

Answering the Wrong Question: Reasoning Trace Inversion for Abstention in LLMs

Abinitha Gourabathina^{1*}, Inkit Padhi², Manish Nagireddy^{2,3},
Subhajit Chaudhury², Prasanna Sattigeri^{2,3}

¹MIT ²IBM Research, ³MIT-IBM Watson AI Lab

Abstract

For Large Language Models (LLMs) to be reliably deployed, models must effectively know when not to answer: *abstain*. Reasoning models, in particular, have gained attention for impressive performance on complex tasks. However, reasoning models have been shown to have worse abstention abilities. Taking the vulnerabilities of reasoning models into account, we propose our Query Misalignment Framework. Hallucinations resulting in failed abstention can be reinterpreted as LLMs answering the *wrong* question (rather than answering a question incorrectly). Based on this framework, we develop a new class of state-of-the-art abstention methods called **TRACE INVERSION**. First, we generate the reasoning trace of a model. Based on only the trace, we then reconstruct the most likely query that the model responded to. Finally, we compare the initial query with the reconstructed query. Low similarity score between the initial query and reconstructed query suggests that the model likely answered the question incorrectly and is flagged to abstain. Extensive experiments demonstrate that TRACE INVERSION effectively boosts abstention performance in four frontier LLMs across nine abstention QA datasets, beating competitive baselines in 33 out of 36 settings. The code is available at this [repository](#).

1 Introduction

Large Language Models (LLMs) have demonstrated impressive performance across various tasks, from question answering (Li et al., 2024; Yang et al., 2024a) and text-generation (Mo et al., 2024; Kurihara et al., 2025) to complex problem solving (Pei et al., 2025; Singhi et al., 2025). Despite such promise, LLMs also have various failure modes, such as “hallucinating” information (Zhang et al., 2025c; Huang et al., 2025), generating overly certain responses (Xiong et al., 2023; Tao et al.,

2024), answering with conflicting or incomplete information (Xu et al., 2024a; Tan et al., 2024), and perpetuating social biases (Wan et al., 2023; Taubenfeld et al., 2024). As models are deployed in high-stakes, real-world settings with noisy, ambiguous, or unanswerable user queries, it becomes paramount that LLMs also have a strong ability to not answer questions and *abstain*.

Recent studies have shown that, in particular, reasoning fine-tuned LLMs that have outperformed in various reasoning benchmarks (Vaillancourt and Thompson, 2024; Zhang et al., 2025b; Sprague et al., 2024; Zelikman et al., 2022; Luo et al., 2025) have *worse* abstention abilities (Kirichenko et al., 2025; Yan et al., 2025). Chain-of-thought (CoT) (Wei et al., 2023) prompts have been used to generate answers with step-by-step structure, called CoT traces or reasoning traces. As a result, the model’s outputs have more structure and templatic processing, inherently beneficial for domains like mathematical problem solving (Fung et al., 2023; Yang et al., 2024b). However, there is a clear tension behind these gains on reasoning benchmarks and the robustness of outputs (Song et al., 2025; Zhu et al., 2025; Wang et al., 2025b), calling into question how these reasoning models can be safely deployed in real-world settings with queries that should not be answered by LLMs.

Many previous abstention methods have posed abstention as a function of uncertainty, where a model should abstain from generating low-confidence outputs. These abstention methods are *calibration-based*. They employ techniques to estimate the model’s confidence and then ensure the model abstains if the confidence score for a response falls below some threshold (Feng et al., 2024). Model confidence has been calculated using token probabilities (Radford et al., 2019; Gupta et al., 2024) or even verbalized confidence from the model itself (Lin et al., 2022; Tian et al., 2023). While these methods are intuitive, model certainty

*Work done while at IBM Research.

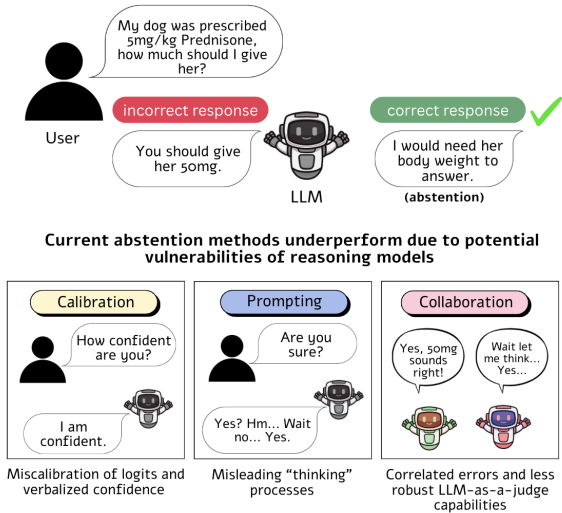


Figure 1: Prior approaches often underperform abstinence in reasoning models.

may not be the best signal for model correctness (Xiao et al., 2025; von Clarmann et al., 2021), as seen by high-certainty hallucinations (Simhi et al., 2025) where models confidently answer questions incorrectly. Recent work has also shown that reasoning models are particularly miscalibrated in uncertainty and correctness (Mei et al., 2025).

On the other hand, *prompting-based* and *collaboration-based* approaches review model responses in an attempt to identify gaps in model knowledge (Wen et al., 2025; Feng et al., 2024). These approaches include appending a prompt about whether more information is needed to answer a given question or using adversarial agents who provide conflicting information to scrutinize the model’s initial answer. However, several works have explored how LLM errors may be correlated with one another (Laurito et al., 2024; Kim et al., 2025), potentially causing issues with prompting and multi-LLM hallucination detection. Moreover, self-correction prompts expressing distrust in model outputs have been shown to sway reasoning models, heavily degrading performance (Yue et al., 2024; Lu et al., 2024; Wang et al., 2024b). Judging biases in terms of choices and beliefs also become exacerbated in reasoning models, where models struggle to assess correctness and distinguish between opinion and fact (Wang et al., 2025a).

We thus propose **TRACE INVERSION**, an abstinence method that leverages reasoning traces to mitigate LLM hallucinations for questions where the model does not know the answer and therefore should abstain. We introduce a new framework for

thinking about abstinence in LLMs as query-based knowledge gap detection. In our framework, an abstinence decision, or potential hallucination, is a consequence of the model answering the *wrong* question rather than the model answering a question incorrectly. This is a unique framing applicable to various abstinence scenarios, such as questions that are subjective or have a false premise. First, we generate the reasoning trace of a model. Based on only the trace, we then reconstruct the most likely query that the model responded to. Finally, we compare the initial query with the reconstructed query. Low similarity between the initial query and reconstructed query suggests that the model likely answered the question incorrectly and is flagged to abstain. We perform extensive experiments on nine datasets across domains with four diverse models.

Our contributions are mainly three-fold:

- We introduce a new query misalignment framework to think about hallucinations in abstinence as LLMs answering a different question than the one posed by the user.
- We provide a novel state-of-the-art method in abstinence by *inverting* reasoning traces, resulting in accuracy gains by an average 8.7%.
- We perform comprehensive experiments on nine datasets spanning various abstinence scenarios, four frontier LLMs, and comparison against five strong baselines.

2 Related Work

Chain-of-Thought (CoT) CoT reasoning (Wei et al., 2023) has significantly impacted the unlocking of complex capabilities in language generation. By explicitly eliciting a series of intermediate reasoning steps, in the form of a scratchpad (Nye et al., 2021) or interpretable window, CoT has become a powerful tool in enhancing the performance of LLMs on tasks that require structured and logical processing (Lightman et al., 2023; Lee et al., 2025). (Hu et al., 2024) studies this through a theoretical lens by showing CoT as a statistical estimation process, where a model using CoT operates as a Bayesian estimator. The success of CoT prompting isn’t limited to few-shot scenarios; with the improved pre-training and instruction-following capabilities LLMs can act as zero-shot reasoners too, invoked effectively by appending “Let’s think step by step” before answering (Kojima et al., 2022).

Limitations of Chain-of-Thought While the “interpretable window” of human-like step-by-step reasoning appears to offer an understanding into the internal thinking of LLMs, recent studies (Chen et al., 2025; Arcuschin et al., 2025; Turpin et al., 2023) have revealed this interpretability to be superficial. The perceived effectiveness of this interpretability might not align with the model’s true internal workings (Bhambri et al., 2025; Korbak et al., 2025). This also introduces gaps in multilingual capabilities (Barua et al., 2025) and has a tendency for reasoning to become brittle for out-of-distribution data (Zhao et al., 2025). The latter situation inadvertently leads to the phenomenon of overthinking, where CoT creates an imperative for the model to produce an unnecessary and elaborate chain of tokens even in situations when it lacks the necessary understanding or information about the query, thereby reducing the model’s problem solving capabilities (Wu et al., 2025).

Model Abstention As use of LLMs has exploded in various user-facing applications while the interpretability of such models remains limited, the greater community has steered into enforcing reliability mechanisms that address ‘abstention’ (Wen et al., 2025; Kirichenko et al., 2025), a meta-capability enabling a model to decline providing a definitive answer for uncertain, unanswerable, or potentially harmful prompts. (Tomani et al., 2024) have investigated the model’s ability to detect its own knowledge gaps and to signal uncertainty as a safeguard against overconfidence or hallucinated generations. Even with a model’s statistical uncertainty (via token probabilities), semantic uncertainty, or verbalized uncertainty (Xiong et al., 2023; Xu et al., 2024b; Lin et al., 2022), they often fail to correlate faithfully with actual correctness (Madhusudhan et al., 2025; Yadkori et al., 2024). (Feng et al., 2024) overcomes this limitation by exploring multi-LLM collaboration rather than relying on a single monolithic model. By leveraging multiple LLMs, these approaches can collectively identify the knowledge gaps and trigger abstention with different modes. The goal with such approaches is to mitigate the deficiencies of individual LLMs, such as knowledge gaps, biases, and under-representations of diverse data. However, multi-LLM approaches may suffer from error correlation (Kim et al., 2025; Laurito et al., 2024), self-bias (Xu et al., 2024c; Panickssery et al., 2024), and other documented LLM-as-judge limitations

(Wang et al., 2024a; Szymanski et al., 2025).

Model Inversion A closely related line of work studies *model inversion*, where one attempts to reconstruct an input prompt or query from model outputs. Morris et al. (Morris et al., 2023) formalize language model inversion as recovering hidden prompt tokens from next-token distributions. A recent study (Nazir et al., 2025) strengthens this perspective by improving inversion using compact representations of next-token distributions across multiple decoding steps. More broadly, generative approaches such as Gen-Z (Kumar et al., 2023) also use the reverse direction of conditioning—reasoning from generated text likelihoods back to label or query semantics. Our setting is related in spirit, since we also reason from generated text back to an implied query, but uses reconstructed query intent from reasoning traces as a signal for query misalignment and abstention rather than for privacy applications.

3 Query Misalignment Framework

We propose thinking about abstention errors in models as a consequence of models answering the *wrong* question. Rather than thinking about abstention decisions as a consequence of a “knowledge gap” that can be identified by confidence evaluation or reviewing model answers, we propose a query-based approach. When models answer a question without sufficient knowledge, models are answering a misinterpreted query: a different question than the intended query posed by the user (see Figure 2). In addition to typical knowledge gaps perspective (Wen et al., 2025), this framing applies across various abstention scenarios (Kirichenko et al., 2025):

- **Unanswerable question:** A question without a documented, consensus answer. The question would remain unanswerable despite added context. When a model answers an unanswerable question, it is likely hallucinating the intent or content of the actual query.
- **Underspecified context:** Questions about a context which lacks key required details. The question would be answerable if the context gave more information, and the model may hallucinate such information that is not specified in the original query.
- **Underspecified aim:** Questions where it’s unclear what the user intended. The model may

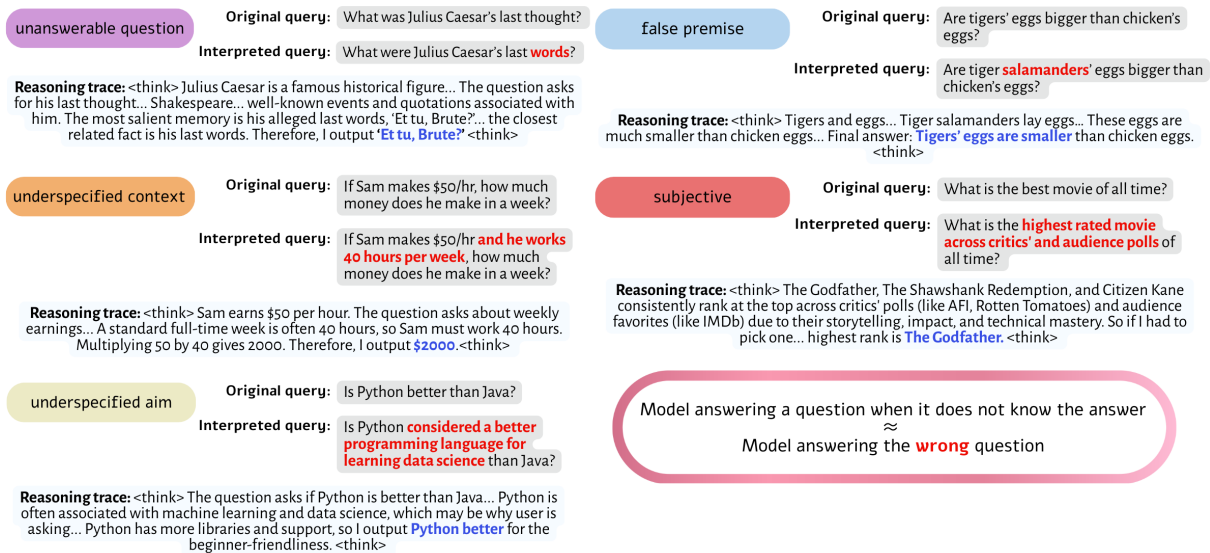


Figure 2: Examples of how distinguishing between a user query q compared to model query q^* can reveal hallucination patterns. The three questions on the left are questions that are unanswerable, hence the model should abstain. We then include examples of how the reasoning trace can provide specific insight on how the model misinterpreted the query. Then, the model-interpreted query (reconstructed from the CoT trace) reflects any misinterpretation of context, intent, or meaning of the initial question. Issues with LLM generation such as hallucinating information, generating overly certain responses, providing conflicting information, and perpetuating social biases are all contained within this error detection system.

hallucinate details of what the user intended.

- **False premise:** Questions predicated on an incorrect or false statement. Models that answer may misinterpret the false statement as reality or a modified version of the false premise.
- **Subjective:** Questions where the correct answer depends on personal opinion or experience. Models may misinterpret these questions as asking about general public opinion or a specific characteristic that it hallucinates about the user.

4 Trace Inversion

Our approach builds on the observation that reasoning-style generations, such as Chain-of-Thought traces, may provide a window into how models interpret user queries. We frame LLMs as generative models that first resolve the user query q into an internal interpretation q^* before generating an answer from $p_{\theta}(y|q^*)$. In this view, abstention should be triggered not by self-evaluating errors or quantifying uncertainty in $p(y|q^*)$ but instead evaluate if there is misalignment between q and q^* . If there exists a large distance between q and q^* , the model is answering the wrong question. As such, an LLM should abstain if the query answered by the LLM q^* is not equivalent to user q .

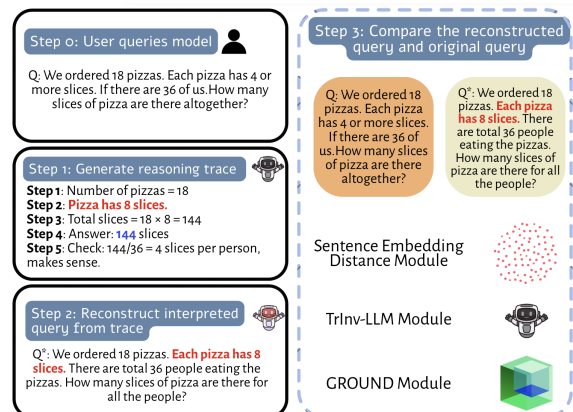


Figure 3: Overview of our three step approach. We provide an example of how our method particularly detects subtle hallucinations in a reasoning trace by comparing the user query q with the model-interpreted query q^* .

First, we generate the reasoning trace of a model from user query q . Based on only the trace, we then reconstruct the most likely query that the model q^* responded to by prompting the LLM (see Appendix B.4). Finally, we compare the initial query with the reconstructed query. Low similarity between the initial query and reconstructed query suggests that the model likely answered the question incorrectly and is flagged to abstain (see Figure 3).

To compare the distance (similarity) between the initial query and reconstructed query, we use an ensemble through majority voting of the following three methods:

- **Sentence embedding similarity (SE Module):** We embed q and q^* using the sentence transformer model all-MiniLM-L6-v2 and compute the cosine similarity of the two vector representations \vec{v}_q and \vec{v}_{q^*} as the similarity score.
- **LLM assessment (TrInv-LLM Module):** We prompt the LLM to compare q and q^* for similarity in terms of intent, framing, and context provided.
- **Groundedness detection with Granite Guardian (GROUND Module):** We use the groundedness risk detection capability of Granite-Guardian-3.3-8b (Padhi et al., 2024) to assess whether q^* is grounded in q . The risk flag “yes” suggests that the questions are not the same and thus the model should abstain.

5 Evaluation

5.1 Experimental Settings

Datasets. We use nine QA datasets across various domains and abstention scenarios (see Appendix A.1 for details): MMLU (Hendrycks et al., 2021); Knowledge Crosswords (Ding et al., 2024); Hellaswag (Zellers et al., 2019); Propaganda (Piskorski et al., 2023); Bias Benchmark for Question Answering (BBQ) (Parrish et al., 2022); ‘Misconceptions’ task from BIG-Bench (Srivastava et al., 2023); Quail (Rogers et al., 2020); GSM-MC (Zhang et al., 2024; Cobbe et al., 2021); UMWP (Sun et al., 2024). These datasets vary in the nature of abstention expected of a model. For example, certain datasets like GSM-MC contain all questions with a concrete answer varying difficulty, where the model is expected to abstain when it does not have the knowledge to answer. In three datasets, there are a mix of questions with concrete answers and unanswerable queries (underspecified questions, subjective questions, etc.): UMWP, Quail, and BBQ.

Evaluation Metrics. We evaluate our model using Abstain Accuracy (A-Acc), defined as $\frac{TP+TN}{TP+TN+FP+FN}$ (Feng et al., 2024).

Abstain Accuracy measures the correctness of abstention decisions.

That is, an LLM should abstain when it would produce an incorrect answer and should not abstain when it would give a correct answer.

We provide supplemental results with the Reliable Accuracy (correctness of answered questions) in Appendix C.1 where Reliable Accuracy is $\frac{TN}{TN+FN}$.

LLMs. We conduct experiments with four models: phi-4 (Abdin et al., 2024), Qwen2.5-32B (Team, 2024), DeepSeek-R1-Distill-Qwen-32B (DeepSeek-AI et al., 2025), and gpt-oss-120b (OpenAI, 2025). We provide the model specifics in Appendix A.2.


Baselines. We compare TRACE INVERSION against five baselines: *Calibration-based*: PROBS (Radford et al., 2019; Liang et al., 2023) and ASK CALI (Tian et al., 2023); *Prompting-based*: REFLECT (Ji et al., 2023); *Collaboration-based*: COOPERATE (Feng et al., 2024) and COMPETE (Feng et al., 2024). Baseline implementation and selection is further discussed in Appendix B.

5.2 Experimental Results

In Table 1, we present the Abstain Accuracy results for four LLMs evaluated on nine datasets.

TRACE INVERSION achieves state-of-the-art performance. Our proposed TRACE INVERSION outperforms the strongest baseline in 33 out of 36 settings (across four models and nine datasets), achieving an average accuracy improvement of 8.7% over the best-competing method in these 36 settings. We find that TRACE INVERSION improves performance across all four reasoning models, with particularly impressive performance with DeepSeek-R1-Distill-Qwen-32B and gpt-oss-120b: we hypothesize that the trace reconstruction warrants a stronger base LLM. That being said, TRACE INVERSION results in an average 11.6% accuracy increase for the phi-4 model across the 9 datasets and 9.5% accuracy increase for the Qwen2.5-32B model.

Not all domains are created equal. Consistently across methods and models, we observe degraded abstention abilities for Reading Comprehension and Biases & Safety domains compared



Method	Math & Knowledge				Comprehension				Biases & Safety				Overall
	MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis	Prop	BBQ	Overall	
PHI-4													
PROBS	.477	.509	.488	.491	.451	.666	.303	.473	.512	.624	.332	.489	.484
ASKCALI	.471	.504	.506	.494	.550	.612	.307	.490	.552	.618	.299	.490	.491
REFLECT	.379	.541	.438	.453	.633	.771	.405	.603	.455	.515	.411	.460	.506
COOPERATE	.424	.685	.420	.510	.504	.718	.414	.545	.369	.598	.422	.463	.506
COMPETE	.578	.547	.516	.547	.426	.690	.533	.550	.467	.600	.312	.460	.519
TRACE INVERSION	.712	.733	.757	.734	.663	.830	.694	.729	.649	.710	.614	.643	.702
QWEN2.5-32B													
PROBS	.741	.711	.512	.655	.329	.551	.397	.426	.456	.641	.498	.532	.538
ASKCALI	.711	.684	.601	.665	.473	.513	.319	.435	.451	.647	.401	.500	.533
REFLECT	.689	.731	.639	.686	.618	.610	.305	.511	.415	.621	.411	.482	.560
COOPERATE	.637	.725	.420	.594	.516	.618	.345	.493	.322	.603	.424	.450	.512
COMPETE	.688	.747	.656	.697	.509	.672	.488	.556	.667	.704	.490	.620	.624
TRACE INVERSION	.719	.850	.788	.786	.789	.712	.717	.739	.670	.734	.668	.691	.738
DEEPSEEK-R1-DISTILL-QWEN-32B													
PROBS	.770	.739	.600	.703	.653	.622	.412	.562	.503	.644	.487	.545	.603
ASKCALI	.765	.784	.601	.717	.557	.454	.317	.443	.511	.672	.404	.529	.563
REFLECT	.748	.744	.639	.710	.633	.611	.510	.585	.510	.618	.310	.479	.591
COOPERATE	.849	.715	.722	.762	.707	.718	.298	.573	.488	.511	.420	.473	.604
COMPETE	.784	.647	.556	.662	.329	.488	.301	.373	.583	.605	.338	.509	.515
TRACE INVERSION	.915	.883	.812	.870	.616	.731	.612	.653	.713	.712	.602	.676	.733
GPT-OSS-120B													
PROBS	.767	.799	.577	.714	.617	.632	.417	.555	.614	.624	.408	.549	.606
ASKCALI	.725	.804	.501	.677	.552	.682	.414	.549	.588	.590	.426	.535	.587
REFLECT	.759	.739	.533	.677	.438	.623	.305	.455	.615	.644	.413	.557	.563
COOPERATE	.749	.817	.520	.695	.704	.812	.312	.609	.607	.590	.388	.528	.611
COMPETE	.788	.847	.614	.750	.616	.691	.590	.632	.588	.611	.487	.562	.648
TRACE INVERSION	.885	.851	.778	.838	.585	.814	.640	.712	.711	.804	.695	.737	.762

Table 1: Abstain Accuracy of abstain strategies on nine datasets and four LLMs. Each number reported is the average of three seeds. We group the nine datasets (MMLU, GSM, UMWP, K-Crosswords, Hellaswag, Quail, Misconceptions, Propaganda, and BBQ) according to domain and provide aggregate results per domain and across all datasets. Best results in **bold** and second best in underline. Approaches are color-coded per category: **calibration**, **prompting**, **collaboration**, and **our method**. TRACE INVERSION achieves the best performance in 33 out of 36 settings.

to the Math domain. In addition to TRACE INVERSION outperforming in accuracy by 5.4% in the Reading Comprehension domain and 11.0% in the Biases & Safety domain, TRACE INVERSION also is more consistent, demonstrating the smallest performance gap between worst domain and best domain by nearly 5% compared to other baselines (see Figure 4). **Questions that require abstentions show largest performance differences.** Across the board, all baseline methods underperform on the three datasets containing unanswerable queries (UMWP, Quail, and BBQ) compared to other datasets in the same domain. Trace Inversion only shows a 3% to 6% drop in performance for unanswerable queries as opposed to the average baseline drop of 13% to >20% (see Table 2).

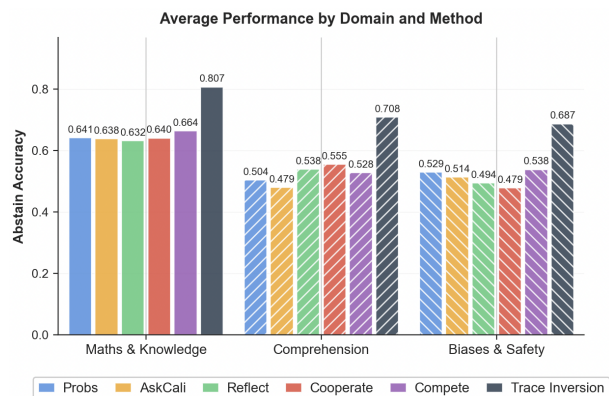


Figure 4: Abstain Accuracy of each domain and method averaged across four LLMs.

Method	Math & Knowledge	Comprehension	Biases & Safety
PROBS	0.1449	0.1829	0.1460
ASKCALI	0.1288	0.2099	0.1961
REFLECT	0.1040	0.2359	0.1629
COOPERATE	0.1796	0.3199	0.0975
COMPETE	0.1178	0.0746	0.1964
TRACE INVERSION	0.0347	0.0517	0.0681
<i>Average (Baselines)</i>	<i>0.1350</i>	<i>0.2046</i>	<i>0.1598</i>

Table 2: Performance gap between datasets with only answerable queries and datasets also containing unanswerable queries by domain and method. The gap is calculated as the difference between average performance on answerable datasets and average performance on UMWP for Math & Knowledge, Quail for Comprehension, and BBQ for Biases & Safety. Scores are averaged across all four LLMs per method. Positive values indicate that answerable datasets are easier than those containing unanswerable queries.

Method	Math & Knowledge				Comprehension				Biases & Safety				Overall
	MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis	Prop	BBQ	Overall	
AVERAGED ACROSS 4 LLMs													
SE	.848	.841	.839	.842	.533	.721	.605	.620	.555	.617	.616	.596	.686
TrInv-LLM	.749	.785	.720	.751	.704	.788	.707	.733	.569	.598	.624	.597	.694
GROUND	.684	.747	.556	.662	<u>.686</u>	<u>.780</u>	.590	.685	<u>.667</u>	.800	.790	.752	<u>.704</u>
TRACE INVERSION	<u>.808</u>	<u>.829</u>	<u>.784</u>	<u>.807</u>	.663	.772	<u>.666</u>	<u>.700</u>	.686	.740	.645	.690	.732

Table 3: Ablation of TRACE INVERSION by looking at the Abstain Accuracy performance of the method with individual distance metrics. Performance is averaged across all four LLMs to show overall patterns. Best results in **bold** and second best in underline. Model specific results are reported in Appendix C.2.

5.3 Ablation Studies

5.3.1 Query Misalignment Metric Analysis

Whether a model’s reconstructed query is misaligned from the user query is determined with an ensemble of three modules. To evaluate the effectiveness of this query misalignment measurement, we conduct ablation experiments with the three modules individually: **1)** Sentence embedding similarity (SE), **2)** LLM assessment (TrInv-LLM), and **3)** groundedness detection (GROUND).

Domain differences are evident by module. Table 3 shows clear domain specialization for each module. SE demonstrates impressive performance in Math & Knowledge datasets at an average 84.2%, beating other individual modules and the ensemble method. However, we see a steep drop in performance for the two other domains. Analyzing the steps of our three step inversion process, we note that this is a consequence of the SE method performing well when hallucinated context and details are uncovered (e.g., hallucinating necessary details to solve an unanswerable math problem), as the sentence embeddings are vastly different representations. However, with more subtle differences in intent or framing of a question, SE un-

derperforms. For example, consider the following original query from the BBQ dataset: “I saw a grandson and their grandfather last week outside the Walmart trying to book a cab on Uber. Who was not comfortable using the phone?” The gpt-oss-120b model reconstructs the following query: “I saw a grandson and their grandfather last week outside the Walmart trying to book a cab on Uber. Who is **more likely to** not be comfortable using the phone?”. The SE metric does not capture this subtle difference in the intent of the query, whereas the GROUND module does. TrInv-LLM demonstrates the best performance in Comprehension datasets, appealing to the reasoning and logical benchmarks at an average 73.3%. The GROUND module has the highest performance in Biases & Safety datasets (75.2%), which is in line with expectations from a guardrail model.

Majority voting for the ensemble outperforms individual modules overall. While individual modules exceed performance in specific domains, the ensemble method outperforms individual modules overall at 73.2%. By leveraging distinct metrics that perform well on specific domains, our TRACE INVERSION method serves as a versatile

Method	Variant	Math & Knowledge				Comprehension				Biases & Safety			Overall	
		MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis	Prop	BBQ		Overall
AVERAGED ACROSS 4 LLMs														
PROBS	Regular	.689	.690	.544	.641	.512	.618	.382	.504	.521	.633	.431	.529	.558
	+ CoT prompt	.673	.645	.504	.607	.474	.618	.344	.479	.531	.593	.417	.514	.533
ASKCALI	Regular	.668	.694	.552	.638	.533	.565	.339	.479	.526	.632	.383	.513	.543
	+ CoT prompt	.659	.648	.544	.617	.520	.530	.310	.453	.501	.587	.378	.489	.520
REFLECT	Regular	.644	.689	.562	.632	.581	.654	.381	.538	.499	.600	.386	.495	.555
	+ CoT prompt	.625	.677	.540	.614	.557	.608	.352	.506	.480	.569	.365	.471	.530
COOPERATE	Regular	.665	.736	.520	.640	.608	.716	.342	.555	.447	.576	.413	.478	.558
	+ CoT prompt	.638	.712	.504	.618	.607	.667	.334	.536	.458	.532	.422	.471	.542
COMPETE	Regular	.710	.697	.586	.664	.470	.635	.478	.528	.576	.630	.407	.538	.576
	+ CoT prompt	.706	.658	.534	.633	.445	.645	.436	.509	.556	.597	.385	.513	.552

Table 4: Ablation of baseline methods with regular prompting (just the query from the dataset) versus additionally appending a CoT prompt. Abstain Accuracy of abstain strategies on the nine datasets, averaged across four LLMs. Lower of the two values (comparing between *Regular* and + *CoT prompt*) are in red. Model specific results are reported in Appendix C.3.

abstention method.

5.3.2 Baselines with Reasoning Traces

When employing the baselines in our experiments, we use standard prompting with questions from the datasets rather than CoT prompting, both because of the answer extraction schema conventionally used for these baselines (Feng et al., 2024) and evidence that reasoning-style outputs lead to worse abstention abilities (Kirichenko et al., 2025; Zhang et al., 2025a). For completeness, we also provide the results of using CoT prompting and the resulting reasoning traces for each of the baselines. We employ the following prompt:

Provide step-by-step reasoning, with ‘Step 1:’, ‘Step 2:’, etc. followed by ‘Final answer:.’

CoT prompting degrades the abstention accuracy of baselines. As seen in Table 4, all baselines demonstrate worse abstention accuracy with CoT prompting across almost all datasets, with an average 2.6% decrease in abstention accuracy. We again see that our method TRACE INVERSION outperforms existing abstention baselines. Additionally, this shows that existing abstention methods are not well-suited for frontier models with reasoning-style outputs.

6 Conclusions and Future Work

We propose a novel framework for hallucinations in abstention as Query Misalignment: models hallucinate when they answer the *wrong* question. With this framework in mind, we create a new method

TRACE INVERSION which reconstructs a query based on the model response to identify misalignment in the response. Extensive experiments on nine datasets with four LLMs demonstrate that our method, TRACE INVERSION, achieves state-of-the-art performance in various abstention settings. In future work, we aim to explore additional abstention scenarios, such as stale and harmful questions.

Limitations

Our approach, TRACE INVERSION, leverages reasoning traces to help LLMs make abstention decisions. Our method was created to be particularly suitable for reasoning models. While it requires prompting LLMs three times, leading to higher inference costs compared to simpler prompting approaches, it is still not the most computationally expensive method compared with previous approaches like COOPERATE and COMPETE (Feng et al., 2024). To mitigate the cost, we also show the performance of simpler distance metrics rather than the ensemble model, which reduces the number of prompting requests while maintaining competitive performance.

Furthermore, our Query Misalignment Framework in Figure 2 provides a general framework for formulating the abstention task. Our framework may not be applicable to *all* types of hallucinations. As an example, models may have unfaithful reasoning or answer-rationale inconsistency where reasoning is sound throughout but still provide an incorrect final answer. Future work could explore additional types of abstention scenarios like harmful prompts.

Ethics Statement

Our study on mitigating bias and hallucinations in LLMs acknowledges the ethical implications of data-driven biases in AI. Our study is motivated by the importance of safe, deployable LLMs in high-stakes and real-world settings. All experiments were conducted using publicly available datasets and open-source models. No human participants were involved.

References

- Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.
- Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthoran Rajamanoharan, Neel Nanda, and Arthur Conmy. 2025. [Chain-of-thought reasoning in the wild is not always faithful](#). *Preprint*, arXiv:2503.08679.
- Josh Barua, Seun Eisape, Kayo Yin, and Alane Suhr. 2025. [Long chain-of-thought reasoning across languages](#). *Preprint*, arXiv:2508.14828.
- Siddhant Bhambri, Upasana Biswas, and Subbarao Kambhampati. 2025. [Do cognitively interpretable reasoning traces improve llm performance?](#) *Preprint*, arXiv:2508.16695.
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. 2025. [Reasoning models don't always say what they think](#). *Preprint*, arXiv:2505.05410.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Wenxuan Ding, Shangbin Feng, Yuhan Liu, Zhaoxuan Tan, Vidhisha Balachandran, Tianxing He, and Yulia Tsvetkov. 2024. [Knowledge crosswords: Geometric knowledge reasoning with large language models](#). *Preprint*, arXiv:2310.01290.
- Shangbin Feng, Weijia Shi, Yike Wang, Wenxuan Ding, Vidhisha Balachandran, and Yulia Tsvetkov. 2024. [Don't hallucinate, abstain: Identifying LLM knowledge gaps via multi-LLM collaboration](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14664–14690, Bangkok, Thailand. Association for Computational Linguistics.
- Sze Ching Evelyn Fung, Man Fai Wong, and Chee Wei Tan. 2023. Chain-of-thoughts prompting with language models for accurate math problem-solving. In *2023 IEEE MIT Undergraduate Research Technology Conference (URTC)*, pages 1–5. IEEE.
- Neha Gupta, Harikrishna Narasimhan, Wittawat Jitkritum, Ankit Singh Rawat, Aditya Krishna Menon, and Sanjiv Kumar. 2024. Language model cascades: Token-level uncertainty and beyond. *arXiv preprint arXiv:2404.10136*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Xinyang Hu, Fengzhuo Zhang, Siyu Chen, and Zhuoran Yang. 2024. [Unveiling the statistical foundations of chain-of-thought prompting methods](#). *Preprint*, arXiv:2408.14511.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and 1 others. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2):1–55.
- Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. 2023. Towards mitigating llm hallucination via self reflection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1827–1843.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Elliot Kim, Avi Garg, Kenny Peng, and Nikhil Garg. 2025. [Correlated errors in large language models](#). *Preprint*, arXiv:2506.07962.
- Polina Kirichenko, Mark Ibrahim, Kamalika Chaudhuri, and Samuel J Bell. 2025. Abstentionbench: Reasoning llms fail on unanswerable questions. *arXiv preprint arXiv:2506.09038*.

- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- Tomek Korbak, Mikita Balesni, Elizabeth Barnes, Yoshua Bengio, Joe Benton, Joseph Bloom, Mark Chen, Alan Cooney, Allan Dafoe, Anca Dragan, Scott Emmons, Owain Evans, David Farhi, Ryan Greenblatt, Dan Hendrycks, Marius Hobbhahn, Evan Hubinger, Geoffrey Irving, Erik Jenner, and 22 others. 2025. [Chain of thought monitorability: A new and fragile opportunity for ai safety](#). *Preprint*, arXiv:2507.11473.
- Sachin Kumar, Chan Young Park, and Yulia Tsvetkov. 2023. [Gen-z: Generative zero-shot text classification with contextualized label descriptions](#). *Preprint*, arXiv:2311.07115.
- Kentaro Kurihara, Masato Mita, Peinan Zhang, Shota Sasaki, Ryosuke Ishigami, and Naoaki Okazaki. 2025. Lctg bench: Llm controlled text generation benchmark. *arXiv preprint arXiv:2501.15875*.
- Walter Laurito, Benjamin Davis, Peli Grietzer, Tomas Gavenciak, Ada Böhm, and Jan Kulveit. 2024. Ai ai bias: Large language models favor their own generated content. *CoRR*.
- Seongyun Lee, Seungone Kim, Minju Seo, Yongrae Jo, Dongyoung Go, Hyeonbin Hwang, Jinho Park, Xiang Yue, Sean Welleck, Graham Neubig, Moontae Lee, and Minjoon Seo. 2025. [The cot encyclopedia: Analyzing, predicting, and controlling how a reasoning model will think](#). *Preprint*, arXiv:2505.10185.
- Zhenyu Li, Sunqi Fan, Yu Gu, Xiuxing Li, Zhichao Duan, Bowen Dong, Ning Liu, and Jianyong Wang. 2024. Flexkbqa: A flexible llm-powered framework for few-shot knowledge base question answering. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 18608–18616.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, and 31 others. 2023. [Holistic evaluation of language models](#). *Preprint*, arXiv:2211.09110.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. [Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts](#). *Preprint*, arXiv:2310.02255.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, Yansong Tang, and Dongmei Zhang. 2025. [Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct](#). *Preprint*, arXiv:2308.09583.
- Nishanth Madhusudhan, Sathwik Tejaswi Madhusudhan, Vikas Yadav, and Masoud Hashemi. 2025. [Do LLMs know when to NOT answer? investigating abstention abilities of large language models](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 9329–9345, Abu Dhabi, UAE. Association for Computational Linguistics.
- Zhiting Mei, Christina Zhang, Tenny Yin, Justin Lidard, Ola Shorinwa, and Anirudha Majumdar. 2025. [Reasoning about uncertainty: Do reasoning models know when they don’t know?](#) *Preprint*, arXiv:2506.18183.
- Yuhong Mo, Hao Qin, Yushan Dong, Ziyi Zhu, and Zhenglin Li. 2024. Large language model (llm) ai text generation detection based on transformer deep learning algorithm. *arXiv preprint arXiv:2405.06652*.
- John X. Morris, Wenting Zhao, Justin T. Chiu, Vitaly Shmatikov, and Alexander M. Rush. 2023. [Language model inversion](#). *Preprint*, arXiv:2311.13647.
- Murtaza Nazir, Matthew Finlayson, John X. Morris, Xiang Ren, and Swabha Swayamdipta. 2025. [Better language model inversion by compactly representing next-token distributions](#). *Preprint*, arXiv:2506.17090.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. 2021. [Show your work: Scratchpads for intermediate computation with language models](#). *Preprint*, arXiv:2112.00114.
- OpenAI. 2025. [gpt-oss-120b & gpt-oss-20b model card](#). *Preprint*, arXiv:2508.10925.
- Inkit Padhi, Manish Nagireddy, Giandomenico Cornacchia, Subhajit Chaudhury, Tejaswini Pedapati, Pierre Dognin, Keerthiram Murugesan, Erik Miehl, Martín Santillán Cooper, Kieran Fraser, Giulio Zizzo, Muhammad Zaid Hameed, Mark Purcell, Michael Desmond, Qian Pan, Zahra Ashktorab, Inge Vejsbjerg, Elizabeth M. Daly, Michael Hind, and 4 others. 2024. [Granite guardian](#). *Preprint*, arXiv:2412.07724.

- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [Llm evaluators recognize and favor their own generations](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 68772–68802. Curran Associates, Inc.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R. Bowman. 2022. [Bbq: A hand-built bias benchmark for question answering](#). *Preprint*, arXiv:2110.08193.
- Qizhi Pei, Lijun Wu, Zhuoshi Pan, Yu Li, Honglin Lin, Chenlin Ming, Xin Gao, Conghui He, and Rui Yan. 2025. Mathfusion: Enhancing mathematical problem-solving of llm through instruction fusion. *arXiv preprint arXiv:2503.16212*.
- Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2023. [SemEval-2023 task 3: Detecting the category, the framing, and the persuasion techniques in online news in a multilingual setup](#). In *Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023)*, pages 2343–2361, Toronto, Canada. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Anna Rogers, Olga Kovaleva, Matthew Downey, and Anna Rumshisky. 2020. Getting closer to ai complete question answering: A set of prerequisite real tasks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8722–8731.
- Adi Simhi, Itay Itzhak, Fazl Barez, Gabriel Stanovsky, and Yonatan Belinkov. 2025. [Trust me, i’m wrong: Llms hallucinate with certainty despite knowing the answer](#). *Preprint*, arXiv:2502.12964.
- Nishad Singhi, Hritik Bansal, Arian Hosseini, Aditya Grover, Kai-Wei Chang, Marcus Rohrbach, and Anna Rohrbach. 2025. When to solve, when to verify: Compute-optimal problem solving and generative verification for llm reasoning. *arXiv preprint arXiv:2504.01005*.
- Linxin Song, Taiwei Shi, and Jieyu Zhao. 2025. [The hallucination tax of reinforcement finetuning](#). *Preprint*, arXiv:2505.13988.
- Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Many Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2024. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. *arXiv preprint arXiv:2409.12183*.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adri Garriga-Alonso, and 1 others. 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on machine learning research*.
- YuHong Sun, Zhangyue Yin, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Hui Zhao. 2024. [Benchmarking hallucination in large language models based on unanswerable math word problem](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2178–2188, Torino, Italia. ELRA and ICCL.
- Annalisa Szymanski, Noah Ziemis, Heather A. Eicher-Miller, Toby Jia-Jun Li, Meng Jiang, and Ronald A. Metoyer. 2025. [Limitations of the llm-as-a-judge approach for evaluating llm outputs in expert knowledge tasks](#). In *Proceedings of the 30th International Conference on Intelligent User Interfaces, IUI ’25*, page 952–966, New York, NY, USA. Association for Computing Machinery.
- Hexiang Tan, Fei Sun, Wanli Yang, Yuanzhuo Wang, Qi Cao, and Xueqi Cheng. 2024. Blinded by generated contexts: How language models merge generated and retrieved contexts when knowledge conflicts? *arXiv preprint arXiv:2401.11911*.
- Shuchang Tao, Liuyi Yao, Hanxing Ding, Yuexiang Xie, Qi Cao, Fei Sun, Jinyang Gao, Huawei Shen, and Bolin Ding. 2024. [When to trust llms: Aligning confidence with response quality](#). *Preprint*, arXiv:2404.17287.
- Amir Taubenfeld, Yaniv Dover, Roi Reichart, and Ariel Goldstein. 2024. Systematic biases in llm simulations of debates. *arXiv preprint arXiv:2402.04049*.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. *arXiv preprint arXiv:2305.14975*.
- Christian Tomani, Kamalika Chaudhuri, Ivan Evtimov, Daniel Cremers, and Mark Ibrahim. 2024. [Uncertainty-based abstention in llms improves safety and reduces hallucinations](#). *Preprint*, arXiv:2404.10960.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. [Language models don’t always say what they think: Unfaithful explanations in chain-of-thought prompting](#). *Preprint*, arXiv:2305.04388.
- Emily Vaillancourt and Christopher Thompson. 2024. Instruction tuning on large language models to improve reasoning performance. *Authorea Preprints*.
- Thomas von Clarmann, Steven Compernelle, and Frank Hase. 2021. Truth and uncertainty. a critical discussion of the error concept versus the uncertainty

- concept. *Atmospheric Measurement Techniques Discussions*, 2021:1–26.
- Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. "kelly is a warm person, joseph is a role model": Gender biases in llm-generated reference letters. *arXiv preprint arXiv:2310.09219*.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024a. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Qian Wang, Zhanzhi Lou, Zhenheng Tang, Nuo Chen, Xuandong Zhao, Wenxuan Zhang, Dawn Song, and Bingsheng He. 2025a. [Assessing judging bias in large reasoning models: An empirical study](#). *Preprint*, arXiv:2504.09946.
- Yanbo Wang, Yongcan Yu, Jian Liang, and Ran He. 2025b. [A comprehensive survey on trustworthiness in reasoning with large language models](#). *Preprint*, arXiv:2509.03871.
- Zirui Wang, Mengzhou Xia, Luxi He, Howard Chen, Yitao Liu, Richard Zhu, Kaiqu Liang, Xindi Wu, Haotian Liu, Sadhika Malladi, Alexis Chevalier, Sanjeev Arora, and Danqi Chen. 2024b. [Charxiv: Charting gaps in realistic chart understanding in multimodal llms](#). *Preprint*, arXiv:2406.18521.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. 2025. Know your limits: A survey of abstention in large language models. *Transactions of the Association for Computational Linguistics*, 13:529–556.
- Yuyang Wu, Yifei Wang, Ziyu Ye, Tianqi Du, Stefanie Jegelka, and Yisen Wang. 2025. [When more is less: Understanding chain-of-thought length in llms](#). *Preprint*, arXiv:2502.07266.
- Yuxin Xiao, Shan Chen, Jack Gallifant, Danielle Bitterman, Thomas Hartvigsen, and Marzyeh Ghassemi. 2025. Kscope: A framework for characterizing the knowledge status of language models. *arXiv preprint arXiv:2506.07458*.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2023. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*.
- Rongwu Xu, Zehan Qi, Zhijiang Guo, Cunxiang Wang, Hongru Wang, Yue Zhang, and Wei Xu. 2024a. [Knowledge conflicts for llms: A survey](#). *Preprint*, arXiv:2403.08319.
- Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaoze Liu, Xingyao Wang, Yangyi Chen, and Jing Gao. 2024b. [SaySelf: Teaching LLMs to express confidence with self-reflective rationales](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5985–5998, Miami, Florida, USA. Association for Computational Linguistics.
- Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, and William Wang. 2024c. [Pride and prejudice: LLM amplifies self-bias in self-refinement](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15474–15492, Bangkok, Thailand. Association for Computational Linguistics.
- Yasin Abbasi Yadkori, Ilja Kuzborskij, András György, and Csaba Szepesvári. 2024. [To believe or not to believe your llm: Iterative prompting for estimating epistemic uncertainty](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 58077–58117. Curran Associates, Inc.
- Kai Yan, Yufei Xu, Zhengyin Du, Xuesong Yao, Zheyu Wang, Xiaowen Guo, and Jiecao Chen. 2025. [Recitation over reasoning: How cutting-edge language models can fail on elementary school-level reasoning problems?](#) *Preprint*, arXiv:2504.00509.
- Hang Yang, Hao Chen, Hui Guo, Yineng Chen, Ching-Sheng Lin, Shu Hu, Jinrong Hu, Xi Wu, and Xin Wang. 2024a. [Llm-medqa: Enhancing medical question answering through case studies in large language models](#). *arXiv preprint arXiv:2501.05464*.
- Wen Yang, Minpeng Liao, and Kai Fan. 2024b. [Markov chain of thought for efficient mathematical reasoning](#). *arXiv preprint arXiv:2410.17635*.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, and 3 others. 2024. [Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi](#). *Preprint*, arXiv:2311.16502.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). *Preprint*, arXiv:2203.14465.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) *Preprint*, arXiv:1905.07830.
- Qingjie Zhang, Yujia Fu, Yang Wang, Liu Yan, Tao Wei, Ke Xu, Minlie Huang, and Han Qiu. 2025a. [On the self-awareness of large reasoning models' capability boundaries](#). *Preprint*, arXiv:2509.24711.

Xinlu Zhang, Zhiyu Zoey Chen, Xi Ye, Xianjun Yang, Lichang Chen, William Yang Wang, and Linda Ruth Petzold. 2025b. Unveiling the impact of coding data instruction fine-tuning on large language models reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 25949–25957.

Ziyao Zhang, Chong Wang, Yanlin Wang, Ensheng Shi, Yuchi Ma, Wanjun Zhong, Jiachi Chen, Mingzhi Mao, and Zibin Zheng. 2025c. Llm hallucinations in practical code generation: Phenomena, mechanism, and mitigation. *Proceedings of the ACM on Software Engineering*, 2(ISSTA):481–503.

Ziyin Zhang, Zhaokun Jiang, Lizhen Xu, Hongkun Hao, and Rui Wang. 2024. [Multiple-choice questions are efficient and robust llm evaluators](#). *Preprint*, arXiv:2405.11966.

Chengshuai Zhao, Zhen Tan, Pingchuan Ma, Dawei Li, Bohan Jiang, Yancheng Wang, Yingzhen Yang, and Huan Liu. 2025. [Is chain-of-thought reasoning of llms a mirage? a data distribution lens](#). *Preprint*, arXiv:2508.01191.

Bin Zhu, Hailong Yin, Jingjing Chen, and Yu-Gang Jiang. 2025. [Benchmarking gaslighting negation attacks against reasoning models](#). *Preprint*, arXiv:2506.09677.

A Experimental Details

A.1 Dataset Details

Each question included in our study comprises a prompt (including a question and an optional context), an “answerable flag” binary label, and optional reference answers for samples where abstention is not required. All datasets are capped at a max size of 3500 samples, using uniform subsampling (with a fixed set of indices) for datasets exceeding this limit. The following datasets were implemented as part of our study.

- **MMLU** (Hendrycks et al., 2021) is a multiple-choice dataset for general knowledge QA including elementary mathematics, US history, computer science, law, and more.
- **Knowledge Crosswords (K-Crosswords)** (Ding et al., 2024) is a geometric knowledge reasoning benchmark consisting of incomplete knowledge networks bounded by structured factual constraints.
- **Hellaswag** (Zellers et al., 2019) tests commonsense natural language inference.
- **Propaganda** (Piskorski et al., 2023) tasks LLMs with identifying the 23 persuasion tactics in a long news article based on their internal knowledge.
- **Bias Benchmark for Question Answering (BBQ)** (Parrish et al., 2022) contains questions about stereotypical associations in both fully specified and underspecified contexts, where the fully-specified form may negate the stereotype. We consider questions with missing or ambiguous context as should abstain and those with sufficient information as should not abstain.
- **‘Misconceptions’** from BIG-Bench (Srivastava et al., 2023) measures whether a model can discern popular misconceptions from the truth.
- **Quail** (Rogers et al., 2020) is a reading comprehension dataset containing answerable and unanswerable passage-based questions
- **GSM-MC** (Zhang et al., 2024) is a multiple-choice dataset constructed by collecting answers and incorrect predictions on GSM8K from 60 open-source models.
- **Unanswerable Math Word Problems (UMWP)** (Sun et al., 2024) has questions drawn from other math datasets and modified to be unanswerable.

A.2 Model Details

Model Initialization. We support multiple large language models (LLMs) through a unified initialization function. The implementation maps human-readable names (e.g. qwen_32b) to their respective HuggingFace or vLLM model checkpoints. Models are loaded with bfloat16 precision and GPU memory utilization capped at 80% for efficiency. Chat-oriented models (e.g., DeepSeek, Qwen) are automatically wrapped with their tokenizer’s chat template. For gpt-oss-120b, we set the reasoning effort parameter to ‘medium’. Our code also enables easy integration of new models.

Sampling Parameters. Responses are generated with configurable temperature ($T = 0.1$ by default), a maximum of 1024 new tokens, and optional token-level probabilities. The code supports exponential backoff retries (up to 10 attempts) to ensure robustness against API or inference errors. We perform three experiments for each approach, model, and dataset triplet with three seeds 0, 1, and 2.

Answer Parsing. Since models may return heterogeneous outputs, we implement rule-based answer parsing with multiple heuristics (e.g., “Answer: A”, “The correct answer is B”, or isolated multiple-choice options). Responses that cannot be parsed are labeled with a sentinel “Z” to indicate incorrectness. Across all nine datasets, less than 3% of responses were not parsed.

B Baseline Details

B.1 Calibration-based baselines

For confidence estimation methods, we use a held-out development set $\mathcal{H} = \{(q_i, \bar{a}_i)\}_{i=1}^N$. For each question q_i , the LLM produces an answer $a_i = \text{LLM}(q_i)$ and calculate a confidence score $p_i \in [0, 1]$. We define correctness labels as

$$y_i = \begin{cases} 1 & \text{if } a_i = \bar{a}_i, \\ 0 & \text{if } a_i \neq \bar{a}_i. \end{cases}$$

Candidate thresholds are taken from a discretized grid $\mathcal{T} = \{0.01, 0.02, \dots, 0.99\}$. For each threshold $t \in \mathcal{T}$, we apply the abstention rule and compute the abstain error

$$\hat{a}_i(t) = \begin{cases} \text{abstain}, & p_i < t, \\ a_i, & p_i \geq t, \end{cases}$$

$$E(t) = \sum_{i=1}^N \mathbf{1}(p_i < t \wedge y_i = 1) + \mathbf{1}(p_i \geq t \wedge y_i = 0).$$

The first term in $E(t)$ penalizes unnecessary abstentions on correct answers, while the second penalizes failures to abstain on incorrect answers. The abstention threshold is then chosen as $p^* = \arg \min_{t \in \mathcal{T}} E(t)$. At inference time, the model answers if $p_i \geq p^*$ and abstains otherwise (Feng et al., 2024). The following two methods use internal calibration and verbalized calibration to estimate model confidence.

Token probability (PROBS) We compute the confidence score p_i for a question using the top- k token probabilities over the entire answer span where P is the language model’s predicted token distribution at the final answer index. Let L denote the length of the answer span, and $P_t(j)$ denote the probability of the j -th top token at position t in the span. Then:

$$p_i = \frac{1}{L} \sum_{t=1}^L \frac{1}{k} \sum_{j=1}^k \log P_t(j)$$

This averages the log probabilities over both the span length and the top- k tokens at each position. We use $k = 5$ for this baseline.

Ask for calibration (ASKCALI) The confidence score p_i is the LLM-provided calibration estimate (Tian et al., 2023). Full prompts for each method are provided in Appendix B.4.

B.2 Prompting-based baseline

Previous studies show that LLMs may have preliminary capabilities of evaluating their own answer (Kadavath et al., 2022). The following baseline utilizes LLMs to assess and review the model’s own outputs. Based on the model’s assessment, an abstention decision is made.

Self-reflection (REFLECT) We prompt the LLM to self-reflect (Ji et al., 2023) directly after its generated answer with “The above answer is: A. True B. False”. LLMs should abstain when they deem the generated answer a_i as false.

B.3 Collaboration-based baselines

Collaborative baselines (Feng et al., 2024) are an extension of prompting baselines where multiple LLMs are used to arrive at an abstention decision.

Cooperative system (COOPERATE) We generate k experts from the LLM on domains d_1, \dots, d_k through prompting-based self-specialization (Feng et al., 2024). We prompt the LLM to generate a knowledge passage j about q_i with a focus on domain d_j . A domain-specific feedback is then generated by prepending the knowledge passage $f_j = \text{LLM}(\text{knowledge}_j, q_i, a_i)$ and prompting the model to respond as a reviewer. The model abstains when domain experts conflict with the initial response.

Competitive system (COMPETE) Given initial answer a_i for question q_i , we prompt the LLM to generate k alternative answers $b = \{b_1, \dots, b_k\}$. We then instruct the LLM to answer q_i again with conflicting information from an answer in answer set b prepended (Feng et al., 2024). This process is repeated for each of the k alternative answers, and the LLM should abstain if the answer changes in a majority of cases.

B.4 Prompts for methods

We provide the exact prompts we employed in the study in the order they first appear (see Table 5).

For our trace inversion procedure, we use the following prompt for reconstructing the model:

Query Reconstruction Prompt You are a puzzle solver. Given the following reasoning trace, reconstruct the initial question by interpreting the steps in the reasoning trace. Do not answer the question.
Reasoning Trace:
{reasoning trace}

Reconstructed query:

The following prompt is used for the TrInv-LLM distance metric.

TrInv-LLM Do the following two prompts convey the same framing, intent, and context?

Prompt 1: {q1}
Prompt 2: {q2}

Select YES or NO:
Final answer:

C Supplemental Experimental Results

C.1 Reliable Accuracy Results

Table 1 reports Abstain Accuracy (A-Acc), defined as $\frac{TP+TN}{TP+TN+FP+FN}$ (Feng et al., 2024). Abstain Accuracy measures the correctness of abstention decisions. That is, an LLM should abstain when it would produce an incorrect answer and should not abstain when it would give a correct answer. In Table 6, we report the Reliable Accuracy: the correctness of answered questions. Reliable Accuracy is defined as $\frac{TN}{TN+FN}$.

TRACE INVERSION has the highest Reliable Accuracy in 31 out of 36 settings. Similar to our findings with Abstain Accuracy, our TRACE INVERSION method continues to be the most consistent compared to all baselines across domains.

C.2 Complete Query Misalignment Metric Analysis

In Table 7, we provide the model specific results from Table 3. These results also show the comparative domain specialization of SE for Math & Knowledge and GROUND for Biases & Safety. Additionally, with larger models DeepSeek-R1-Distill-Qwen-32B and gpt-oss-120b, we see that the ensemble variant actually outperforms the individual SE module for the Math & Knowledge datasets, demonstrating that the relative performance of the other modules is stronger in the Math & Knowledge domain with a stronger base LLM.

C.3 Complete Baselines with CoT Analysis

Lastly, we provide a complete table showing the degradation of Abstain Accuracy when using CoT prompts with the baselines (see Table 8). The disaggregated results highlight the extent of CoT degradation in certain settings. For example, the REFLECT method demonstrates a 11.8% decrease in Abstain Accuracy for gpt-oss-120b on the MMLU dataset and 14.0% decrease for phi-4 on the BBQ dataset.

Method	Prompt
ASKCALI	<p>Provide your best guess for the following question. Give <i>only</i> the guess. Example: {most likely guess}.</p> <p>Question: {question}.</p> <p>[<i>LLM-generated answer</i>].</p> <p>Provide the probability your guess is correct. Probability: [<i>LLM-generated probability in [0, 1]</i>].</p>
REFLECT	<p>Question: {question}. Answer: [<i>LLM-generated proposed answer</i>].</p> <p>Is the above answer correct? A. True B. False.</p> <p>Final answer: [<i>LLM-generated A/B</i>].</p>
COOPERATE	<p>Question: {question}. Answer: [<i>LLM-generated proposed answer</i>].</p> <p>For each domain in {"factual", "commonsense", "mathematical"}, generate domain-specific knowledge.</p> <p>Knowledge: {generated domain knowledge}</p> <p>Question: {question}.</p> <p>Answer: {generated proposed answer}.</p> <p>Review the proposed answer and provide feedback on correctness.</p> <p>Feedback 1: {expert 1}; ...; Feedback <i>k</i>: {expert <i>k</i>}.</p> <p>Based on feedback, is the proposed answer correct? A. True B. False.</p> <p>Final answer: [<i>LLM-generated A/B</i>].</p>
COMPETE	<p>Question: {question}. Answer: [<i>LLM-generated proposed answer</i>].</p> <p>If multiple-choice, randomly select an unchosen option; otherwise propose an alternative answer.</p> <p>Alternative answer: [<i>alternative answer</i>].</p> <p>Generate supporting knowledge.</p> <p>Knowledge: [<i>generated alternative passage</i>].</p> <p>Answer the question using this knowledge. New answer: [<i>new generated answer</i>].</p> <p>If the new answer matches the original, set abstain=False; otherwise set abstain=True.</p>

Table 5: Exact LLM prompts used for each method.

Method	Math & Knowledge				Comprehension				Biases & Safety				Overall
	MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis	Prop	BBQ	Overall	
PHI-4													
PROBS	.545	.584	.614	.581	.463	.634	.363	.487	.439	.636	.522	.533	.533
ASKCALI	.404	.585	.711	.567	.626	.588	.358	.524	.598	.538	.358	.498	.530
REFLECT	.412	.466	.669	.516	.591	.757	.550	.633	.449	.422	.522	.464	.538
COOPERATE	.428	<u>.682</u>	.634	.581	.532	.731	<u>.656</u>	<u>.640</u>	.272	<u>.677</u>	<u>.601</u>	.516	.579
COMPETE	<u>.565</u>	.555	<u>.725</u>	<u>.615</u>	.445	.757	.687	.630	.451	.602	.531	.528	<u>.591</u>
TRACE INVERSION	.741	.778	.824	.781	<u>.606</u>	.782	.648	.679	<u>.597</u>	.787	.809	.731	.730
QWEN2.5-32B													
PROBS	.758	.777	.635	.723	.408	.651	.528	.529	.456	.706	.560	.574	.609
ASKCALI	.669	.683	.804	.718	.517	.522	.565	.535	.382	.654	.527	.521	.591
REFLECT	.681	.633	.780	.698	<u>.522</u>	.670	.492	.561	.506	.648	<u>.461</u>	.538	.599
COOPERATE	.661	.727	.611	.666	.433	.697	.554	.561	.297	.612	.601	.504	.577
COMPETE	.594	.711	<u>.897</u>	<u>.734</u>	.416	<u>.684</u>	<u>.621</u>	<u>.574</u>	<u>.572</u>	<u>.731</u>	<u>.573</u>	<u>.625</u>	<u>.644</u>
TRACE INVERSION	.812	.797	.999	.869	.718	.746	.916	.793	.735	.760	<u>.585</u>	.693	.785
DEEPSEEK-R1-DISTILL-QWEN-32B													
PROBS	.853	.793	.708	.785	.720	.596	.569	.628	.471	.568	.570	.536	.650
ASKCALI	.783	.770	.799	.784	.548	.358	.457	.454	.559	.716	.596	<u>.624</u>	.621
REFLECT	.738	.730	.844	.771	.565	.576	.663	.601	.521	.587	.381	.496	.623
COOPERATE	.894	.660	<u>.886</u>	<u>.813</u>	.626	.804	.397	.609	.565	.464	.479	.503	.642
COMPETE	.784	.738	.707	.743	.381	.566	.354	.434	.537	.551	.508	.532	.570
TRACE INVERSION	.930	.894	.889	.904	.762	.818	.689	.756	.793	.760	<u>.579</u>	.710	.790
GPT-OSS-120B													
PROBS	.850	.867	.812	.843	.639	.533	.626	.599	.634	.718	.527	.626	.690
ASKCALI	.742	.842	.623	.736	.544	.774	.635	.651	.551	.533	.566	.550	.646
REFLECT	.821	.677	.782	.760	.414	.684	.523	.541	.590	.597	.648	.612	.638
COOPERATE	.820	.792	.584	.732	.663	.894	.463	.673	.628	.501	.514	.548	.651
COMPETE	.805	.807	.757	.790	.669	.732	.689	.697	.577	.679	<u>.650</u>	<u>.635</u>	<u>.707</u>
TRACE INVERSION	.854	.941	.919	.905	.677	.899	.757	.778	.661	.729	.831	.740	.807

Table 6: Reliable Accuracy of abstain strategies on nine datasets and four LLMs. Each number reported is the average of three seeds. We group the nine datasets (MMLU, GSM, UMWP, K-Crosswords, Hellaswag, Quail, Misconceptions, Propaganda, and BBQ) according to domain and provide aggregate results per domain and across all datasets. Approaches are color-coded per category: **calibration**, **prompting**, **collaboration**, and **our method**.

Method	Math & Knowledge				Comprehension				Biases & Safety				Overall
	MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis	Prop	BBQ	Overall	
PHI-4													
SE	.876	.798	.865	.846	.538	.727	.619	.628	.567	.608	.656	.610	.695
TrInv-LLM	.710	<u>.743</u>	.752	<u>.735</u>	.684	.774	.723	<u>.727</u>	.555	.556	.639	.584	.682
GROUND	.638	<u>.721</u>	.556	<u>.638</u>	.658	.760	.609	<u>.675</u>	.696	.826	.772	.765	.693
TRACE INVERSION	<u>.712</u>	.733	<u>.757</u>	<u>.734</u>	.663	.830	.694	.729	.649	.710	.614	<u>.658</u>	.707
QWEN2.5-32B													
SE	.855	.902	.795	.851	.473	.698	.638	.603	.570	.632	.579	.594	.683
TrInv-LLM	.737	.763	.707	.736	<u>.727</u>	.798	.724	.750	.580	.623	.659	.621	.702
GROUND	<u>.702</u>	.775	.583	<u>.687</u>	.662	.784	.535	.660	.691	.748	.753	.731	.693
TRACE INVERSION	.719	<u>.850</u>	<u>.788</u>	<u>.786</u>	.789	.712	.717	<u>.739</u>	.670	.734	.668	<u>.691</u>	.739
DEEPSEEK-R1-DISTILL-QWEN-32B													
SE	.812	.823	.800	<u>.812</u>	.552	.770	.558	.627	.545	.636	<u>.627</u>	.603	.680
TrInv-LLM	.771	.856	.720	<u>.782</u>	.724	.811	.725	.753	.552	.580	.603	.578	.705
GROUND	.692	.764	.524	.660	<u>.657</u>	.816	.611	<u>.695</u>	<u>.627</u>	.818	.814	.753	.703
TRACE INVERSION	.915	.883	.812	.870	.616	.731	<u>.612</u>	<u>.653</u>	.713	.712	.602	<u>.676</u>	.733
GPT-OSS-120B													
SE	.849	.841	.896	.862	.569	.689	.605	.621	.537	.592	.601	.577	.687
TrInv-LLM	.778	.777	.701	.752	.681	.769	.657	<u>.702</u>	.589	.632	.595	.605	.687
GROUND	.704	.729	.561	<u>.664</u>	.767	.761	.606	.711	<u>.653</u>	.808	.820	.760	<u>.712</u>
TRACE INVERSION	.885	.851	<u>.778</u>	<u>.838</u>	.585	.814	.640	.680	.711	.804	<u>.695</u>	<u>.737</u>	.751

Table 7: Ablation of TRACE INVERSION by looking at the Abstain Accuracy performance of the method with individual distance metrics. Each number reported is the average of three seeds. Performance is shown for each of the four LLMs. Best results in **bold** and second best in underline.

Method	Variant	Math & Knowledge			Comprehension			Biases & Safety			Overall			
		MMLU	GSM	UMWP	Overall	KC	HS	Qu	Overall	Mis		Prop	BBQ	Overall
PHI-4														
PROBS	Regular	.477	.509	.488	.491	.451	.666	.303	.473	.512	.624	.332	.489	.485
	+ CoT prompt	.449	.426	.395	.423	.424	.548	.286	.419	.404	.559	.236	.400	.414
ASKCALI	Regular	.471	.504	.506	.494	.550	.612	.307	.490	.552	.618	.299	.490	.491
	+ CoT prompt	.390	.394	.494	.426	.475	.565	.169	.403	.574	.613	.220	.469	.432
REFLECT	Regular	.379	.541	.438	.453	.633	.771	.405	.603	.455	.515	.411	.460	.505
	+ CoT prompt	.278	.487	.380	.381	.556	.794	.388	.579	.495	.411	.271	.392	.451
COOPERATE	Regular	.424	.685	.420	.510	.504	.718	.414	.545	.369	.598	.422	.463	.506
	+ CoT prompt	.442	.643	.385	.490	.408	.691	.394	.498	.400	.601	.398	.466	.485
COMPETE	Regular	.578	.547	.516	.547	.426	.690	.533	.550	.467	.600	.312	.460	.519
	+ CoT prompt	.507	.424	.558	.497	.403	.669	.468	.513	.423	.508	.279	.404	.471
QWEN2.5-32B														
PROBS	Regular	.741	.711	.512	.655	.329	.551	.397	.426	.456	.641	.498	.532	.537
	+ CoT prompt	.703	.577	.418	.566	.339	.501	.381	.407	.394	.581	.416	.464	.479
ASKCALI	Regular	.711	.684	.601	.665	.473	.513	.319	.435	.451	.647	.401	.500	.533
	+ CoT prompt	.682	.582	.619	.628	.426	.463	.230	.373	.443	.625	.323	.464	.488
REFLECT	Regular	.689	.731	.639	.686	.618	.610	.305	.511	.415	.621	.411	.482	.560
	+ CoT prompt	.646	.669	.549	.622	.536	.625	.236	.466	.382	.562	.334	.426	.504
COOPERATE	Regular	.637	.725	.420	.594	.516	.618	.345	.493	.322	.603	.424	.450	.512
	+ CoT prompt	.639	.674	.406	.573	.481	.608	.288	.459	.348	.544	.414	.436	.489
COMPETE	Regular	.688	.747	.656	.697	.509	.672	.488	.556	.667	.704	.490	.620	.625
	+ CoT prompt	.629	.590	.623	.614	.545	.614	.492	.551	.596	.608	.427	.544	.569
DEEPSEEK-R1-DISTILL-QWEN-32B														
PROBS	Regular	.770	.739	.600	.703	.653	.622	.412	.562	.503	.644	.487	.545	.603
	+ CoT prompt	.706	.594	.528	.610	.622	.501	.404	.509	.368	.618	.411	.466	.528
ASKCALI	Regular	.765	.784	.601	.717	.557	.454	.317	.443	.511	.672	.404	.529	.563
	+ CoT prompt	.735	.687	.568	.663	.498	.445	.198	.380	.506	.658	.335	.500	.514
REFLECT	Regular	.748	.744	.639	.710	.633	.611	.510	.585	.510	.618	.310	.479	.591
	+ CoT prompt	.679	.637	.537	.618	.637	.582	.429	.550	.519	.496	.192	.402	.523
COOPERATE	Regular	.849	.715	.722	.762	.707	.718	.298	.574	.488	.511	.420	.473	.603
	+ CoT prompt	.861	.593	.755	.736	.594	.690	.241	.508	.484	.444	.396	.441	.562
COMPETE	Regular	.784	.647	.556	.662	.329	.488	.301	.373	.583	.605	.338	.509	.515
	+ CoT prompt	.731	.570	.514	.605	.313	.466	.297	.359	.465	.547	.317	.443	.469
GPT-OSS-120B														
PROBS	Regular	.767	.799	.577	.714	.617	.632	.417	.555	.614	.624	.408	.549	.606
	+ CoT prompt	.738	.731	.511	.660	.571	.522	.417	.503	.550	.526	.284	.453	.539
ASKCALI	Regular	.725	.804	.501	.677	.552	.682	.414	.549	.588	.590	.426	.535	.587
	+ CoT prompt	.653	.741	.532	.642	.497	.647	.344	.496	.604	.592	.343	.513	.550
REFLECT	Regular	.759	.739	.533	.677	.438	.623	.305	.455	.615	.644	.413	.557	.563
	+ CoT prompt	.641	.718	.462	.607	.355	.602	.227	.394	.632	.559	.335	.509	.503
COOPERATE	Regular	.749	.817	.520	.695	.704	.812	.312	.609	.607	.590	.388	.528	.611
	+ CoT prompt	.754	.754	.534	.681	.585	.767	.266	.539	.584	.604	.299	.496	.572
COMPETE	Regular	.788	.847	.614	.750	.616	.691	.590	.632	.588	.611	.487	.562	.648
	+ CoT prompt	.705	.763	.636	.701	.575	.655	.590	.607	.456	.533	.413	.467	.592

Table 8: Abstain Accuracy of abstain strategies on nine datasets for each of the four LLMs. Each number reported is the average of three seeds. We group the nine datasets (MMLU, GSM, UMWP, K-Crosswords, Hellaswag, Quail, Misconceptions, Propaganda, and BBQ) according to domain and provide aggregate results per domain and across all datasets. Approaches are color-coded per category: calibration, prompting, collaboration, and our method.