# TRAINING-FREE LLM VERIFICATION VIA RECYCLING FEW-SHOT EXAMPLES

**Anonymous authors**Paper under double-blind review

#### **ABSTRACT**

Although LLMs have achieved remarkable performance, the inherent stochasticity of their reasoning process and varying conclusions present significant challenges. Majority voting or Best-of-N with external verification models has been explored to find the most promising solution among multiple LLM outputs. However, these approaches have certain limitations, such as limited applicability or the cost of an additional training step. To address this problem, we propose a novel and effective framework that Recycles Few-shot examples to verify LLM outputs (ReFeri). Our key idea is to additionally utilize the given few-shot examples to evaluate the candidate outputs of the target query, not only using them to generate outputs as the conventional few-shot prompting setup. Specifically, ReFeri evaluates the generated outputs by combining two different scores, designed motivated by Bayes' rule, and subsequently selects the candidate that is both confidently determined and contextually coherent through a few additional LLM inferences. Experiments with three different LLMs and across seven diverse tasks demonstrate that our framework significantly improves the accuracy of LLM-achieving an average gain of 4.5%-through effective response selection, without additional training.

#### 1 Introduction

Recently, large language models (LLMs) have shown remarkable performance in many real-world tasks involving complex reasoning, such as math, coding, and robotics (Anthropic, 2024; Dubey et al., 2024; OpenAI, 2024c; Team et al., 2023). To enhance the reasoning capacity of LLMs, various approaches have been proposed, ranging from in-context learning at test time (Wei et al., 2022; Kojima et al., 2022) to recent RL training method (Qu et al., 2024; Guo et al., 2025). Despite these improvements, the inherent stochastic nature of LLM still presents significant challenges, since different reasoning paths can be generated for the same input and can lead to varying conclusions (Kadavath et al., 2022; Wang & Zhou, 2024; Qiu & Miikkulainen, 2024). Majority voting approaches, such as self-consistency (Wang et al., 2023b; Aggarwal et al., 2023), have been widely adopted to reduce such randomness by aggregating multiple LLM outputs and determining a single prediction. However, this approach is only applicable when the answer can be easily extracted from the output and aggregated. Consequently, it is difficult to apply to open-ended text generation tasks such as summarization and personalized chatbot (Stiennon et al., 2020; Salemi et al., 2024).

To address this challenge, finding the most promising one among multiple LLM outputs using a specific selection method, often called *Best-of-N*, has recently gained attention (Snell et al., 2024; Gui et al., 2024). For instance, one of the most representative approaches is to score each output using external verification models such as Outcome Reward Models (ORMs) (Cobbe et al., 2021; Uesato et al., 2022) or Process Reward Models (PRMs) (Lightman et al., 2024; Wang et al., 2024a), and then selecting the highest-scoring output. However, to obtain these reward models, training with a large amount of task-specific labeled data is often necessary; therefore, applying this framework to specific target domain, which is far from well-explored domains such as math and coding, is challenging. Prompting LLM to select the most promising output—such as *LLM-as-judge*—is considerable to remove the reliance on the verification model (Chen et al., 2023; Zheng et al., 2023). However, this approach is only effective when the given LLM has sufficient intrinsic knowledge for the target domain; consequently, it often requires separate training steps and datasets again to achieve sufficient performance (Yuan et al., 2024; Mahan et al., 2024; Zhang et al., 2025).

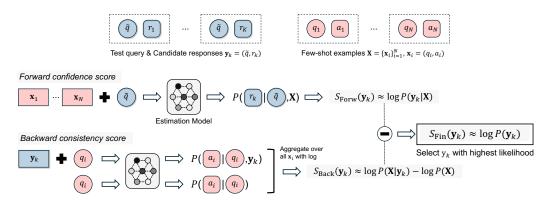


Figure 1: **An overview of ReFeri.** For *K* candidate responses from LLMs, ReFeri assigns each candidate a forward confidence score (*how likely candidate is to be generated conditioned on few-shot examples*) and a backward consistency score (*how well candidate explains the answers of few-shot examples*). Then, the response with the best joint score is selected as the final answer.

Motivated by this, we suggest a new perspective: *utilization of few-shot examples to verify and select among multiple LLM outputs*. As recent LLMs have been trained with an extensive instruction tuning step, they often exhibit better performance without few-shot examples (Guo et al., 2025; Sprague et al., 2025), and hence using these examples at test time is recently losing attention (see results in Table 1). However, we argue that using few-shot examples is still one of the easiest and most direct ways to let LLMs know how to solve the given task with human prior knowledge, even if LLMs have not encountered it before. Therefore, in this work, we provide a new framework that enables better exploitation of few-shot examples by using them not only for generating multiple outputs, but also for selecting the most promising one.

**Contribution.** In this work, we propose **ReFeri**, a novel and effective framework that **Re**cycles **Few**-shot examples to ve**ri**fy LLM outputs. The core idea of ReFeri is additionally utilizing the given few-shot examples to evaluate the candidate outputs of the target query, not only using them to generate outputs as conventional few-shot in-context learning. Specifically, ReFeri estimates the likelihood of the generated outputs by decomposing it into two different scores (*forward confidence score* and *backward consistency score*) conditioned on few-shot examples, which are derived from *Bayes' rule*. The forward confidence score measures the likelihood of candidate outputs given the few-shot examples and the test query, favoring more confident ones. On the other hand, the backward consistency score measures whether conditioning on the candidate output well explains the likelihood of the few-shot examples compared to conditioning on their queries alone. By combining these scores, ReFeri selects the candidate that is both confidently determined and contextually coherent through a few additional LLM inferences. Consequently, ReFeri does not require additional model training to select the most promising output, and allows better leverage of both intrinsic knowledge of LLM and human prior within the provided few-shot examples. See Figure 1 for the illustration.

We validate the effectiveness of ReFeri across three different LLMs (GPT-4o, GPT-4o-mini, and LLaMA-3.1-8B) and seven different benchmarks. When selecting one response among five candidates generated by few-shot chain-of-thought (CoT) prompting, ReFeri consistently outperforms other training-free selection across all tasks, with an average gain of 4.5% over random selection and 2.4% over prompt-based selection methods (see Figure 2). ReFeri also scales reliably with the number of candidate responses, demonstrating its practical utility in test-time scaling. To better understand the behavior of ReFeri, we conduct more complementary analyses, showing that our method is robust to variations in few-shot example selection, prompt template choices, and the choice of model used for likelihood estimation; ReFeri yields consistent improvements without reliance on specific prompt templates or few-shot ex-

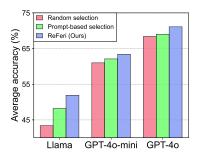


Figure 2: **Summary of results.** Average accuracy across seven benchmarks with training-free selection methods. ReFeri shows consistent effectiveness (see Section 3.2).

<sup>&</sup>lt;sup>1</sup>In-context learning uses given few-shot examples as additional input context upon the target query.

amples. Moreover, when combined with smaller likelihood estimators, ReFeri significantly reduces the cost per query compared to larger baseline models while outperforming them.

## 2 Training-free LLM Verification via Recycling Few-shot Data

## 2.1 PROBLEM FORMULATION

Let us denote LLM as  $\mathcal{M}$  and a given test query as  $\widetilde{q}$ . We assume that we have N-shot examples  $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N, \mathbf{x}_i = (q_i, a_i)$  where  $q_i$  is another input query from the same task and  $a_i$  is the ground-truth answer, which can be provided by human annotator or generated by LLM itself. Then, few-shot prompting incorporates the few-shot examples  $\mathbf{x}_i$  in  $\mathbf{X}$  as additional input context to obtain the response  $r_k$ , which is expected to be improved thanks to the in-context learning capability of LLMs:

$$r_k \sim \mathcal{M}(\widetilde{q}, \mathbf{X}),$$
 (1)

where multiple non-identical predictions  $r_k, k=1,\ldots,K$  can be sampled. Then, our goal is to find the most appropriate response  $r_{k^*}$  among them. For example, the self-consistency method (Wang et al., 2023b) simply applies majority voting to determine the single prediction. On the other hand, the best-of-K method uses the external verifier such as reward models (Cobbe et al., 2021; Lightman et al., 2024) to score the predictions and select the highest scored one. Formally, with the external verifier  $R_{\phi}$ , it can be described as below:

$$r_{k^*} = \arg\max_{k=1,\dots,K} R_{\phi}(\mathbf{y}_k), \tag{2}$$

where  $\mathbf{y}_k = (\widetilde{q}, r_k)$ . While these approaches are widely used in practice, there are certain challenges due to the limited applicability and the need for a verification model for the target task.

#### 2.2 REFERI: VERIFYING LLM OUTPUTS WITH BAYES' RULE WITH FEW-SHOT DATA

In this section, we introduce a framework that selects candidates from LLM by **Re**cycling **Few**-shot examples for the verification (**ReFeri**). The core idea of ReFeri is to leverage few-shot examples not only for generation but also for validation, thereby recycling them to score and select answers without additional training. Specifically, ReFeri estimates the plausibility of each answer candidate by combining two complementary signals: (1) a *forward confidence score* which captures how likely the model is to generate response  $r_k$  given test query  $\tilde{q}$ , few-shot examples  $\mathbf{X}$ , and (2) a *backward consistency score*, measuring how  $r_k$  is effective to correctly answer the queries  $q_i$  in  $\mathbf{X}$ .

**Problem setup.** Let us assume that we have an estimation model P which can measure the likelihood  $P(\mathbf{y}_k) = P(r_k \mid \widetilde{q})$  of the response  $r_k$  conditioned on the given query  $\widetilde{q}$ . Then, our goal is to select the response  $r_{k^*}$  which yields the highest likelihood if the estimation is accurate:

$$k^* = \arg\max_{k=1,\dots,K} P(\mathbf{y}_k). \tag{3}$$

We note that the likelihood has shown effectiveness to find high-quality reasoning path (Wang & Zhou, 2024). However, selecting based on the estimated  $P(\mathbf{y}_k)$  could be ineffective in practice, as it entirely depends on the estimation model's intrinsic knowledge, which can be limited in unfamiliar or challenging domains. Furthermore, when there is a mismatch between  $\mathcal{M}$  and P, the estimated likelihoods can be unreliable as minor syntactic variations in response can make large deviations. To address this, we propose to reinterpret  $P(\mathbf{y_k})$  with few-shot examples  $\mathbf{X}$ , through Bayes' rule:

$$P(\mathbf{y}_k) = \frac{P(\mathbf{y}_k \mid \mathbf{X}) \cdot P(\mathbf{X})}{P(\mathbf{X} \mid \mathbf{y}_k)}.$$
 (4)

Then, in the log form, this can be decomposed into two intuitive forward and backward scores:

$$\log P(\mathbf{y}_k) = \underbrace{\log P(\mathbf{y}_k \mid \mathbf{X})}_{\text{forward}} - \underbrace{\left(\log P(\mathbf{X} \mid \mathbf{y}_k) - \log P(\mathbf{X})\right)}_{\text{backward}}.$$
 (5)

While Eq. 5 holds mathematically, discrepancies between the left- and right-hand sides can arise in practice due to the limitations of the estimation model. To address this, the core idea of ReFeri is to

<sup>&</sup>lt;sup>2</sup>For the experiments in Section 3, we use pre-trained LLM as the estimation model.

estimate the forward and backward scores separately, as each can be more accurately approximated by the estimation model with the help of few-shot examples. Then, ReFeri combines these two estimated scores to yield the final selection score. Overall algorithm is presented in Algorithm 1.

Forward confidence score. Intuitively,  $\log P(\mathbf{y}_k|\mathbf{X})$  captures the confidence of candidate response  $r_k$  to test query  $\widetilde{q}$ ; this score is high when  $r_k$  well-aligns with the reasoning patterns in the few-shot examples  $\mathbf{X}$ . This forward score has certain advantages over direct estimation of  $P(\mathbf{y}_k)$ , as it allows the estimation to be grounded in the few-shot examples and hence reduces the reliance on its prior knowledge alone. As a result, the forward score provides a more context-aware and robust estimation, especially important in unfamiliar or domain-shifted scenarios. When the estimation model P is equal to generation LLM  $\mathcal{M}$ , the forward score can be freely obtained during generation of  $r_k$ . Formally, under the autoregressive assumption for estimation model P, the forward score is derived as below:

$$S_{\text{Forw}}(\mathbf{y_k}) := \log P(\mathbf{y_k} \mid \mathbf{X}) = \sum_{t=1}^{T} \log P(r_{k,t} \mid \widetilde{q}, \mathbf{X}, r_{k, < t}), \tag{6}$$

where each candidate response is a sequence of T tokens  $r_k = (r_{k,1}, \dots, r_{k,T})$ .

**Backward consistency score.** The backward score,  $\log P(\mathbf{X}|\mathbf{y}_k) - \log P(\mathbf{X})$ , evaluates how well the test query  $\widetilde{q}$  and candidate response  $r_k$  explains the few-shot examples  $\mathbf{X}$ . At a high level, this score serves as a form of consistency check between the response and the given few-shot examples. Under the assumption of mutual independence between few-shot examples, the backward score can also be derived similar to Eq. 6. To better utilize given few-shot examples, we refine the backward term using a leave-one-out strategy (Perez et al., 2021; Izacard et al., 2023) through prompt replacement; namely, we construct new demonstration  $\widetilde{\mathbf{X}}_i$  by replacing i-th example  $\mathbf{x}_i = (q_i, a_i)$  with a pair of test query and candidate response  $(\widetilde{q}, r_k)$ :

$$\widetilde{\mathbf{X}}_i := \mathbf{X}_{-i} \cup \{(\widetilde{q}, r_k)\},\tag{7}$$

where  $\mathbf{X}_{-i}$  denotes the few-shot examples excluding  $\mathbf{x}_i$ . Then, by including  $\widetilde{\mathbf{X}}_i$  during the estimation for  $\mathbf{x}_i$  as additional input context similar to forward term, we define the modified backward score:

$$S_{\text{Back}}(\mathbf{y_k}) := \log P(\mathbf{X} \mid \mathbf{y}_k) - \log P(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^{N} \left( \log P(a_i \mid q_i, \widetilde{\mathbf{X}}_i) - \log P(a_i \mid q_i) \right). \tag{8}$$

This inclusion of remaining examples  $\mathbf{X}_{-i}$  enables more accurate estimation of the likelihood of target example  $\mathbf{x}_i$  by leveraging the in-context learning capability of P (see more discussions in Appendix B.7). Similar to Eq. 6,  $\log P(a_i|q_i,\widetilde{\mathbf{X}}_i)$  and  $\log P(a_i|q_i)$  can be calculated through a token-level decomposition using the autoregressive nature of P.

While the backward consistency score provides a reliable consistency signal, one may concern its computational cost as it grows linearly with the number of few-shot examples. To alleviate this, we propose a lightweight approximation; instead of iterating over all N few-shot examples, we select only the single most relevant example to the test query  $\tilde{q}$ . Specifically, we employ a pre-trained embedding model E to encode both  $\tilde{q}$  and each  $q_i$ , and identify the most relevant example  $\mathbf{x}_{i^{\dagger}}$ :

$$i^{\dagger} = \arg\max_{i=1,\dots,N} \cos\left(E(\widetilde{q}), E(q_i)\right). \tag{9}$$

Then, we define the approximated backward score  $\widetilde{S}_{\text{Back}}$ :

$$\widetilde{S}_{\texttt{Back}}(\mathbf{y_k}) := \log P(a_{i^\dagger} \mid q_{i^\dagger}, \widetilde{\mathbf{X}}_{i^\dagger}) - \log P(a_{i^\dagger} \mid q_{i^\dagger}) \tag{10}$$

**Final score.** By combining forward and backward scores following Eq. 5, we design our main selection score  $S_{\text{Fin}}$  to find the most promising output  $r_{k^*}$  as below:

$$k^{\star} = \arg \max_{k=1,\dots,K} S_{\text{Fin}}(\mathbf{y_k}), \quad S_{\text{Fin}}(\mathbf{y_k}) := S_{\text{Forw}}(\mathbf{y_k}) - \widetilde{S}_{\text{Back}}(\mathbf{y_k}). \tag{11}$$

#### 3 EXPERIMENTS

In this section, we design our experiments to investigate the following questions:

- Is ReFeri effective to select the correct output across various tasks and LLMs? (Table 1)
- Can ReFeri enable test-time scaling without external reward model and training? (Figure 3)
- What is the effect of each component, and how robust is ReFeri? (Tables 2, 3, 4)
- How does the estimation model affect cost and performance of ReFeri? (Figure 4, Table 15)

#### 3.1 Setups

**Datasets.** We evaluate our method on seven benchmarks encompassing diverse reasoning paradigms, including symbolic-numeric, expertise-based, and multi-hop textual reasoning tasks. (1) *MATH500* (Lightman et al., 2024); a 500-problem subset of MATH (Hendrycks et al., 2021b), focused on symbolic manipulation and multi-step mathematical reasoning. (2) *MMLU-pro* (Wang et al., 2024b); 4200 examples, including 300 randomly sampled questions per domain (e.g., physics, law, computer science) extends the original MMLU benchmark (Hendrycks et al., 2021a) by adding reasoning-focused questions and expanding the choice set from four to ten. (3) *HotpotQA* (Yang et al., 2018); 500 samples from (Kim et al., 2024) a multi-hop question-answering benchmark requiring reasoning across multiple Wikipedia paragraphs with annotated supporting facts. (4) *DROP* (Dua et al., 2019); 500 randomly sampled questions from this reading comprehension benchmark, demanding discrete numerical reasoning (e.g., addition, counting, sorting) over paragraphs. (5) GPQA-diamond (Rein et al., 2024) (*GPQA*); 198 graduate-level questions assessing complex reasoning in biology, physics, and chemistry. (6,7) MuSR (Sprague et al., 2024); 256 examples in Object Placement (*MuSR-op*) and 250 examples in Team Allocation (*MuSR-ta*) tasks assessing spatial and relational reasoning.

Notably, prior work (Sprague et al., 2025) has shown that few-shot Chain-of-Thought (CoT) prompting yields significant gains over zero-shot CoT in MuSR, highlighting the role of in-context examples in complex reasoning. As few-shot examples are necessary for some baselines and ReFeri, we collect them following the previous works. MATH500: 5 examples from (Yang et al., 2024) (GPTs), 4 examples from (Lewkowycz et al., 2022) (LLaMA). MMLU-Pro: 5 examples from (Wang et al., 2024b). HotpotQA: 6 examples from (Yao et al., 2023). DROP: 3 examples following (Zhou et al., 2022). GPQA-Diamond: 5 examples from (Rein et al., 2024). MuSR: 3 examples from (Sprague et al., 2025). Complete prompt templates are available in Appendix A.1.

**Baselines.** We compare ReFeri against five widely-used prompt-based methods that require no additional training, with some reflecting different uses of few-shot examples: (1) *Zero-shot CoT* appends a trigger phrase ("Let's think step by step.") to each query without providing exemplars, instead relying on LLM's intrinsic reasoning capabilities. (2) *Few-shot CoT* prepends a fixed set of few examples, enabling LLM to generalize from few in-context demonstrations. (3) *LEAP* (Zhang et al., 2024) improves few-shot prompting by intentionally inducing mistakes on few examples. Then extracting generalizable task-specific principles through self-reflection without human annotations, and prompting the model to apply these principles to unseen questions. Specific prompts for each baseline are in Appendix A.2. (4) *USC* asks LLM to select the best answer from multiple CoT outputs, by following (Chen et al., 2023). (5) *CoT-WP* (Wang & Zhou, 2024) scores each candidate response using token-level probabilities from LLM conditioned on the same few-shot examples. Specifically, the score is a confidence gap between top-1 and top-2 tokens at answer positions.

**Implementation details.** For the experiments, we use (1) gpt-4o-2024-08-06 (GPT-4o) (OpenAI, 2024a), (2) gpt-4o-mini-2024-07-18 (GPT-4o-mini) (OpenAI, 2024b), and (3) LLaMA-3.1-8B-Instruct (LLaMA-3.1-8B) (OpenAI, OpenAI) as target OpenAI, OpenAI,

<sup>&</sup>lt;sup>3</sup>(1) Using the same prompt as GPT results in significantly lower accuracy, and (2) LLaMA-based models provide their own optimized prompt templates (see meta-llama/Llama-3.2-3B-Instruct-evals).

Table 1: **Main results.** Overall performance on seven reasoning benchmarks comparing the proposed **ReFeri** with different baselines not require additional training, under three different state-of-the-art LLMs. The best and second-best scores are highlighted in **bold** and underline, respectively.

Models	Methods	MuSR-ta (Acc.)	MuSR-op (Acc.)	GPQA (Acc.)	MATH500 (Acc.)	DROP (EM / F1)	HotpotQA (EM / F1)	MMLU-PRO (Acc.)	Avg.
3B	Zero-shot CoT	43.0	50.6	21.6	44.2	60.4 / 66.4	15.2 / 21.2	39.8	39.3
	Few-shot CoT	64.8	53.3	24.0	42.9	61.4 / 67.3	19.0 / 25.1	38.7	43.4
-3.	LEAP	69.2	51.6	27.8	42.3	58.2 / 64.1	19.9 / 26.8	37.3	43.8
4A	USC	67.2	52.3	28.8	<u>49.6</u>	<u>69.6</u> / <b>75.8</b>	24.4 / 32.5	<u>45.6</u>	48.2
LLaMA-3.1-8B	CoT-WP	72.4	<u>54.7</u>	29.3	47.8	71.6 / 75.8	25.8 / 33.4	46.0	<u>49.7</u>
I	ReFeri	79.6	57.8	35.4	51.2	69.4 / <u>75.7</u>	<u>25.0 / 33.2</u>	45.1	51.9
-=	Zero-shot CoT	56.2	58.1	43.0	76.4	77.6 / <b>85.6</b>	31.5 / 41.4	63.0	58.0
·ii	Few-shot CoT	77.0	59.4	41.3	75.2	76.8 / 83.1	34.0 / 45.1	63.0	61.0
GPT-40-mini	LEAP	74.4	59.8	43.9	74.5	75.8 / 83.0	34.0 / 45.1	63.2	60.8
7	USC	74.4	60.9	46.0	77.8	76.8 / 83.8	35.0 / 47.2	63.7	62.1
Ĕ	CoT-WP	<u>78.8</u>	56.3	42.4	77.8	76.4 / 82.5	35.8 / 46.7	64.6	61.7
0	ReFeri	82.8	61.3	41.9	77.8	<b>79.2</b> / <u>84.9</u>	36.2 / 48.0	64.9	63.4
	Zero-shot CoT	66.6	61.7	48.8	77.5	75.1 / 85.3	37.6 / 49.9	73.9	63.0
0	Few-shot CoT	87.0	69.7	47.8	75.6	80.6 / 89.2	44.6 / 58.4	73.7	68.4
7	LEAP	87.2	66.8	45.5	75.6	81.5 / 89.8	45.1 / 58.4	74.0	68.0
GPT-40	USC	85.2	<u>71.1</u>	47.0	77.4	82.2 / 90.2	45.6 / 59.7	74.5	69.0
Ö	CoT-WP	88.0	68.8	47.5	78.4	83.4 / <b>91.4</b>	<b>47.2</b> / <u>60.2</u>	74.1	<u>69.6</u>
	ReFeri	90.4	71.9	51.5	<u>77.8</u>	<b>83.6</b> / <u>91.1</u>	<u>47.0</u> / <b>60.7</b>	75.4	71.1

## 3.2 Main results

270

271

272

273274275276277278279

281

284

287

289 290 291

292293

295

296

297

298

299

300

301

302

303

304

305

306

307

308 309

310

311

312

313

314

315

316

317

318

319

320

321

322

323

Table 1 summarizes the experimental results across seven different reasoning benchmarks and three different LLMs. For instance, across all LLMs and benchmarks, ReFeri improves average accuracy by 4.5% over Few-shot CoT, which corresponds to apply random selection instead. Compared to the second-best method, CoT-WP, ReFeri achieves an average improvement of 1.8% across all benchmarks. Notably, CoT-WP relies solely on the forward likelihood of each candidate, while ReFeri combines both forward and backward signals via a Bayes-derived scoring function. This bidirectional formulation allows ReFeri to capture not just the confidence of an answer, but also its consistency with few-shot examples upon the LLM's intrinsic knowledge about the task; consequently, it enables a better selection across various tasks. We note that performance of prompt-based selection, USC, largely varies depending on the task and used LLMs, which reveals the limitation of solely relying on LLM's intrinsic knowledge. In addition, as mentioned in Section 3.1, MuSR is a benchmark where few-shot examples play a critical role (Sprague et al., 2025) and our results also support this with 21.0% average improvement by Few-shot CoT over Zero-shot CoT. Here, we find that ReFeri further enlarges the improvement with the largest gain, outperforming the second-best method by 4.5%. This result shows that ReFeri is particularly effective in new domains where LLM has little prior knowledge and need to heavily rely on a few examples without additional training or reward models.

Next, to assess whether ReFeri scales effectively with the number of candidate outputs similar to the conventional reward-based best-of-K selection, we evaluate performance as the candidate pool grows. Specifically, we test  $K = \{1, 5, 10, 15, 20\}$  candidates on three representative tasks—MATH500, GPQA, and MuSR-ta by using GPT-40-mini as the generation model under Few-shot CoT. Across the three tasks, ReFeri yields consistent improvements as K increases. On MATH500, while the accuracy of random selection decreases as the number of generated samples increases, ReFeri consistently selects higher-quality responses, improving from 75.8% at K=1 to 79.4% at K=20. On GPQA, where ReFeri raises performance from 41.4% to 45.5% as the candidate pool grows. Consistently, the largest gain is observed on MuSR-ta, which saw a sharp jump in accuracy from 75.6% to 86.0%, an improvement of 10.4%. In contrast, CoT-WP and USC exhibit unstable accuracy under the test-time scaling. Their performance even degrades as the number of candidates increases, suggesting that these methods do not capture what is truly plausible among the candidates. Notably, USC demonstrates strong performance on GPQA when K=5, but its accuracy declines as K increases, highlighting sensitivity to the candidate set size. In addition, we observe an inherent ordering bias in USC: selections come from the first two responses regardless of correctness (see Appendix B.1), indicating a limitation of prompt-based approach. Overall, these results confirm that ReFeri scales well with more candidates, demonstrating effectiveness and reliability in practical test-time scaling.

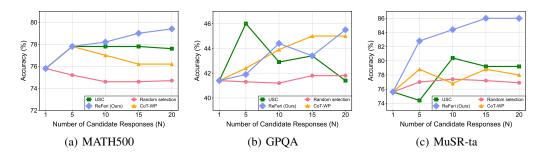


Figure 3: **Test-time scaling with ReFeri.** Accuracy of ReFeri versus other training-free selection methods (Random selection, CoT-WP and USC) on MATH500, GPQA, and MuSR-ta. GPT-4o-mini generate different numbers of candidate responses (K=1,5,10,15,20) using Few-shot CoT.

Table 2: **Ablation study.** Evaluation of scoring variants averaged over three generation LLMs, comparing the contribution of metric term (forward and backward) on MATH500 and GPQA.

	Forw.	Back.	MATH500	GPQA
	X / X	X X ✓	64.6 68.9 64.6	37.7 42.2 34.7
ReFeri (Full) ReFeri	1	1	69.0 68.9	42.7 42.9

Table 3: **Different few-shot examples.** Accuracy across three different choices of few-shot examples on MATH500 (top) and GPQA (bottotm) using GPT-40-mini to generate responses.

Methods	1st	2nd	3rd	Avg.
Few-shot CoT	75.2	74.5	75.0	74.9
ReFeri (Ours)	<b>77.8</b>	<b>79.0</b>	<b>77.8</b>	<b>78.2</b>
Few-shot CoT	41.3	41.5	38.9	40.6
ReFeri (Ours)	<b>41.9</b>	<b>43.4</b>	<b>41.9</b>	<b>42.4</b>

#### 3.3 Additional analyses

In this section, we present additional analyses of ReFeri. We primarily perform experiments with GPT-40-mini on MATH500 and GPQA datasets. More results are presented in Appendix B.

**Ablation study.** To better understand which components contribute to the effectiveness of ReFeri, we conduct an ablation study on each part of the proposed scoring function (Eq. 11), which is grounded in Bayes' rule (Eq. 5). In Table 2, we report performance on MATH500 and GPQA, averaged across the three LLMs used in Table 1. First, one can observe that using the combined score yields better results compared to solely using the forward confidence score (Eq. 6) or the backward consistency score (Eq. 10). This complementary effect comes from their different natures; while the forward score focuses on model-generated response which may contain noise, backward score utilizes given few-shot examples, which are well-curated inputs and ground-truth labels, and thus relatively clean. Next, it is also observed that using the cost-efficient variant of backward term (*ReFeri*) does not compromise the performance compared to *ReFeri* (*Full*) which uses original backward score (Eq. 8), when combined with forward term (more results are in Table 16). This result mitigates concerns regarding the additional computations incurred by original backward term.

Robustness to few-shot examples. ReFeri highly relies on few-shot examples for scoring of both forward and backward scores (Section 2.2). This raises the question of how sensitive the method is to the choice of few-shot examples. To answer this, we conduct a sensitivity study on MATH500 and GPQA using GPT-40-mini, where we use three different few-shot examples with one original and two newly sampled. As shown in Table 3, both Few-shot CoT and ReFeri show variation across these different sets. Nevertheless, ReFeri consistently outperforms Few-shot CoT which corresponds to random selection, and the average gap remains 2.6%. These results indicate that ReFeri remains robust to exemplar choice and is consistently effective, rather than overfitted to specific demonstrations.

Moreover, in practical applications, the clarity of few-shot example might be not always guaranteed. To verify the effectiveness of ReFeri under this scenario, we first synthesize low-quality few-shot examples by converting the original examples via prompting GPT-40-mini to degrade the quality of reasoning in data. The degradation of quality is indeed confirmed through LLM-as-judge framework (results and judgments are in Appendix B.4) As shown in Table 11, ReFeri maintains consistent improvements even under degraded exemplars, indicating that recycled few-shot examples as verification

remains effective without well curated examples; for instance, on MATH500 with LLaMA-3.1-8B, accuracy improves from 39.5% to 47.6% (+8.1), demonstrating ReFeri's robustness.

Additionally, to investigate the impact of various prompt choices, we conduct new experiments with two alternative prompting techniques, following prior work planning and role-playing (Wang et al., 2023a; Kong et al., 2024). Specifically, we assess the robustness of ReFeri by varying prompts during both the generation stage and the verification stage by adapting different prompting styles (orig, plan, and role). For plan and role prompting at the generation, we newly sample five responses similar to Table 1. Table 4 shows the performance on MATH500 and GPQA under

Table 4: **Ablation on generation/evaluation prompts.** Evaluation on MATH500 and GPQA with generation/evaluation prompt variants (orig/plan/role).

		1 1	`	C I		
Gen   Eval		MATH		<b>GPQA</b> Few-shot ReFeri		
		Few-shot	ReFeri	rew-snot	Keren	
	Orig	75.2	77.8	41.3	41.9	
Orig	Plan	75.2	78.0	41.3	42.4	
	Role	75.2	77.8	41.3	41.9	
Plan	Plan	74.6	78.2	42.6	47.5	
Flan	Orig	74.6	78.4	42.6	47.5	
Role	Role	74.5	78.2	43.5	47.5	
Kole	Orig	74.5	78.2	43.5	47.0	

different configurations. A key observation is that verification performance remains highly consistent across different verification prompt styles for a given generation prompt style. For instance, on GPQA, performance for "plan  $\rightarrow$  orig" and "plan  $\rightarrow$  plan" conditions is identical (47.5 vs. 47.5), with similar consistency observed for the "role" condition (47.5 vs. 47.0). This indicates that ReFeri is inherently robust to variations in prompt formatting during the evaluation stage.

However, we also observe that the initial quality of the generated candidate set varies depending on the prompt style. For relatively simple tasks like MATH500, the quality of generated responses is similar across prompts. Conversely, on the more challenging GPQA, prompts offering structured guidance (*e.g.*, plans or roles) tend to generate higher-quality seeds, reflected in slightly higher accuracy. Consequently, ReFeri performs better when the initial candidates are of higher quality.

Estimation models and computational cost. To examine whether ReFeri is robust to the choice of estimation model  $p_{\theta}$ , we evaluate its performance using three LLMs with diverse scales and backbones: LLaMA-3.2-1B-Inst, Qwen-2.5-7B-Inst, and LLaMA-3.1-70B-Inst. The generation model is fixed (either GPT-4o-mini, GPT-4o, or LLaMA-3.1-8B), and we apply each estimation models to two tasks on MATH500 and GPQA. The average accuracy of three generation LLMs is presented in Figure 4 (Full results are in Appendix B.6). Here, ReFeri consistently improves Few-shot CoT across all settings, with an average gain of 4.9% on MATH500 and 5.1% on GPQA. Notably, the smallest model (LLaMA-3.2-1B) performs competitively, and even achieves competitive

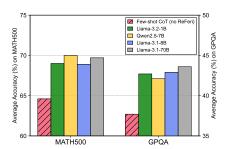


Figure 4: **Estimation model.** Each bar shows the average accuracy of three generation LLMs on MATH500 and GPQA.

performance on MATH500. We attribute this to the relative simplicity of MATH benchmark, as recent small LLMs often exhibit reasonable performance; hence, they can make reliable likelihood estimates for selection. In contrast, GPQA requires more complex reasoning; therefore, using the large estimation model could be more beneficial. Indeed, LLaMA-3.1-70B achieves the best performance on this case. Despite these task-specific differences, the overall improvements are consistent across all estimation models. This suggests that the effectiveness of ReFeri primarily stems from its validation strategy with few-shot examples, rather than the specific choice of estimation model.

In addition, this consistent effectiveness offers better cost-accuracy trade-off. To show this, we conduct experiments using a small verifier for ReFeri (LLaMA-3.2-1B) and compared it against baselines that rely on a larger model (USC and CoT-WP with LLaMA-3.1-8B). Accuracy and latency per query (seconds per instance on a single GPU with identical configuration) are reported in Table 15. Here, ReFeri with a 1B estimator outperforms the strong 8B CoT-WP baseline while substantially reducing latency. For instance, on MATH500, ReFeri (1B) is approximately 60–65% faster than the CoT-WP (8B) baseline (*e.g.*, 3.0s vs. 8.3s on MATH500). Furthermore, ReFeri exhibits robust performance regardless of estimation model size, whereas baselines often suffer significant degradation when scaled down. This demonstrates that combining ReFeri with a small-scale estimator provides a highly advantageous, delivering robust validation at remarkably low computational cost.

## 4 RELATED WORKS

432

433 434

435

436

437

438

439

440

441

442

443

444

445

446

447

448

449

450 451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468 469 470

471 472

473

474

475

476 477

478

479

480

481

482

483 484 485

**Few-shot in-context learning of LLM.** Few-shot in-context learning (ICL) revealed that LLMs can generalize to unseen tasks with just a handful of input-output demonstrations (Brown et al., 2020). To handle complex reasoning problems, chain-of-thought (CoT) prompting was proposed to append intermediate steps to the few-shot examples, leading to substantial gains in tasks such as arithmetic, commonsense reasoning, and symbolic manipulation (Wei et al., 2022; Fu et al., 2023; Jin et al., 2024). To further enhance ICL, various strategies have been developed to retrieve better examples using semantic similarity or entropy-based selection (Wu et al., 2023; Peng et al., 2024). However, some studies have shown that few-shot ICL does not always guarantee improvements. For instance, label shuffling or format changes can often leave performance unaffected (Min et al., 2022), and the performance gap between zero-shot and few-shot CoT is narrowing in several benchmarks as instruction tuning becomes more effective (Sprague et al., 2025). In particular, recent LLMs such as DeepSeek-R1, which are trained with reinforcement learning-based reasoning steps, sometimes even show performance degradation when few-shot CoT examples are added (Guo et al., 2025). Nonetheless, carefully selected demonstrations are still effective (Huang et al., 2024). For example, (Ge et al., 2025) show that few-shot examples can reduce overconfidence in multi-step reasoning, and (Yan et al., 2025) show that they help mitigate hallucinations and memory-based mistakes in complex tasks. These observations motivate us to go beyond using few-shot examples solely for generation, and recycling them to evaluate multiple LLM responses and to select the most promising one.

**Selection of diverse LLM outputs.** Due to the probabilistic nature of LLM decoding, LLM can provide diverse outputs for a single input, each reflecting different reasoning paths (Kadavath et al., 2022; Wang & Zhou, 2024; Qiu & Miikkulainen, 2024). To handle this variability, self-consistency (Wang et al., 2023b) samples K independent reasoning paths and selects the majority answer to improve accuracy. However, it assumes that the model produces a single, well-formatted answer, and this assumption is often violated in open-ended tasks such as summarization or free-form dialogue (Stiennon et al., 2020; Salemi et al., 2024). Alternatively, recent Best-of-N approaches aim to directly select the best output among candidates, often using external verification models. For instance, Outcome Reward Models (ORMs) grade final outputs (Cobbe et al., 2021; Uesato et al., 2022), while Process Reward Models (PRMs) assess intermediate reasoning steps to provide finer supervision (Lightman et al., 2024; Wang et al., 2024a). Despite their successes, these models require large-scale, task-specific annotations or domain-specific checkers, limiting their scalability to new domains or unseen tasks. To eliminate the need for external verification models, prompting-based methods such as LLM-as-Judge ask LLM to evaluate its own outputs (Chen et al., 2023; Zheng et al., 2023). However, their effectiveness heavily depends on the model's prior knowledge in the target domain. When this knowledge is lacking, these methods require additional fine-tuning with curated evaluation datasets for sufficient performance, which reintroduces the need for supervision (Yuan et al., 2024; Mahan et al., 2024; Zhang et al., 2025). In contrast, ReFeri is training-free and task-agnostic, offering a more scalable and generalizable approach by recycling a few-shot examples for verification.

#### 5 Conclusion

We propose ReFeri, a training-free framework to find promising LLM output by reusing few-shot data not only for generation but also for verification. In experiments, ReFeri performs consistently effective in various LLMs and tasks, demonstrating robustness across few-shot data and prompt variations. It suggests that ReFeri is a practical way to find the reliable LLM output with minimal human involvement, opening future directions to reconsider the broader utility of few-shot examples.

**Limitation and future works.** Since the selection by ReFeri is determined by likelihoods produced by an estimation model, it does not explain why a response is incorrect, unlike PRMs, which offer step-level feedback, or LLM-as-judge, which can easily generate explanations by prompting. However, we believe that ReFeri can potentially provide a certain level of interpretability; for example, we visualize the token-level uncertainty of candidate responses and observe that it reveals potentially untrustworthy tokens (see Appendix D). This kind of token-level consideration not only provides the interpretability but also can improve the effectiveness of ReFeri, suggesting a future direction.

## ETHICS STATEMENT

ReFeri provides a training-free method for selecting promising outputs from LLMs. This makes it particularly valuable in scenarios where labeled data is scarce or where model fine-tuning is impractical such as limited access to data, or applications in emerging domains where predefined labels are unavailable. In addition, ReFeri reduces the barrier to deploying LLMs in real-world settings without additional supervision. This may contribute to broader and more efficient adoption of LLMs in resource-constrained environments. All datasets used are public and widely adopted.

## REPRODUCIBILITY STATEMENT

For reproducibility, we provide detailed prompts, datasets, and experimental setups in Appendix A. In Section 3 and Appendix B, we report extensive experiments that demonstrate the robustness of our approach. In addition, we will release our code to ensure transparency and facilitate further research.

#### REFERENCES

- Pranjal Aggarwal, Aman Madaan, Yiming Yang, et al. Let's sample step by step: Adaptive-consistency for efficient reasoning and coding with llms. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- Anthropic. Claude 3.5 sonnet. https://www.anthropic.com/news/claude-3-5-sonnet, 2024.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. Universal self-consistency for large language model generation. *arXiv preprint arXiv:2311.17311*, 2023.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 2019.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. arXiv preprint arXiv:2407.21783, 2024.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. Complexity-based prompting for multi-step reasoning. In *International Conference on Learning Representations (ICLR)*, 2023.
- Yuyao Ge, Shenghua Liu, Yiwei Wang, Lingrui Mei, Lizhe Chen, Baolong Bi, and Xueqi Cheng. Innate reasoning is not enough: In-context learning enhances reasoning large language models with less overthinking. *arXiv preprint arXiv:2503.19602*, 2025.
- Lin Gui, Cristina Gârbacea, and Victor Veitch. Bonbon alignment for large language models and the sweetness of best-of-n sampling. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2024.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations (ICLR)*, 2021a.
  - Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021b.
  - Xijie Huang, Li Lyna Zhang, Kwang-Ting Cheng, Fan Yang, and Mao Yang. Fewer is more: Boosting math reasoning with reinforced context pruning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
  - Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. Atlas: Few-shot learning with retrieval augmented language models. *Journal of Machine Learning Research*, 24(251):1–43, 2023.
  - Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and Mengnan Du. The impact of reasoning step length on large language models. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
  - Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.
  - Jaehyung Kim, Jaehyun Nam, Sangwoo Mo, Jongjin Park, Sang-Woo Lee, Minjoon Seo, Jung-Woo Ha, and Jinwoo Shin. Sure: Summarizing retrievals using answer candidates for open-domain qa of Ilms. In *International Conference on Learning Representations (ICLR)*, 2024.
  - Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems* (NeurIPS), 2022.
  - Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. Better zero-shot reasoning with role-play prompting. 2024.
  - Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative reasoning problems with language models. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2022.
  - Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations (ICLR)*, 2024.
  - Dakota Mahan, Duy Van Phung, Rafael Rafailov, Chase Blagden, Nathan Lile, Louis Castricato, Jan-Philipp Fränken, Chelsea Finn, and Alon Albalak. Generative reward models. *arXiv preprint arXiv:2410.12832*, 2024.
  - Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.
  - OpenAI. Hello gpt-4o. https://openai.com/index/hello-gpt-4o/, 2024a.
  - OpenAI. Gpt-4o mini: advancing cost-efficient intelligence. https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/, 2024b.
  - OpenAI. Learning to reason with llms. https://openai.com/index/learning-to-reason-with-llms/, 2024c.
  - Keqin Peng, Liang Ding, Yancheng Yuan, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. Revisiting demonstration selection strategies in in-context learning. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.

- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. True few-shot learning with language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
  - Xin Qiu and Risto Miikkulainen. Semantic density: Uncertainty quantification for large language models through confidence measurement in semantic space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
  - Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching language model agents how to self-improve. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
  - David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024.
  - Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
  - Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling Ilm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024.
  - Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. Musr: Testing the limits of chain-of-thought with multistep soft reasoning, 2024. In *International Conference on Learning Representations (ICLR)*, 2024.
  - Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. In *International Conference on Learning Representations (ICLR)*, 2025.
  - Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
  - Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
  - Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*, 2022.
  - Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023a.
  - Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024a.
  - Xuezhi Wang and Denny Zhou. Chain-of-thought reasoning without prompting. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
  - Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *International Conference on Learning Representations (ICLR)*, 2023b.
  - Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024b.

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- Zhiyong Wu, Yaoxiang Wang, Jiacheng Ye, and Lingpeng Kong. Self-adaptive in-context learning: An information compression perspective for in-context example selection and ordering. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- Kai Yan, Yufei Xu, Zhengyin Du, Xuesong Yao, Zheyu Wang, Xiaowen Guo, and Jiecao Chen. Recitation over reasoning: How cutting-edge language models can fail on elementary school-level reasoning problems? *arXiv preprint arXiv:2504.00509*, 2025.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W Cohen, Ruslan Salakhutdinov, and Christopher D Manning. Hotpotqa: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models, 2023. In *International Conference on Learning Representations (ICLR)*, 2023.
- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. Self-rewarding language models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.
- Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction. In *International Conference on Learning Representations (ICLR)*, 2025.
- Tianjun Zhang, Aman Madaan, Luyu Gao, Steven Zheng, Swaroop Mishra, Yiming Yang, Niket Tandon, and Uri Alon. In-context principle learning from mistakes. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. In *International Conference on Learning Representations (ICLR)*, 2022.

## A MORE DETAILS OF EXPERIMENTAL SETUPS

This section covers more details about the experiments from Section 3.

#### A.1 DATASETS

This subsection provides more information about the dataset and the few-shot examples we used.

- MATH500. The MATH benchmark (Hendrycks et al., 2021b) consists of 12,500 LaTeX-formatted competition-level math problems, with topics ranging from algebra and geometry to number theory. Each problem includes a step-by-step solution and expects the model to generate a boxed final answer (e.g., an integer or simplified expression). We use MATH500, a 500-question subset introduced in (Lightman et al., 2024), uniformly sampled from the test split to preserve subject and difficulty distribution. For few-shot examples, we follow (Yang et al., 2024)<sup>4</sup> for GPT-based models and (Lewkowycz et al., 2022)<sup>5</sup> for LLaMA-based models. The reason for this choice is based on our empirical observation: Simply adding "Please think step by step and put your final answer within \boxed{} \partial \text{" as done in GPT-style few-shot prompts led to a significant drop in accuracy. Namely, LLaMA-based models require prompt formats that are aligned with their own instructions and are sensitive to deviations from the learned template. This benchmark evaluates symbolic reasoning ability in mathematical domains.
- MMLU-Pro. MMLU-Pro (Wang et al., 2024b) is an extension of the original MMLU benchmark (Hendrycks et al., 2021a), which evaluates broad knowledge and reasoning over 57 subjects using 14k 4-way multiple-choice questions. MMLU-Pro introduces 12k curated 10-way multiple-choice questions across 14 professional domains, increasing task difficulty and emphasizing complex, multi-step reasoning. Instead of using the full test set, we subsample 300 questions per subject (totaling 4,200) using random seed 42 and we will share the used indices at the code. For few-shot examples, we follow the format used in (Wang et al., 2024b). This benchmark is used to assess domain-specific and robust reasoning performance.
- **GPQA.** GPQA (Rein et al., 2024) is a graduate-level QA benchmark consisting of 448 expertauthored multiple-choice questions in domains such as physics, chemistry, and biology. Designed to be "Google-proof," it focuses on evaluating complex scientific reasoning that cannot be answered through simple retrieval. We evaluate on GPQA-Diamond, a curated subset of 198 especially difficult questions selected by the authors. Few-shot examples are taken directly from the official release (Rein et al., 2024). This task measures deep scientific understanding.
- **DROP.** The DROP benchmark (Dua et al., 2019) contains 96k question-answer pairs requiring discrete reasoning over Wikipedia passages (e.g., numerical operations, counting, or date comparison). Answers may include spans, numbers, or dates. We evaluate on a 500-sample subset randomly selected from the dev set, and we will share the selected indices at the code. We use 3-shot examples from (Zhou et al., 2022) and report both EM and F1 metrics following the official implementation. This benchmark evaluates models' symbolic reasoning grounded in natural language passages.
- HotpotQA. HotpotQA (Yang et al., 2018) consists of 113k multi-hop QA pairs requiring reasoning over multiple Wikipedia documents. The model must retrieve at least two relevant passages and combine facts to answer each question. We follow the (Kim et al., 2024), which uses 500 samples from the dev set. Few-shot examples are taken from (Yao et al., 2023). This task tests compositional reasoning and the ability to aggregate distributed information across documents.
- MuSR. MuSR (Sprague et al., 2024) is a benchmark for multi-step reasoning over long-form narratives (800–1000 words), constructed via neuro-symbolic generation to embed logical dependencies into natural language. It includes structured tasks such as TeamAllocation (constraint-based planning) and ObjectPlacement (spatial consistency reasoning). We evaluate on the 256 TeamAllocation and 250 ObjectPlacement examples from the official release (Sprague et al., 2024), using 3-shot prompts tailored to each task (Sprague et al., 2025). MuSR requires understanding of narrative flow, contextual logic, and physical feasibility. As demonstrated in (Sprague et al., 2025), ICL plays a critical role in model performance on MuSR, and demonstrates the effectiveness of ReFeri.

<sup>4</sup>https://github.com/QwenLM/Qwen2.5-Math

https://huggingface.co/datasets/meta-llama/Llama-3.2-3B-Instruct-evals

#### A.2 BASELINES

Here, we provide the template used for our baseline, using MATH500 as a representative task among multiple benchmarks. (see list 1–7).

#### A.3 IMPLEMENTATION

This section provides the detailed information needed to implement the main experiment.

**Resource details.** To avoid out-of-memory, we used two NVIDIA H100 GPUs for evaluation with the LLaMA-3.1-70B-Instruct model. All other experiments were performed on a single A6000 GPU.

**Response generation.** We use lm-eval-harness<sup>6</sup> to generate responses from LLaMA-based models, with temperature set to 1.0 and 5 responses per input. The prompt was written in chat template format using vllm.<sup>7</sup> For GPT-family models, we use the official OpenAI API to generate completions under the same sampling configuration. The remaining settings follow the GPT API default settings. During evaluation, we report the average score across the 5 generations. All evaluations are conducted using our custom evaluation scripts to ensure consistent scoring and formatting across models.

**Algorithm of ReFeri.** In algorithm 1, we present the formal algorithm for ReFeri. We generate multiple candidate responses  $\{r_1, ..., r_K\}$  for each test query using Few-shot CoT, as it exhibit the better quality on the average (see Table 1).

## Algorithm 1 ReFeri algorithm

```
Input: estimation model P, embedding model E, test-query \widetilde{q}, K candidate responses \{r_1, \ldots, r_K\}, N few-shot examples \mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N, replaced prompt \widetilde{\mathbf{X}}_i
```

```
\begin{split} i^\dagger &= \arg\max_{i=1,\ldots,N} \cos\left(E(\widetilde{q}), E(q_i)\right) \\ & \textbf{for } k = 1 \textbf{ to } K \textbf{ do} \\ & S_{\texttt{Forw}} \leftarrow \texttt{Compute forward score with } r_k \text{ as label, using } P \text{ and context } (\widetilde{q}, \mathbf{X}) \text{ (Eq. 6)} \\ & \texttt{Construct } \widetilde{\mathbf{X}}_{i^\dagger} \leftarrow \texttt{using a leave-one-out strategy (Eq. 7)} \\ & S_{\texttt{Back}} \leftarrow \texttt{backward score with } a_i \text{ as label, using } P \text{ and } \widetilde{\mathbf{X}}_{i^\dagger} \text{ (Eq. 10)} \\ & S_{\texttt{Fin}} \leftarrow S_{\texttt{Forw}} - S_{\texttt{Back}} \text{ (Eq. 11)} \\ & S_k \leftarrow S_{\texttt{Fin}} \\ & \textbf{end for} \\ & r_{k^*} \leftarrow \arg\max_k S_k \text{ (Eq. 3)} \\ & \textbf{return } r_{k^*} \end{split}
```

## B More Quantitative Results

### B.1 ADDITIONAL COMPARISON WITH FEW-SHOT PROMPTING-BASED SELECTION METHODS

Among the multiple answer selection methods, the simplest and most accessible approach (e.g., learning overhead, domain specificity, etc.) is arguably LLM-as-Judge Chen et al. (2023); Zheng et al. (2023). It uses the LLM itself to score and select answers via in-context learning without any additional training or external verifiers. In particular, the addition of few-shot examples to LLM-as-Judge might be most closely aligned with the core motivation of ReFeri, which is to use demonstrations not only for generation but also for validation. Therefore, in this section, we compare ReFeri and (1) the original *USC* (Chen et al., 2023), (2) *USC* with few-shot (our adaptation), and (3) *LLM-as-Judge with few-shot* created with our optimized prompt (see list 8 and 9).

As shown in Table 5, ReFeri consistently achieves the best or second-best accuracy across all LLMs and benchmarks. Interestingly, we observe that adding few-shot demonstrations to USC often degrades performance (*e.g.*, on GPQA and DROP with GPT-40-mini and LLaMA-3.1-8B), which is likely due to the sensitivity of LLMs to prompt format and positional bias of the responses.

<sup>6</sup>https://github.com/EleutherAI/lm-evaluation-harness

https://huggingface.co/datasets/meta-llama/Llama-3.1-8B-Instruct-evals

Table 5: **Comparison with prompting-based selection.** Overall performance on seven reasoning benchmarks comparing the proposed ReFeri with different prompting-based baselines not require additional training, under three different state-of-the-art LLMs.

Models	Methods	MuSR-ta	MuSR-op	GPQA	MATH500	DROP	HotpotQA	MMLU-PRO	Avg.
	USC	67.2	52.3	28.8	49.6	69.6 / 75.8	24.4 / 32.5	45.6	48.2
Llama-3.1-8B-It	USC-w/Fewshot	70.0	53.9	28.3	47.8	69.0 / 75.3	<b>25.2</b> / 32.3	<u>45.1</u>	48.5
Liama-3.1-8B-II	LLM-as-Judge	65.2	55.1	21.2	46.0	67.7 / 74.0	$23.4 / \overline{31.2}$	44.1	46.1
	ReFeri	79.6	57.8	35.4	51.2	<u>69.4 / 75.7</u>	<u>25.0</u> / <b>33.2</b>	<u>45.1</u>	51.9
	USC	74.4	60.9	46.0	77.8	76.8 / 83.8	35.0 / 47.2	63.7	62.1
GPT-40-mini	USC-w/Fewshot	76.4	63.3	39.9	78.2	77.2 / 84.0	34.8 / 46.6	63.2	61.9
GP 1-40-IIIIII	LLM-as-Judge	75.6	60.6	34.3	77.0	77.4 / 84.4	35.0 / 46.7	63.3	60.5
	ReFeri	82.8	<u>61.3</u>	<u>41.9</u>	<u>77.8</u>	79.2 / 84.9	36.2 / 48.0	64.9	63.4
	USC	85.2	71.1	47.0	77.4	82.2 / 90.2	<u>45.6</u> / 59.7	74.5	69.0
GPT-40	USC-w/Fewshot	88.8	69.1	46.0	77.4	82.0 / 89.9	45.4 / 60.1	74.1	69.0
	LLM-as-Judge	86.0	68.0	46.5	77.8	82.8 / 91.0	45.6 / 59.8	73.3	68.6
	ReFeri	90.4	71.9	51.5	77.8	83.6 / 91.1	47.0 / 60.7	75.4	71.1

Table 6: Response selection distribution per task (GPT-40-mini).

Task	Method	#1	#2	#3	#4	#5	Fail (-1)
	USC	90.2	3.8	1.0	2.8	2.2	0.0
MATH500	USC-w/ Fewshot	81.0	12.6	0.6	3.2	2.6	0.0
	LLM-as-Judge	51.4	2.2	8.0	12.2	25.4	0.8
	USC	34.2	19.1	7.2	24.0	15.5	0.0
MMLU-Pro	USC-w/ Fewshot	18.8	31.0	7.8	22.7	19.7	0.0
	LLM-as-Judge	22.1	8.0	7.8	13.5	48.3	1.2
	USC	21.7	15.2	13.1	23.7	26.3	0.0
GPQA	USC-w/ Fewshot	19.7	17.7	10.6	23.7	28.3	0.0
	LLM-as-Judge	30.8	9.1	10.1	7.6	42.4	0.0
	USC	73.8	21.8	1.8	1.2	1.4	0.0
DROP	USC-w/ Fewshot	78.2	16.6	3.0	1.0	1.2	0.0
	LLM-as-Judge	68.8	9.6	4.6	4.4	12.4	0.2
	USC	77.0	13.4	2.8	5.4	1.4	0.0
HotpotQA	USC-w/ Fewshot	68.6	20.6	2.8	6.0	2.0	0.0
	LLM-as-Judge	65.0	15.0	4.8	6.0	9.2	0.0
	USC	51.6	18.0	11.7	7.4	11.3	0.0
MuSR-op	USC-w/ Fewshot	36.7	40.2	10.2	6.2	6.6	0.0
	LLM-as-Judge	26.9	21.5	15.2	20.7	15.2	0.4
	USC	34.0	3.2	0.8	9.6	9.6	42.8
MuSR-ta	USC-w/ Fewshot	46.4	35.2	2.4	3.6	12.4	0.0
	LLM-as-Judge	27.6	2.8	0.4	1.2	50.4	17.6

Notably, we observe that both prompt-based selection methods, USC and LLM-as-Judge, are highly sensitive to the order of candidate responses. In our experiments, USC frequently selections were made from the first two responses regardless of correctness; on multiple choice question tasks this pattern is less extreme but skew toward early positions is still visible. Moreover, since USC requires explicit answer extraction, tasks such as MuSR-ta revealed many failure cases (e.g., over 40% failures in Table 6), further highlighting its fragility. This highlights a critical weakness in prompt-based selection: the output is often determined more by position than content. Based on these observations, we conducted additional experiments where we randomly rearranged the order of candidate responses. Indeed, we observed this issue in Table 7; on GPQA, for example, USC's accuracy varied notably across different permutations (e.g.,  $46.0 \rightarrow 41.9$ ), demonstrating its sensitivity to presentation order. In contrast, our approach mitigates such ordering artifacts by decoupling few-shot demonstrations from the selection prompt and using them only for scoring. Furthermore, LLM-as-Judge does not perform reliably on more complex tasks (e.g., GPQA showing a noticeable accuracy degradation compared to other methods). These results emphasize that naively incorporating a few examples into prompts does not guarantee consistent gains, and that ReFeri is more robust and scalable.

Table 7: **Evaluation of USC ordering with GPT-4o-mini.** Two random permutations (*perm-A*, *perm-B*) of the candidate order versus the original order.

Methods	MuSR-ta	MuSR-op	GPQA	MATH500	DROP	HotpotQA	MMLU-PRO	Avg.
USC (perm-A)	75.6	59.8	41.9	78.0	76.6 / 83.5	35.6 / 47.2	63.2	61.5
USC (perm-B)	77.2	56.6	45.0	77.2	76.6 / 83.4	35.2 / 46.9	63.6	61.6
USC (default)	74.4	60.9	46.0	77.8	76.8 / 83.8	35.0 / 47.2	63.7	62.1

Table 8: Performance comparison between Zero-shot and ReFeri under zero-shot setting.

Models	Methods	MATH500	GPQA	MuSR-ta
GPT-4o-mini	Zero-shot	76.4	43.0	56.2
	ReFeri	78.2	43.9	58.8
GPT-40	Zero-shot	77.5	48.8	66.6
	ReFeri	80.8	54.0	69.6
Llama-3.1-8B-It	Zero-shot	44.2	21.6	39.6
	ReFeri	50.8	24.2	41.2

Finally, we note that the application of prompt-based approach could be limited due to inherent input context-window length.

## B.2 APPLICATION REFERI TO ZERO-SHOT RESPONSE

As shown in Table 1, Zero-shot CoT often achieves higher accuracy than Few-shot CoT, reflecting the intrinsic knowledge of the model. However, as described in Eq. 3, ReFeri is also applicable to selecting reasoning paths of Zero-shot CoT, although we primarily apply it to Few-shot CoT since it usually yields better reasoning paths (Table 1). With the experiments in Table 8, we verify that applying ReFeri to Zero-shot CoT yields improvements. These results further suggest that the few-shot exemplars in ReFeri mainly function as a post-hoc validation pipeline, rather than as generation guidance as in conventional Few-shot CoT. Also, this effectiveness of ReFeri under decoupling between generation and selection suggests a robust alternative to conventional few-shot prompting strategies, particularly in settings where few-shot examples are ineffective with LLMs.

#### B.3 MULTIPLE RUN

Table 9: Multiple runs of GPT-40-mini across methods (Few-shot, CoT-WP, ReFeri).

	Methods	MATH500	GPQA	MuSR-ta	Avg
	Few-shot	74.1	41.0	77.8	64.3
Run 1	CoT-WP	76.0	43.9	78.4	66.1
	ReFeri	77.2	43.4	82.8	67.8
	Few-shot	74.6	42.9	77.0	64.8
Run 2	CoT-WP	76.4	40.4	82.8	66.5
	ReFeri	78.4	43.9	82.8	68.4
	Few-shot	74.9	41.8	75.8	64.2
Run 3	CoT-WP	76.6	45.0	80.0	67.2
	ReFeri	79.2	44.9	80.8	68.3
	Few-shot	0.11	0.64	0.68	0.07
Variance	CoT-WP	0.06	4.86	5.02	0.21
	ReFeri	0.73	0.39	1.07	0.09

To assess the stability of ReFeri, we conduct multiple runs of GPT-4o-mini and reported in Table 9. By design, our method is deterministic for a given set of candidate responses. However, since the candidates themselves are subject to sampling stochasticity, we perform multiple runs to evaluate the consistency of the performance gains. In these runs, we compare ReFeri with the few-shot and CoT-WP baselines. As the table illustrates, ReFeri consistently outperforms both alternatives

Table 10: Results of GPT-40-mini across different few-shot examples and ReFeri.

Methods	MATH500	GPQA	MuSR-ta
Few-shot 1	75.2	41.3	77.0
ReFeri 1	77.8	41.9	82.8
Few-shot 2	74.5	41.5	57.8
ReFeri 2	79.0	43.4	59.2
Few-shot 3	75.0	38.9	60.1
ReFeri 3	77.8	41.9	62.8

in all runs. Notably, ReFeri also proves to be significantly more stable. Whereas CoT-WP shows considerable variance in tasks like GPQA (4.9) and MuSR-ta (5.0), ReFeri maintains low variance, indicating that its improvements are reliable despite the stochastic nature of candidate generation. This demonstrates the robustness of ReFeri to sampling variations.

#### B.4 More results with different few-shot examples

In addition to the results reported in Section 3.3, we provide extended experiments in Table 10 including MuSR-ta benchmark. Interestingly, MuSR-ta once again highlights the importance of example quality; when synthesizing new data according to (Sprague et al., 2024) to use as few-shot examples, baseline accuracy significantly degrades. Nevertheless, ReFeri demonstrates consistent performance improvements and confirming the robustness.

Table 11: Performance of ReFeri in weak few-shot settings.

Models	Methods	MATH500	GPQA	MuSR-ta
GPT-4o-mini	Few-shot	73.7	41.1	57.5
	ReFeri	76.4	43.4	58.0
GPT-40	Few-shot	75.2	46.1	71.0
	ReFeri	79.0	47.0	75.6
Llama-3.1-8B-It	Few-shot	39.5	26.6	38.7
	ReFeri	47.6	33.3	41.6

We believe that constructing accurate few-shot examples is a minimal effort that one should invest to guide LLMs (even humans) toward a proper behavior for a target task. Still, to evaluate robustness, we conducted experiments using intentionally synthesized "weak few-shot" (Table 11) by GPT-40-mini. Even under these weaker conditions, ReFeri continues to improve performance relative to the Few-shot CoT, confirming that verification remains effective even with suboptimal examples and can provide meaningful gains in more practical, less curated scenarios.

Table 12: Judgment scores (1–10) by GPT-40 for weak fewshot quality.

Judge by GPT-4o (1–10)	MATH500	GPQA	MuSR-ta
Few-shot	8	8	8
Low quality	3	5	4

To verify the degradation, we asked GPT-40 to evaluate the quality of the original set and the weak example set. The evaluation was conducted in random order, and information about each set was not provided to avoid bias. As shown in Table 12, the weak set consistently received significantly lower scores (3–5 points) compared to the original examples (8 points). This further demonstrates that ReFeri maintains its effectiveness even when the quality of the provided examples is low. See list 10-11 for exact prompts used in the generation and evaluation assessment.

Table 13: Ablation on generation/evaluation prompts.

Gen Prompt	Eval Prompt	MATH500		GPQA		MuSR-ta	
		Few-shot	ReFeri	Few-shot	ReFeri	Few-shot	ReFeri
Orig	Orig	75.2	77.8	41.3	41.9	77.0	82.8
	Plan	75.2	78.0	41.3	42.4	77.0	82.8
	Role	75.2	77.8	41.3	41.9	77.0	82.4
Plan	Plan	74.6	78.2	42.6	47.5	77.0	82.4
	Orig	74.6	78.4	42.6	47.5	77.0	82.4
Role	Role	74.5	78.2	43.5	47.5	75.8	81.6
	Orig	74.5	78.2	43.5	47.0	75.8	81.6

Table 14: Full results with different estimation models across three benchmarks.

(a) MATH500 Estimation   GPT-4o-mini GPT-4o Llama-3,1-8B-It   Avg								
II 22.1D	1 70.0	77.6	51.4					
Llama-3.2-1B	78.0	77.6	51.4	69.0				
Llama-3.1-8B	77.8	77.8	51.2	68.9				
Qwen-2.5-7B	78.8	79.2	52.0	70.0				
Llama-3.1-70B	77.8	77.6	53.6	69.7				
(b) GPQA								
Estimation	GPT-40-mini	GPT-40	Llama-3.1-8B-It	Avg				
Llama-3.2-1B	43.9	50.5	33.8	42.7				
Llama-3.1-8B	41.9	51.5	35.4	42.9				
Qwen-2.5-7B	41.4	50.5	34.3	42.1				
Llama-3.1-70B	42.4	53.5	34.8	43.6				
	(c) MuSR-ta							
Estimation	GPT-40-mini	GPT-40	Llama-3.1-8B-It	Avg				
Llama-3.2-1B	83.2	90.8	80.0	84.7				
Llama-3.1-8B	82.8	90.4	79.6	84.3				
Qwen-2.5-7B	82.0	90.8	81.6	84.8				
Llama-3.1-70B	83.6	91.2	81.6	85.5				

#### MORE RESULTS ON GENERATION/EVALUATION PROMPTS

In addition to the prompt style (see 12, 13) ablation study reported in Section 3.3, Table 13 extends the results to include MuSR-ta. As mentioned above, ReFeri demonstrates stable performance across various combinations of generation and evaluation prompts (orig, plan, role), indicating robustness to changes in prompt style. The accuracy of responses generated by Few-shot CoT varies depending on the generation style, but ReFeri consistently shows improved performance across all configurations.

## B.6 FULL RESULTS WITH DIFFERENT ESTIMATION MODELS

Table 15: Computational cost. Evaluation cost of GPT-4o-mini. Costs are measured in actual processing time (seconds) per instance on a single GPU using the same model configuration.

Size	Methods	MATH500 (Acc. / Time)	GPQA (Acc. / Time)	MuSR-ta (Acc. / Time)	
1B	USC	75.0 / 0.6	44.9 / 0.1	75.6 / 0.7	
	CoT-WP	76.0 / 1.5	43.4 / 2.0	77.6 / 5.0	
	ReFeri	78.0 / 3.0	43.9 / 4.0	83.2 / 8.0	
8B	USC	77.8 / 3.7	46.0 / 3.7	74.4 / 3.9	
	CoT-WP	77.8 / 8.3	42.4 / 11.0	78.8 / 25.6	
	ReFeri	77.8 / 16.6	41.9 / 22.1	82.8 / 41.8	

Table 14 provides full results for all estimation model combinations of MATH500, GPQA and MuSRta. This complements the average performance across different generation LLMs (GPT-4o-mini, GPT-40, and LLaMA3.1-8B) shown in Figure 4. Across all three tasks, ReFeri shows consistent performance gains regardless of the estimation model used, emphasizing its robustness. There are

some model-specific trends; for example, smaller models (LLaMA-3.2-1B) perform competitively on (relatively) simple tasks like MATH500, as discussed in Section 3.2.

Moreover, we provide further results by including MuSR-ta in a cost-accuracy analysis (Table 15), which complements the discussion in Sec. 3.3. On this benchmark, name with 1B estimator achieves 83.2% accuracy while requiring only 8s per query, but clearly outperforms the robust 8B CoT-WP baseline, which achieves 78.8% but consumes more than three times the latency (25.6s). This result illustrates that ReFeri with a smaller estimator can still effectively utilizes a few-shot examples to provide robust validation at a much lower cost, making it particularly attractive for scenarios where latency and resource budgets are critical.

## B.7 ADDITIONAL ABLATION

Table 16: Additional ablation study on GPT-4o-mini

Methods	MuSR-ta (Acc.)	MuSR-op (Acc.)	GPQA (Acc.)	MATH500 (Acc.)	DROP (EM / F1)	HotpotQA (EM / F1)	MMLU-PRO (Acc.)	Avg.
No replace (full) No replace	<b>82.8</b> 82.4	60.2 60.2	42.4 <b>42.9</b>	<b>78.0</b> 77.6	78.4 / 84.2 78.4 / 84.1	<b>36.2 / 48.0</b> 35.8 / 47.6	<b>65.0</b> 64.7	63.3
ReFeri (Full) ReFeri	82.8 82.8	61.3 61.3	42.4 41.9	77.8 77.8	<b>79.6 / 85.3</b> 79.2 / 84.9	35.8 / 47.9 36.2 / 48.0	<b>65.0</b> 64.9	<b>63.5</b> 63.4

Here, we conduct the additional experiments to provide comprehensive ablation study for ReFeri. We first evaluate the *Full* variant (Eq. 8), which generally achieves the strongest results across benchmarks (Table 16). This is expected, as using the complete set of examples provides the most faithful estimate of backward consistency. However, as discussed in Sec. 2.2, the computational overhead increases linearly with the number of few-shot examples, which renders the *Full* variant less appealing for large-scale or resource-constrained scenarios.

To further analyze this trade-off, we examine the effectiveness of the proposed prompt replacement (Eq. 7) for better estimation of backward score. To this end, we consider a simplified variant of our backward score, termed No replace, where each few-shot example  $\mathbf{x}_i = (q_i, a_i)$  is evaluated in a one-shot manner using the test query  $\widetilde{q}$  and the candidate response  $r_k$  as additional context. Specifically, this variant modifies the backward score in Eq. 8 by replacing the leave-one-out prompt  $\widetilde{\mathbf{X}}_i$  with a single pair  $\mathbf{y}_k = (\widetilde{q}, r_k)$ :

$$S_{\text{Back}}'(r_k) := \log P(\mathbf{X} \mid \mathbf{y}_k) - \log P(\mathbf{X}) = \sum_{i=1}^N \left[ \log P(a_i \mid q_i, \widetilde{q}, r_k) - \log P(a_i \mid q_i) \right],$$

We note that, as in our main method, a cost-efficient variant can be obtained by incorporating the  $i^{\dagger}$  selection strategy (Eq. 9), which adaptively chooses the most relevant exemplar to the test query.

$$S'_{\text{Back}}(r_k) := \log P(a_{i^{\dagger}} \mid q_{i^{\dagger}}, \widetilde{q}, r_k) - \log P(a_{i^{\dagger}} \mid q_{i^{\dagger}}),$$

This formulation can be interpreted as the most straight-forward implementation of backward score (see Eq. 5) under the assumption of mutual independence between few-shot examples. As shown in Table 16, the accuracy under *No replace* is consistently less or equal than ReFeri (6 of 7). We attribute this to the fact that using full leave-one-out prompts better reflects the consistency of  $\mathbf{y}_k$  with the original in-context reasoning trajectory. Nonetheless, *No replace* could serve as a practical alternative that trades off a small performance drop with the greater simplicity.

## C USAGE OF AI ASSISTANTS

This paper used AI-based writing aids to improve sentence structure, correct grammar, and improve readability. These tools were applied only to language refinement and did not affect the advancement of technical content, research methodology, or experimental analysis. All scientific ideas, results, and conclusions were conceived and written entirely by researchers. The use of AI aids was limited to editorial purposes and did not impair the originality or intellectual contribution of the work.

## D QUALITATIVE EXAMPLES

In this section, we present qualitative examples to further analyze the proposed ReFeri. For better readability, we only present the examples from MATH500, GPQA, and HotpotQA. All the responses are generated by GPT-40-mini, and we use the ReFeri (*Full*) variant for illustration to provide the clearest comparisons.

## D.1 TOKEN LEVEL ANALYSIS

To better understand how ReFeri identifies high-quality response using given few-shot examples, we perform a token-level analysis of following backward consistency score (Eq. 8).

For a given test query  $\widetilde{q}$ , we divide the candidate responses into correct and incorrect groups using ground-truth labels, and calculate the difference in token level score between the two groups. When the backward scores for the tokens in the few-shot examples exhibit lower score in the correct group compared to incorrect one, the tokens are colored red. In other case, the tokens are colored blue. For visual clarity, we only highlight the top 60% of tokens based on the absolute difference in values. The remaining 40% remain uncolored. This visualization highlights the tokens that contributed the most to plausible candidate answers as determined by the backward consistency score. The value in parentheses is the ratio of tokens highlighted in red to the total number of tokens.

Lower token-level scores indicate higher validity in the model, so tokens highlighted in red can be interpreted as those where backward consistency most effectively distinguishes plausible response. For example, in MATH500 and GPQA, we found that numbers, symbols, final answer formatting, or next reasoning steps were often colored red to support the interpretability of the selection criteria for ReFeri. In HotpotQA can also see that red is dominant for most, but not all, meaningful words.

#### Few-shot examples for MATH500 (0.6455):

```
Test-Query:
Twelve 1 by 1 squares form a rectangle, as shown. What is the total area of the shaded region? [asy] size(4cm); defaultpen(linewidth(0.75));

// Filled portions fill((0, 4)-(0, 0)-(2, 0)-cycle, gray(0.75)); fill((0, 4)-(3, 4)-(3, 0)-cycle, gray(0.75));

// grid int j; for (j = 0; j < 4; ++j) draw((j, 0)-(j, 4)); for (j = 0; j < 5; ++j) draw((0, j)-(3, j));

//diagonals draw((0, 4)-(3, 0)); draw((0, 4)-(2, 0)); [/asy]</pre>
```

```
1134
        Few-shot examples for MATH500 (0.6455):
1135
1136
       Few-shot 1 Query:
       Kevin Kangaroo begins hopping on a number line at 0. He wants to get
1137
       to 1, but he can hop only \frac{1}{3} of the distance. Each hop tires him out so
1138
       that he continues to hop \frac{1}{3} of the remaining distance. How far has he hopped after five hops? Express your answer as a common fraction.
1139
1140
       Few-shot 1 Response:
1141
       Let's think step by step
1142
       Kevin hops $1/3$ of the remaining distance with every hop.
       His first hop takes $1/3$ closer.
1143
        For his second hop, he has $2/3$ left to travel, so he hops forward
        $(2/3)(1/3)$.
1145
       For his third hop, he has (2/3)^2 left to travel, so he hops forward
1146
        $ (2/3) ^2 (1/3) $.
1147
        In general, Kevin hops forward (2/3)^{k-1}(1/3) on his kth hop.
       We want to find how far he has hopped after five hops.
1148
       This is a finite geometric series with first term $1/3$, common ratio
1149
        $2/3$, and five terms.
1150
       Thus, Kevin has hopped \frac{1}{3}\left(1-\left(\frac{2}{3}\right)\right)
        right)^5 \left[1-\frac{2}{3}\right] = \frac{211}{243}
1152
       The answer is \frac{211}{243}
1153
       Few-shot 2 Query:
1154
       What is the area of the region defined by the equation x^2+y^2-7=4y-14x+
1155
       3?
       Few-shot 2 Response:
1156
       Let's think step by step
1157
       We rewrite the equation as x^2 + 14x + y^2 - 4y = 10$ and then
1158
       complete the square,
1159
        resulting in (x+7)^2-49 + (y-2)^2-4=10,
1160
       or (x+7)^2+(y-2)^2=63.
        This is the equation of a circle with center (-7, 2) and radius
1161
        $\sqrt{63},$
1162
        so the area of this region is \pi^2 = \c \{63\pi\}.
1163
       The answer is 63\pi
1164
       Few-shot 3 Query:
1165
       If x^2 + y^2 = 1, what is the largest possible value of |x| + |y|?
       Few-shot 3 Response:
1167
        Let's think step by step
1168
        If (x,y) lies on the circle,
        so does (x,-y), (x,-y), and (-x,-y), (which all give the same
1169
       value of |x| + |y|,
1170
        so we can assume that x \ge 0 and y \ge 0.5Then |x| + |y| = x + y.5
1171
        Squaring, we get (x + y)^2 = x^2 + 2xy + y^2 = 1 + 2xy
1172
       Note that (x - y)^2 \ge 0.$Expanding, we get x^2 - 2xy + y^2 \le 0.$
1173
        so 2xy \le x^2 + y^2 = 1.5 Hence, 1 + 2xy \le 2, which means x + y
        \left(2\right).$Equality occurs when x = y = \frac{1}{\sqrt{2}},$
1174
        so the maximum value of |x| + |y| is \frac{\sqrt{2}}{.}
1175
        The answer is \sqrt{2}
1176
1177
1178
1179
1180
```

```
1188
        Few-shot examples for MATH500 (0.6455):
1189
1190
        Few-shot 4 Query:
        If f(x) = \frac{ax + \overline{b}}{cx + d}, abcd \neq 0 and f(f(x)) = x for all x in the domain of f, what is
1191
        the value of a+d?
1192
       Few-shot 4 Response:
1193
        Let's think step by step
1194
        The condition f(f(x)) means that f is the inverse of itself,
1195
        so its graph is symmetrical about the line y = x.
1196
       With a rational function of this form, we will have two asymptotes:
        a vertical one at x=-d/c if cx+d does not divide ax+b,
1197
        and a horizontal one at y=a/c,
1198
        if we take the limit of f(x) as x goes to \mbox{pm\infty}.
1199
        In order for $f$ to be its own inverse, the intersection of the
1200
        asymptotes must lie on the line $y=x$
1201
        so that it and its asymptotes reflect onto themselves.
       This means that -d/c=a/c,
1202
        and therefore -d=a and a+d=\boxed{0}.
1203
        The answer is 0
1204
       Few-shot 5 Query: Expand (2z^2+5z-6)(3z^3-2z+1).
1205
1206
        Few-shot 5 Response:
1207
        Let's think step by step
1208
        $$\begin{array}{crrrrrrr}
        & & & 3z^3 & & -2z & + 1 & \\
1209
        \times & & & & 2z^2 & +5z & -6 \\
1210
        \cline{1-7}\rule{0pt}{0.17in}
1211
        & & & -18z^3 & & +12z & -6 & \\
1212
        & & +15z^4 & & -10z^2 & +5z & & \\
1213
       + \& 6z^5 \& & -4z^3 \& +2z^2 \& & & \\
1214
        \cline{1-7}\rule{0pt}{0.17in}
        & 6z^5 & +15z^4 & -22z^3 & - 8z^2 &+17z & -6 &
1215
        \end{array} The answer is 6z^5+15z^4-22z^3-8z^2+17z-6
1216
1217
1218
1219
1220
        Few-shot examples for GPQA (0.6970):
1221
1222
        Test-Query:
       A chemist performs two reactions:
1223
1224
        Reaction 1: (E)-oct-4-ene is treated with one equiv. of mCPBA,
1225
        followed by aqueous acid.
1226
1227
        Reaction 2: (Z)-oct-4-ene is treated with one equiv. of mCPBA,
        followed by aqueous acid.
1228
1229
        Both reactions go to 100% completion. The chemist combines the
1230
        products of both reactions, and then runs the product mixture on both a
1231
        standard (achiral) reverse-phase HPLC column and a chiral HPLC column.
1232
       Assuming that the chromatographic resolution is as high as
1233
        theoretically possible in both cases, what does the chemist observe
1234
        in each of these chromatograms?
```

```
1242
       Few-shot examples for GPQA (0.6970):
1243
1244
       Few-shot 1 Query:
       In a given population, 1 out of every 400 people has a cancer caused
1245
       by a completely recessive allele, b. Assuming the population is in
1246
       Hardy-Weinberg equilibrium, which of the following is the expected
1247
       proportion of individuals who carry the b allele but are not expected
1248
       to develop the cancer?
1249
       Few-shot 1 Response:
1250
       The expected proportion of individuals who carry the b allele but
       are not expected to develop the cancer equals to the frequency of
1251
       heterozygous allele in the given population.
1252
       According to the Hardy-Weinberg equation p^2 + 2pq + q^2 = 1, where
1253
       p is the frequency of dominant allele frequency, q is the frequency
1254
       of recessive allele frequency, p^2 is the frequency of the homozygous
1255
       dominant allele, q^2 is the frequency of the recessive allele, and 2pq
       is the frequency of the heterozygous allele.
1256
       Given that q^2=1/400, hence, q=0.05 and p=1-q=0.95.
1257
       The frequency of the heterozygous allele is 2pq=2*0.05*0.95=38/400.
1258
       Few-shot 2 Query:
1259
       A Fe pellet of 0.056~\mathrm{g} is first dissolved in 10~\mathrm{mL} of hydrobromic acid
1260
       HBr (0.1 M). The resulting solution is then titrated by KMnO4 (0.02 M).
       How many equivalence points are there?
1262
       Few-shot 2 Response:
1263
       HBr reacts with Fe to produce Fe2+.
       MnO4- initially reacts with Fe2+ followed by Br-.
1264
       There are two equivalence points at 25 ml and 35 ml.
1265
       In the beaker, the present species are Fe2+ and Br-.
1266
       In a titration involving two analytes, it's essential to identify which
1267
       reaction occurs first.
       Given the redox nature of the titration and the reduction potentials:
1268
       E0 (Br2/Br-) = 1.09 V, E0 (Mn04-/Mn2+) = 1.49 V, and E0 (Fe3+/Fe2+) =
1269
       0.77 V.
1270
       With [Fe2+] determined as 0.1M, two reactions are considered.
1271
       Reaction 1: MnO4- reacts with 5Fe2+ and 8H+ to produce Mn2+, 5Fe3+, and
1272
       4H2O.
1273
       Reaction 2: 2MnO4- reacts with 10Br- and 16H+ to produce 2Mn2+ and 5Br2
       with 8H2O as a byproduct.
       MnO4- first reacts with Fe2+ in a 1:5 ratio, making the first
1275
       equivalence point at 10 ml.
1276
       Once Fe2+ is exhausted, MnO4- reacts with Br- in a 2:10 ratio, adding
1277
       another 25 ml for a total second equivalence point at 35 ml.
1278
       Few-shot 3 Query:
1279
       Consider a quantum mechanical system containing a particle of mass
1280
       m moving in an istropic three dimensional potential of the form
       V(r) = 1/2m\omega^2r^2 corresponding to the acted force obeying Hooke's law.
1281
       Here, \omega is the angular frequency of oscillation and r is the radial
1282
       distance of the particle from the origin in spherical polar coordinate.
1283
       What is the value of energy of the third excited state, and how many
1284
       linearly independent eigenfunctions are possible for the same energy
1285
       eigenvalue?
1286
       Few-shot 3 Response:
       This problem is nothing but the three dimensional simple harmonic
1287
       oscillator (SHO) problem.
1288
       The energy spectrum of three dimensional SHO is $E_n= (n+3/2)\hbar
1289
       \omega\$ where n=0,1,2,3....$.
1290
       For third excited state n=3.
1291
       3+3/2=6/2+3/2=9/2.
       Thus the corresponding energy is $(9/2)\hbar \omega$.
1292
       The degeneracy of the state is q_n=(n+1)(n+2)/2.
1293
       For n=3, degeneracy is (3+1)*(3+2)/2=4*5/2=10.
1294
```

## Few-shot examples for GPQA (0.6970):

#### Few-shot 4 Query:

Your overhear two chemists talking to each other as they leave a synthetic organic chemistry lab. One asks the other "So, how did it go?" The second chemist replies, "Not well - my compounds are on top of each other." What is the second chemist most likely referring to?

#### Few-shot 4 Response:

"On top of each other" commonly refers to two compounds that have similar Rf values on chromatography (a common operation in synthetic chemistry).

Similar Rf values arise for compounds with similar polarities.

#### Few-shot 5 Query:

Mitochondria are semi-autonomous cellular organelles in charge of energy production. They encode for a part of their own translational machinery and respiratory complexes. Mitochondrial function is governed by over a thousand proteins imported from the cell, contributing to processes like the transport of proteins, ribosome biogenesis and translation regulation, respiratory oxidation, metabolism, and apoptotic signaling cascade. Mutations in the code for mitochondrial protein networks can cause numerous diseases in humans that are inherited through generations. Mutations of which of the mitochondrial proteins listed below are least likely to be genetically transmitted from a father to his children?

#### Few-shot 5 Response:

The colleague should know that mitochondria from fathers are rarely if ever, transmitted to their offspring.

Therefore, the protein encoded by the paternal mitochondrial genome will most likely not be passed down the generation.

NADH dehydrogenase 2 is the only one encoded by the mitochondrial genome from the MT-ND2 gene among the listed proteins.

Leigh's syndrome, lactic acidosis, and metabolic diseases are all linked to a mutation in the ND2 gene.

ATP binding cassette subfamily B member 8 (ABCB8) is a chromosome 7 encoded gene; Tu translation elongation factor, mitochondrial is chromosome 16 gene TUFM.

Translocase of inner mitochondrial membrane 17B is chromosome X coded gene TIMM17B.

There is no evidence that it is maternally imprinted; hence, daughters may inherit the father's gene copy in a 50:50 ratio.

## Few-shot examples for HotpotQA (0.7883):

#### Test-Query:

 $\mbox{H-63}$  is a highway that runs near the county seat of what Michigan county?.

## Few-shot 1 Query:

What is the elevation range for the area that the eastern sector of the Colorado orogeny extends into?

#### Few-shot 1 Response:

's think step by step. The eastern sector of Colorado orogeny extends into the High Plains. High Plains rise in elevation from around 1,800 to 7,000 ft, so the answer is 1,800 to 7,000 ft. Answer 1,800 to 7,000 ft

## Few-shot 2 Query:

Musician and satirist Allie Goertz wrote a song about the "The Simpsons" character Milhouse, who Matt Groening named after who?

### 1347 Few-shot 2 Response:

's think step by step. Milhouse was named after U.S. president Richard Nixon, so the answer is Richard Nixon. Answer Richard Nixon

Few-shot examples for HotpotQA (0.7883):

#### Few-shot 3 Query:

Which documentary is about Finnish rock groups, Adam Clayton Powell or The Saimaa Gesture?

## Few-shot 3 Response:

's think step by step. Adam Clayton Powell (film) is a documentary about an African-American politician, not Finnish rock groups. So the documentary about Finnish rock groups must instead be The Saimaa Gesture. Answer The Saimaa Gesture

#### Few-shot 4 Query:

What profession does Nicholas Ray and Elia Kazan have in common? Few-shot 4 Response:

's think step by step. Professions of Nicholas Ray are director, screenwriter, and actor. Professions of Elia Kazan are director, producer, screenwriter, and actor. So profession Nicholas Ray and Elia Kazan have in common is director, screenwriter, and actor. Answer director, screenwriter, actor

#### Few-shot 5 Query:

Which magazine was started first Arthur's Magazine or First for Women? Few-shot 5 Response:

's think step by step. Arthur's Magazine was started in 1844. First for Women was started in 1989. 1844 (Arthur's Magazine) < 1989 (First for Women), so Arthur's Magazine was started first. Answer Arthur's Magazine

#### Few-shot 6 Query:

Were Pavel Urysohn and Leonid Levin known for the same type of work? Few-shot 6 Response:

's think step by step. Pavel Urysohn is a mathematician. Leonid Levin is a mathematician and computer scientist. So Pavel Urysohn and Leonid Levin have the same type of work. Answer Yes

## D.2 RESPONSE LEVEL ANALYSIS

For each response selected by ReFeri, we compute the Eq. 11 for both best and worst response. The value in parentheses is the final selection score. The highest score among all candidates corresponds to the best response, and the lowest score represents the worst response. As shown in below examples, the selected response by ReFeri has more accurate reasoning. For instance, in the example of MATH-500, the best-scoring response anchor their reasoning in a coordinate system and follow clean logic. The worst response, on the other hand, ignore spatial cues, misapply subtraction, and over-complicate with lengthy and internally inconsistent steps. Similarly, in GPQA, the best response reasonably combine the results from the two analyses to arrive at the correct answer. The worst response, on the other hand, appears to rely primarily on the idea that "the heavy branching in (A) is consistent with the splitting observed" without making a clear connection to the analyzed results. Thus, one response integrates the data, while the other reduces it to a vague notion of complexity. In HotpotQA, selected answer are consistent with HotpotQA's multi-hop requirements by making intermediate hops explicit and factually correct, while rejected answer provide unsupported single-hop claims.

```
1404
          Responses for MATH500:
1405
          Test-Query:
1406
         Twelve 1 by 1 squares form a rectangle, as shown. What is the total area of the shaded region?
1407
1408
         size(4cm);defaultpen(linewidth(0.75));
1409
1410
         // Filled portions
1411
          fill((0, 4)-(0, 0)-(2, 0)-cycle, gray(0.75));
1412
          fill((0, 4)-(3, 4)-(3, 0)-cycle, gray(0.75));
1413
1414
         // grid
1415
         int j;
          for (j = 0; j < 4; ++j) draw((j, 0)-(j, 4));
1416
          for (j = 0; j < 5; ++j) draw((0, j)-(3, j));
1417
1418
         //diagonals
1419
         draw((0, 4)-(3, 0)); draw((0, 4)-(2, 0));
1420
          [/asy]
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
```

## Responses for MATH500:

#### Worst Response (-1.736):

To find the total area of the shaded region defined by the triangle at the top left and the rectangle on the right, we proceed as follows:

- 1. Identify the dimensions of the figure: The outline shows a rectangle formed by combining the squares. The rectangle consists of 3 squares in width (3 squares wide) and 4 squares in height (4 squares tall). Therefore, the rectangle has a total area of  $3 \times 4 = 12$  square units.
- 2. Determine the area of the upper triangle: The shaded area in the upper left can be identified as a right triangle, located at the top left corner of the rectangle. The base of this triangle is the height of the rectangle (4 squares), and the height is the width of the rectangle (2 squares). The area of a triangle is computed using the formula:

Area of triangle = 
$$\frac{1}{2} \times \text{base} \times \text{height}$$

Area 
$$=\frac{1}{2}\times2\times4=\frac{1}{2}\times8=4$$
 square units

3. Determine the area of the other shaded region: The other shaded area is the rectangle that extends across the top down to the base of the triangle. The rectangle also covers one of the 1x1 squares fully, specifically, it will cover the upper two squares fully which are not part of the triangle. This gives us a square area of 3 (width)  $\times$  4 (height) – Area of triangle = 12 - 4. Thus, the shaded area for consideration is:

$$12 - 4 = 8$$
 square units

- 4. Calculate the total shaded area: Adding these two areas we've established gives: Area of upper triangle: 4 - Area of square: 8
- However, we must ensure we are not double counting any area. Since the rectangle is inclusive, we need only evaluate the total shaded area: The shaded area includes just the triangle area since it encompasses all squares minus the interfered structured area.
- The area accounted for based on shapes left is 4 described above in figures and does not undermine
- 5. Final answer: Thus, the area of the shaded region is confirmed as:

Responses for MATH500: **Best Response (-0.6992):** To find the total area of the shaded region in the rectangle formed by twelve 1 by 1 squares, let's analyze the problem step by step. 1. Understanding the Shape: The given rectangle consists of three columns and four rows of squares, leading to a total of  $3 \times 4 = 12$  unit squares. 2. Identifying the Shaded Areas: The shaded areas consist of two triangular regions: - The first triangular area is defined by points (0, 4), (0, 0), and (2, 0). The second triangular area is defined by points (0, 4), (3, 4), and (3, 0). 3. Calculating the Area of the First Triangle: - Base = 2 (from (0, 0) to (2, 0)) - Height = 4 (from (0,0) to (0,4)) - Area of the triangle =  $\frac{1}{2} \times \text{Base} \times \text{Height} = \frac{1}{2} \times 2 \times 4 = 4$ . 4. Calculating the Area of the Second Triangle: - Base = 3 (from (3, 0) to (3, 4)) - Height = 4 (from (0, 4) to (0, 4)) - Area of the triangle =  $\frac{1}{2} \times \text{Base} \times \text{Height} = \frac{1}{2} \times 3 \times 4 = 6$ . 5. Adding the Areas: - Total shaded area = Area of first triangle + Area of second triangle = 4 + 6 = 10. Thus, the total area of the shaded region is 10 Responses for GPQA: **Test-Ouerv:** Question: You have prepared an unknown compound. To identify the product, you have used the following characterisation techniques: FTIR and 1H NMR. The FTIR spectrum shows a very broad absorption peak at 3000 wavenumbers. A strong absorption peak is also observed at 1700 wavenumbers. Several peaks were observed in the 1H NMR spectrum, none of which are correspond to vinyl-hydrogens. One of the signals in the 1H NMR is a doublet of triplets of quartets whilst a different signal is a doublet of triplets of triplets. Identify the compound as either, CH3CH2C(H)(CH3)C(H)(CH3)COOH, CH3CH2C(H)(C2H5)C(H)(C2H5)COOH, CH3C(H)(CH3)C(H)(CH3)CH2COOH, or CH3C(H)(C2H5)C(H)(C2H5)CH2COOH 

#### 1566 Responses for GPOA: 1567 1568 Worst Response (-1.344): 1569 Let's analyze the information provided systematically. 1570 1571 1. \*\*FTIR Analysis:\*\* 1572 - A broad absorption peak at 3000 wavenumbers indicates the presence of -OH or -NH groups, 1573 which are characteristic of carboxylic acids due to the hydrogen bonding which generates a broad 1574 1575 - A strong absorption peak at 1700 wavenumbers suggests the presence of a carbonyl group (C=O), 1576 also associated with the carboxylic acid functional group. 2. \*\*1H NMR Analysis:\*\* 1578 - The absence of vinyl hydrogen signals suggests that none of the compounds contain a double 1579 bond (i.e., C=C). 1580 - The signal described as a "doublet of triplets of quartets" indicates complex splitting patterns, 1581 likely resulting from hydrogens on adjacent carbon atoms that are not equivalent due to the molecular structure. - The presence of various integrations and patterns suggests alkyl groups and potentially multiple 1584 types of adjacent hydrogens. 3. \*\*Evaluating the Choices:\*\* 1587 - Each choice represents a unique structure. We need to figure out which matches the spectral data: - (A) CH3C(H)(CH3)C(H)(CH3)CH2COOH has the structure for 5 unique hydrogen environments 1589 and matches possible splitting patterns due to the complexity of the methyl and methylene groups. - (B) CH3CH2C(H)(CH3)C(H)(CH3)COOH also has multiple potential splitting but is slightly 1590 less complex than (A). 1591 - (C) CH3C(H)(C2H5)C(H)(C2H5)CH2COOH has an ethyl group that could create distinctive 1592 NMR peaks but still maintains hydrogen integration complexity. 1593 - (D) CH3CH2C(H)(C2H5)C(H)(C2H5)COOH similarly maintains complexity, but having two 1594 ethyl groups may lead to a different integration pattern. 1595 1596 Given this information, we can further deduce that since the FTIR indicated a carboxylic acid and 1597 the NMR shows no vinyl protons, we can confirm the presence of -COOH in any of the compounds. The heavy branching in (A) is consistent with the splitting observed (doublet of triplets of quartets) due to more complex interactions among various hydrogen types. Considering the complexity and integration in the 1H NMR, we can conclude that the most likely compound based on the spectral data is (A) because of the correct splitting patterns associated with the branching substitutions. 1604 The correct answer is (A) CH3C(H)(CH3)C(H)(CH3)CH2COOH. 1608 1609 1610 1611 1612 1613 1614

#### 1620 Responses for GPQA: 1621 1622 Best Response (-1.134): 1623 To identify the unknown compound based on the provided spectroscopic data, let's analyze the 1624 information step by step. 1625 1626 1. \*\*FTIR Analysis\*\*: 1627 - A broad absorption peak at \*\*3000 wavenumbers\*\* indicates the presence of an -OH group, 1628 typical of carboxylic acids or alcohols. 1629 - A strong absorption peak at \*\*1700 wavenumbers\*\* suggests a carbonyl (C=O) group, common 1630 in carboxylic acids and aldehydes. 1631 2. \*\*1H NMR Analysis\*\*: 1632 - The absence of vinyl hydrogens indicates there are no double bonds present in the compound. 1633 - The \*\*doublet of triplets of quartets\*\* and a \*\*doublet of triplets of triplets\*\* indicate a 1634 relatively complex splitting pattern which typically arises from protons in different environments due to adjacent protons. 1637 3. \*\*Analysis of the compounds\*\*: - The compounds listed can be broken down as follows: 1639 - (A) CH3C(H)(CH3)C(H)(CH3)CH2COOH: A branched structure likely leading to complex 1640 splitting. 1641 - (B) CH3CH2C(H)(CH3)C(H)(CH3)COOH: A linear structure with groups that may contribute to 1642 the complexity in splitting. - (C) CH3C(H)(C2H5)C(H)(C2H5)CH2COOH: Also branched, with ethyl groups that could lead 1643 to intricate splitting. 1644 - (D) CH3CH2C(H)(C2H5)C(H)(C2H5)COOH: A linear structure including ethyl groups, likely 1645 resulting in simpler splitting. 1646 1647 4. \*\*Choosing the Compound\*\*: 1648 - The broad absorption at 3000 cm-1 and significant peaks at 1700 cm-1 underscores that we are 1649 likely dealing with a carboxylic acid, given it shows a carbonyl stretch and alcohol absorbance. 1650 - The complexity of the NMR signals suggests a branched structure as those are usually more 1651 complex due to the adjacent hydrogen interactions. 5. \*\*Final Decision\*\*: 1653 - \*\*Compound (C)\*\*, containing ethyl groups and overall complex splitting, aligns with the described spectroscopic data. The branching and additional substituents in compound (C) can 1655 contribute to complex resonance, further supported by the FTIR properties. 1656 1657 Thus, based on the analysis of FTIR and NMR data, the compound can be most accu-1658 rately identified as: 1659 The correct answer is (C). 1661 1662 1663 1664 Responses for HotpotQA: 1665 **Test-Query**: H-63 is a highway that runs near the county seat of what Michigan county?.

1668 1669

Responses for HotpotQA:

#### 1676 Worst Response (-1.324): 1677

1674

1675

1678

1679 1680

1681 1682 1683

1684

1685

1686

1687

1688

1689 1690

1692

Let's think step by step. The H-63 highway runs in Michigan near the county seat of Baraga County, which is L'Anse. So, H-63 is associated with Baraga County.

Answer Baraga County

#### Best Response (-1.033):

Let's think step by step. H-63 is a highway located in Michigan. To determine the county seat it runs near, we can look at its route. H-63 runs primarily through Chippewa County in the Upper Peninsula of Michigan. The county seat of Chippewa County is Sault Ste. Marie. Therefore, H-63 is a highway that runs near the county seat of Chippewa County.

Answer Chippewa County

#### **Listing 1** Few-shot CoT prompt on MATH500

```
1693
1694
        fili
1695
        {system}
       Please reason step by step, and put your final answer within \boxed{{}}.
1697
        {user}
        Kevin Kangaroo begins hopping on a number line at 0. He wants to get to
1699
        \rightarrow 1, but he can hop only frac{1}{3} of the distance. Each hop tires
1700
        \rightarrow him out so that he continues to hop frac{1}{3} of the remaining
1701
        → distance. How far has he hopped after five hops? Express your answer
1702
           as a common fraction.
1703
       Let's think step by step
1704
       Kevin hops $1/3$ of the remaining distance with every hop.
1705
       His first hop takes $1/3$ closer.
1706
       For his second hop, he has $2/3$ left to travel, so he hops forward
        \leftrightarrow $(2/3)(1/3)$.
1707
       For his third hop, he has (2/3)^2 left to travel, so he hops forward
1708
        \leftrightarrow $ (2/3)^2 (1/3)$.
1709
        In general, Kevin hops forward (2/3)^{k-1}(1/3) on his khop.
1710
       We want to find how far he has hopped after five hops.
1711
       This is a finite geometric series with first term $1/3$, common ratio
        \leftrightarrow $2/3$, and five terms.
1712
       Thus, Kevin has hopped
1713
           \frac{f_{cac}}{f_{cac}}  \frac{1}{3}  left \frac{1-\int eft}{\int f_{cac}}  \frac{2}{3}  right) ^5  right) ^5 
1714
        \{1-\sqrt{2}\} = \mathbf{b} oxed \{\sqrt{211}\} \{243\} \}.
1715
       The answer is \{11\}\{243\}
1716
1717
1718
        Convert the point (0,3) in rectangular coordinates to polar
1719
            coordinates. Enter your answer in the form $(r,\theta),$ where $r >
1720
            0$ and $0 \le \theta < 2 \pi.$
1721
        1.1.1
1722
```

## Listing 2 Zero-shot CoT prompt on MATH500

## **Listing 3** Prompt for USC

```
f'''

I have generated the following responses to the question: Convert the

→ point $(0,3)$ in rectangular coordinates to polar coordinates.

→ Enter your answer in the form $(r,\text{theta}),$ where $r > 0$ and $0 \le

→ \theta < 2 \pi.$

Response 0: {response0}

...

Response 4: {response4}

Evaluate these responses.

Select the most consistent response based on majority consensus.

Start your answer with "The most consistent response is Response X"

→ (without quotes).

!!!
```

## **Listing 4** Prompt for LEAP mistakes

```
f'''
{system}
Please reason step by step, and put your final answer within \boxed{{}}.

{user}
Kevin Kangaroo begins hopping on a number line at 0. He wants to get to

→ 1, but he can hop only $\frac{1}{3}$ of the distance. Each hop tires

→ him out so that he continues to hop $\frac{1}{3}$ of the remaining

→ distance. How far has he hopped after five hops? Express your answer

→ as a common fraction.
```

## Listing 5 Prompt for LEAP low-level principles

```
f'''
Question: {question}
Generated Reasoning: {response}
Generated Answer: {generated_answer}
Correct Reasoning: {correct_reasoning}
Correct Answer: {correct_answer}
Instruction: Conduct a thorough analysis of the generated answer in

ightharpoonup comparison to the correct answer. Also observe how the generated
   reasoning differs from the correct reasoning. Identify any
   discrepancies, misunderstandings, or errors. Provide clear insights,
   principles, or guidelines that can be derived from this analysis to
    improve future responses. We are not focused on this one data point,
\rightarrow but rather on the general principle.
Reasoning: <discuss why the generated answer is wrong>
Insights: <what principle should be looked at carefully to improve the
→ performance in the future>
```

## Listing 6 Prompt for LEAP high-level principles

```
Low-level principles:
{low_level_principles}

Create a list of *unique* and insightful principles to improve future

responses based on the analysis above.

Focus on capturing the essence of the feedback while eliminating

redundancies.

Ensure that each point is clear, concise, and directly derived from the

introspection results.

Create a numbered list of principles. Leave specific details in place.

Limit to at most 8 principles.
```

## **Listing 7** Prompt for LEAP generations

```
1849
1850
        fili
1851
        {system}
       Please reason step by step, and put your final answer within \boxed{{}}.
1852
1853
1854
       Please carefully note the following principles:
1855
1856
       Principles: 1. **Meticulous Verification**: Always verify each step in
        \hookrightarrow algebraic processes to prevent errors that can lead to incorrect
1857
        \hookrightarrow conclusions.
1858
1859
        . . .
1860
1861
       8. **Continuous Learning and Adaptation**: Stay open to learning from
        → mistakes and adapting methods to improve future problem-solving
1862
        \rightarrow approaches.
1863
1864
       Kevin Kangaroo begins hopping on a number line at 0. He wants to get to
1865
        \rightarrow 1, but he can hop only frac{1}{3} of the distance. Each hop tires
1866
        \rightarrow him out so that he continues to hop frac{1}{3} of the remaining
1867
        \hookrightarrow distance. How far has he hopped after five hops? Express your answer
        → as a common fraction.
1868
1869
       Let's think step by step
       Kevin hops $1/3$ of the remaining distance with every hop.
1871
       His first hop takes $1/3$ closer.
1872
        . . .
1873
        Convert the point $(0,3)$ in rectangular coordinates to polar
1874
        \leftrightarrow coordinates. Enter your answer in the form (r, \mathbf{t}) where r >
1875
           0$ and $0 \le \theta < 2 \pi.$
1876
1877
```

## Listing 8 Prompt for USC-w/ Fewshot

```
1905
1906
       Kevin Kangaroo begins hopping on a number line at 0. He wants to get to
1907
        \rightarrow 1, but he can hop only frac{1}{3} of the distance. Each hop tires
1908
        \rightarrow him out so that he continues to hop frac{1}{3} of the remaining
1909
        \hookrightarrow distance. How far has he hopped after five hops? Express your answer
1910
           as a common fraction.
1911
       Let's think step by step
1912
       Kevin hops $1/3$ of the remaining distance with every hop.
1913
       His first hop takes $1/3$ closer.
1914
1915
        I have generated the following responses to the question: Convert the
1916
        \rightarrow point (0,3) in rectangular coordinates to polar coordinates.
1917
        \rightarrow Enter your answer in the form (r, \mathbf{t}), where r > 0 and 0 \le
1918

→ \theta < 2 \pi.$
</p>
1919
1920
       Response 0: {response0}
1921
        . . .
1922
1923
       Response 4: {response4}
1924
1925
       Evaluate these responses.
1926
       Select the most consistent response based on majority consensus.
       Start your answer with "The most consistent response is Response X"
1927
        \hookrightarrow (without quotes).
1928
1929
```

## Listing 9 Prompt for LLM-as-Judge

```
1947
       {system}
1948
       Your job is selecting the most accurate response among multiple
1949
       → candidates. You will receive a question and several candidate
1950
       → answers labeled candidate1, candidate2, etc. Please summarize the
1951

ightarrow debate very briefly and then conclude which single candidate is the
1952
       → most plausible. Output exactly in this format:
       Summary: <brief summary>
1953
       Conclusion: candidate<number>
1954
       Remember to choose only one candidate as the final answer.
1955
1956
       {user}
       Please reason step by step, and put your final answer within \bf b
1957
1958
       The below examples are well-constructed gold question and answer pairs
1959
       \hookrightarrow for the same task.
1960
1961
       Kevin Kangaroo begins hopping on a number line at 0. He wants to get to
1962
        \rightarrow 1, but he can hop only frac{1}{3} of the distance. Each hop tires
       \rightarrow him out so that he continues to hop frac{1}{3} of the remaining
1963
       → distance. How far has he hopped after five hops? Express your answer
1964

→ as a common fraction.

1965
1966
       Let's think step by step
1967
       Kevin hops $1/3$ of the remaining distance with every hop.
       His first hop takes $1/3$ closer.
1968
       . . .
1969
1970
       Now, let's select the most proper answer for the given question
1971
       Question: Convert the point (0,3) in rectangular coordinates to polar
1972
       \rightarrow coordinates. Enter your answer in the form (r, \mathbf{t}), where r >
       \rightarrow 0$ and $0 \le \theta < 2 \pi.$
1973
       candidate1: {response 0}
1974
1975
       candidate5: {response 4}
1976
       1.1.1
1977
```

## **Listing 10** Prompt for generate weak few-shot

```
f'''

"You will receive a QUESTION and its original ANSWER.\n"

"Rewrite ONLY the ANSWER; do NOT alter the QUESTION.\n"

"Treat the original as a 10/10 reference. Produce a deliberately

degraded explanation (target quality 1/10):\n"

"- Keep the final answer tokens EXACT (e.g., '\boxed{...}' or 'The

correct answer is (X)').\n"

"- Keep the original CoT style label if present (e.g., 'Let's think step

by step:' / 'Reasoning:').\n"

"- Make reasoning weak: shallow, vague, incomplete; omit steps, avoid

precise formulas/numbers.\n"

"- Prefer generic phrases over concrete derivations. Lower clarity and

rigor compared to the original.\n"

"OUTPUT FORMAT: Return EXACTLY ONE JSON object and NOTHING ELSE:\n"

'{"answer":"<rewritten weaker answer>"}'
```

## Listing 11 Prompt for judging weak few-shot

1998 1999 2000

20012002

2003

2004

2006 2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

2018

20192020

2021

20312032

2033

2035

2036

203820392040204120422043

20442045

2046

2047

2048204920502051

```
"You are judging FEW-SHOT QUALITY only.\n"
"Compare TWO blocks side-by-side. Assume BLOCK A and BLOCK B are

→ candidate few-shot demonstrations.\n\n"

"Ignore question quality entirely -- the question is context only.\n^{n}"
"What "answer quality" means here: \n"
"- clarity, structure, and coherence of the reasoning.\n"
"- specific steps, concrete numbers/equations when relevant, and

→ justified transitions.\n"

"- a single, clearly marked final answer token format (e.g.,
\rightarrow \"\\boxed{\dots}\" or \"The correct answer is (X)\") if present;\n"
"Instructions:\n"
"- Assign an integer score 1-10 to EACH block (higher = better few-shot

    quality).\n"

"- The evaluation should be comparative: scores must reflect their

→ relative quality.\n"

"- Provide brief notes explaining each score.\n\"
"OUTPUT FORMAT:\n"
"Return exactly ONE JSON object with this schema (and nothing else): \n"
"\"A\":{\"score\":int,\"notes\":string},"
"\"B\":{\"score\":int,\"notes\":string},"
"\"comparative_notes\":string"
```

### Listing 12 Prompt for plan-and-sovle on MATH500

```
f'''

"Let's first understand the problem, extract relevant variables and

→ their corresponding numerals, and make a complete plan. Then, let's

→ carry out the plan, calculate intermediate variables (pay attention

→ to correct numerical calculation and commonsense), solve the problem

→ step by step, and put your final answer within \\boxed{{}}.\n"
```

## Listing 13 Prompt for role-playing on MATH500

```
f'''
"From now on, you are an excellent math teacher and always teach your \hookrightarrow students math problems correctly. And I am one of your students. Put \hookrightarrow your final answer within \\boxed{{}}.\n"
```