Yu, and Zhijiang Guo. Towards understanding factual knowledge of large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL `https://openreview.net/forum?id=9OevMUdods`.

[16] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ArXiv*, abs/2311.05232, 2023. URL `https://api.semanticscholar.org/CorpusID:265067168`.

[17] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions, 2023.

[18] Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know, 2022.

[19] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu, editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.550. URL `https://aclanthology.org/2020.emnlp-main.550`.

[20] Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=HklBjCEKvH`.

[21] Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc., 2022. URL `https://proceedings.neurips.cc/paper_files/paper/2022/file/8bb0d291acd4acf06ef112099c16f326-Paper-Conference.pdf`.

[22] Nick Mecklenburg, Yiyou Lin, Xiaoxiao Li, Daniel Holstein, Leonardo Nunes, Sara Malvar, Bruno Silva, Ranveer Chandra, Vijay Aski, Pavan Kumar Reddy Yannam, Tolga Aktas, and Todd Hendry. Injecting new knowledge into large language models via supervised fine-tuning, 2024.

[23] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. Webgpt: Browser-assisted question-answering with human feedback, 2022.

[24] Shiyu Ni, Keping Bi, Jiafeng Guo, and Xueqi Cheng. When do llms need retrieval augmentation? mitigating llms' overconfidence helps retrieval augmentation. *arXiv preprint arXiv:2402.11457*, 2024.

[25] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In Pierre Isabelle, Eugene Charniak, and Dekang Lin, editors, *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics. doi: 10.3115/1073083.1073135. URL https://aclanthology.org/P02-1040.

[26] Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. Explain yourself! leveraging language models for commonsense reasoning. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4932–4942, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1487. URL https://aclanthology.org/P19-1487.

[27] Ruiyang Ren, Yuhao Wang, Yingqi Qu, Wayne Xin Zhao, Jing Liu, Hao Tian, Hua Wu, Ji-Rong Wen, and Haifeng Wang. Investigating the factual knowledge boundary of large language models with retrieval augmentation, 2023.

[28] Joshua Robinson, Christopher Michael Rytting, and David Wingate. Leveraging large language models for multiple choice question answering, 2023.

[29] Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber, and Lijuan Wang. Prompting GPT-3 to be reliable. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=98p5x51L5af.

[30] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.

[31] Yile Wang, Peng Li, Maosong Sun, and Yang Liu. Self-knowledge guided retrieval augmentation for large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10303–10315, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.691. URL https://aclanthology.org/2023.findings-emnlp.691.

[32] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf.

[33] 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.

[34] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9777–9786, June 2021.

[35] Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts, 2024.

[36] Hongshen Xu, Zichen Zhu, Situo Zhang, Da Ma, Shuai Fan, Lu Chen, and Kai Yu. Rejection improves reliability: Training llms to refuse unknown questions using rl from knowledge feedback, 2024.

[37] Yuqing Yang, Ethan Chern, Xipeng Qiu, Graham Neubig, and Pengfei Liu. Alignment for honesty. *arXiv preprint arXiv:2312.07000*, 2023.

[38] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of Thoughts: Deliberate problem solving with large language models, 2023.

[39] Xunjian Yin, Xu Zhang, Jie Ruan, and Xiaojun Wan. Benchmarking knowledge boundary for large language model: A different perspective on model evaluation. *arXiv preprint arXiv:2402.11493*, 2024.

[40] Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. Do large language models know what they don't know? In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8653–8665, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.551. URL `https://aclanthology.org/2023.findings-acl.551`.

[41] Xiang Yue, Boshi Wang, Ziru Chen, Kai Zhang, Yu Su, and Huan Sun. Automatic evaluation of attribution by large language models. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4615–4635, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.307. URL `https://aclanthology.org/2023.findings-emnlp.307`.

[42] Yue Zhang, Leyang Cui, Wei Bi, and Shuming Shi. Alleviating hallucinations of large language models through induced hallucinations. *arXiv preprint arXiv:2312.15710*, 2023.

[43] Yukun Zhao, Lingyong Yan, Weiwei Sun, Guoliang Xing, Chong Meng, Shuaiqiang Wang, Zhicong Cheng, Zhaochun Ren, and Dawei Yin. Knowing what llms do not know: A simple yet effective self-detection method. *arXiv preprint arXiv:2310.17918*, 2023.

A Implimentation Details

We query GPT-4-Turbo through API calls in our experiments. For our RAG-based evaluation in Sec. 3.4, we first employ Microsoft Copilot to search the Internet to find evidence and draw a conclusion. For our evaluation to calculate EM and F1, we hire 11 human annotators with Master's degrees to manually review the responses of Copilot, ensure the credibility of the references, and determine the truthfulness of each answer.

During the generation of our auxiliary model, We apply nucleus sampling (with p=0.9) during generation, setting the generation temperature to 0.7, and the repetition penalty to 1.15.

For our metrics, EM and F1 are common in QA tasks and knowledge boundary detection tasks. The Bleu metric is often used to measure the n-gram overlap between the model response and the ground truth. A lower Bleu score means a more dissimilar response generated by the auxiliary model. Before calculating the above metrics, we normalize answer entities with the NLTK library by lowercase the answers and turning them into the singular form.

B Dataset Information

Table 6: Example of data samples in our dataset. Each question corresponds to many ambiguous answers. Different colors of the answers represent different ground truth truthfulness labels: yellowgreen represents that the answer is verified as being factually correct, red is verified as being incorrect, and salmon is for unverifiable answers.

Domain	Percentage	Question	Ambiguous Answer
Biology	4.93%	Tell me a list of trees that produce fruit with hard shells and are native to tropical regions.	Malabar Chestnut, Miracle Fruit Tree, Guapinol Tree, Peanut Tree, Sapodilla Tree, Desert Date, Castor Oil Plant, Chiclé Tree, Oriental Persimmon, Dika Tree, Langsat Tree, Pequi Tree, Karite Tree, Imbu Tree, Caraipa Tree
Music	4.20%	Tell me a list of films whose plot is centered on a historically significant event, but the narrative is from the perspective of an imagined protagonist.	The Hurt Locker, Zulu, Letters from Iwo Jima, Doctor Zhivago, Black Hawk Down, Les Miserables, Lawrence of Arabia, The Wind that Shakes the Barley, The Longest Day, Anthony Adverse, The Guns of Navarone, Lion of the Desert
Geology	4.09%	Tell me a list of current deserts across the world that were previously covered by an ancient sea.	Qaidam Basin, Dasht-e Kavir, Dasht-e Lut, The Registan Desert, The Franklin Basin, The Makran Desert, The Sonoran Desert, The Jornada Del Muerto, Great Karoo, Tanami Desert
Food and Diet	5.67%	Tell me a list of fruits that are sources of healthy fats, not including avocados and coconuts.	Sacha Inchi Nuts, Hemp Seeds, Chokeberries, Elderberries, Pine Nuts, Passion Fruit Seeds, Tibetan Goji Berries
Literature	1.68%	Tell me a list of female writers from the Victorian era whose work focuses on social reform.	Helen Taylor, Catherine Helen Spence, Octavia Hill, Rhoda Broughton, Elizabeth Missing Sewell, Emmeline Pankhurst, Alice Meynell, Elizabeth Robins, Mary Augusta Ward, Constance Garnett

Our dataset covers the following topics: Environment and Climate, Technology and Industry, Political Science, History and Archaeology, Sociology, Economy and Finance, Philosophy, Languages, Art, Architecture, Music, Physics, Astronomy, Chemistry, Biology, Geology, Computer Science, Anthropology and Cultures, Education, Psychology and Mental Health, Fitness and Physical Health, Literature, Religion, Law and Criminology, Military and War Agriculture, Tourism, Film and Television, Sports and Athletics, Food and Diet, Energy and renewable resources, Mathematics and Statistics, Medicine and Health, Games, Clothing and Fashion.

C Important Instructions

Here we provide important instructions in building our dataset, guiding LLM for self-evaluation, and show guidelines for human annotators in the supplementary materials.

- Domain Generation. *I hope to test my students' knowledge in different domains. Which domains can I use to create questions?*

- Question Generation. *I am a professor of [CATEGORY] and need to test students' under-standing of [CATEGORY] by asking a series of challenging questions. These questions require respondents to list entities that they know meet a series of certain conditions. You need to create more different and diverse challenging questions according to the require-ments. Read the following requirements carefully. I'm going to tip $100 for a perfect list of questions! \n The questions should meet the following criteria: \n 1. Each question should start with "Tell me a list of"; \n 2. To make the question challenging enough, each question should contain multiple limiting conditions. \n 3. The requirement of the question should not involve specific numbers (which makes the question too hard to answer) or vague descriptions (which makes it hard to evaluate the truthfulness of the answer), like "long lifespan", "quick speed", "popular", and "important"; \n 4. The boundaries of the question should be very clear, making it easy to evaluate its truthfulness; \n 5. The answers to the questions should be consistent through a relatively long time and not change frequently, for example, yearly. \n Refer to the style in the following two examples from an exemplary subject, biology. \n Question 1: Tell me a list of land animals unique to Australia. \n Question 2: Tell me a list of fruits that grow on trees in tropical regions.*

- Self-evaluation. *Does [AMBIGUOUS ANSWER] belong to [QUESTION REQUIREMT-NTS]? I'll tip $100 for the factually correct answer. Think step by step and then give your answer.*

- RAG-based evaluation. *Search online for highly credible information related to the following question, and answer the question based on the search results. \n Does [AMBIGUOUS ANSWER] belong to [QUESTION REQUIREMTNTS]? I'll tip $100 for the factually correct answer. Think step by step and then give your answer.*

D   Cases of Different Types of Ambiguous Answers with Misinformation on the Related Questions

We show different types of ambiguous answers and GPT-4's performance on their related questions from Tab.7 to Tab.8.

Table 7: Examples of two ambiguous answers with their related questions. The Malabo, Equatorial Guinea is an answer that neglected by GPT-4, yet supplemented by the auxiliary model, which are correct according to ground truth. However, GPT-4 believes it to be wrong and generates wrong information for another question. The Rome, Italy is an answer that neglected by GPT-4 but generated by the auxiliary model, whose self-evaluation results misalign with the ground truth. However, GPT-4 believes it to be wrong and generates wrong information for another question. Texts in yellowgreen are truthful information, while texts in red are non-factual.

Semi-open-ended Question	Tell me a list of world capitals where more than two languages are considered official.	
GPT-4 Response for Semi-open-ended Question	1. London, United Kingdom \n 2. Montreal, Canada ... 52. Zagreb, Croatia \n 53. Prague, Czech Republic ...	
Auxiliary Model Response	1. Brussels, Belgium \n 2. Ottawa, Canada ... 9. Rome, Italy \n 10. Malabo, Equatorial Guinea ...	
Ambiguous Answer	Malabo, Equatorial Guinea	Rome, Italy
Answer Type	Hidden Correct Answer	Unexpected Wrong Evaluations
Related Question	How many official languages are there in Malabo?	Is Latin the official language of Rome?
GPT-4 Response	There are two official languages in Malabo: Spanish and Pidgin English.	No, the official language of Rome is Italian. ...

Table 8: Examples of two ambiguous answers with their related questions. The Sand Island Light is an unqualified answer generated by GPT-4, which yields a different answer in another question. The Portsmouth Harbor Light is an answer from GPT-4, whose self-evaluation results contradict the ground truth. However, GPT-4 believes it to be true and generates wrong information for another question. Texts in yellowgreen are truthful information, while texts in red are non-factual.

Semi-open-ended Question	Tell me a list of famous lighthouses located on islands in the Atlantic Ocean.	
GPT-4 Response for Semi-open-ended Question	1. Pemaquid Point Light \n 2. Portland Head Light ... 46. Sand Island Light \n 47. Portsmouth Harbor Light \n ...	
Auxiliary Model Response	1. Maine Coast Guard Station and Pier Light \n 2. Minot's Leading Light ... 10. Woodenboat School Light \n 11. Chappaquiddick Light ...	
Ambiguous Answer	Portsmouth Harbor Light	Sand Island Light
Answer Type	Inaccurate Evaluations	Unqualified Answer
Related Question	Is Portsmouth Harbor Light located on which island in the Atlantic Ocean?	Where is Sand Island Light located?
GPT-4 Response	Portsmouth Harbor Light is located on New Castle Island in the Atlantic Ocean, in the state of New Hampshire, USA.	Sand Island Light is located in Alabama, USA, specifically on Sand Island off the coast of Dauphin Island. ...