

LARGE LANGUAGE MODEL CASCADES WITH MIXTURE OF THOUGHT REPRESENTATIONS FOR COST-EFFICIENT REASONING

Murong Yue¹, Jie Zhao², Min Zhang³, Liang Du², Ziyu Yao¹

¹George Mason University ²Microsoft ³Virginia Tech

¹{myue, ziyuyao}@gmu.edu ²{zhaojie, liang.du}@microsoft.com

³{minzhang23}@vt.edu

ABSTRACT

Large language models (LLMs) such as GPT-4 have exhibited remarkable performance in a variety of tasks, but this strong performance often comes with the high expense of using paid API services. In this paper, we are motivated to study building an LLM cascade to save the cost of using LLMs, particularly for performing reasoning (e.g., mathematical, causal) tasks. Our cascade pipeline follows the intuition that simpler questions can be addressed by a weaker but more affordable LLM, whereas only the challenging questions necessitate the stronger and more expensive LLM. To realize this decision-making, we consider “answer consistency” of the weaker LLM as a signal of the question difficulty and propose several methods for the answer sampling and consistency checking, including one leveraging a mixture of two thought representations (i.e., Chain-of-Thought (Wei et al., 2022) and Program-of-Thought (Chen et al., 2022; Gao et al., 2023)). Through experiments on six reasoning benchmark datasets, with GPT-3.5-turbo and GPT-4 being the weaker and stronger LLMs, respectively, we demonstrate that our proposed LLM cascades can achieve performance comparable to using solely the stronger LLM but require only 40% of its cost. Our codes are available at https://github.com/MurongYue/LLM_MoT_cascade.

1 INTRODUCTION

Large language models (LLMs) such as GPT-4 have exhibited remarkable performance in reasoning tasks (Rae et al., 2021; Lewkowycz et al., 2022; Zhong et al., 2023). Because of the intensive computing resources required for training and hosting the LLMs for inference, many such LLMs are only accessible via paid API services, thus leading to high monetary costs. In this work, we are motivated to study strategies for reducing the costs of using LLMs while not sacrificing task performance, particularly for LLMs’ applications to reasoning tasks.

Different types and versions of LLMs often come with different capabilities and costs. Typically, LLMs with better performance (termed “stronger LLMs”) are more expensive than those with relatively worse overall performance (termed “weaker LLMs”). For example, GPT-4 (OpenAI, 2023) is 30 times more expensive than GPT-3.5-turbo for the output tokens.¹ It thus implies a promising solution to cost-saving. That is, simple questions could be answered by the weaker but more affordable LLM, whereas only the difficult questions need to be tackled by the more expensive, stronger LLM. Drawing inspirations from here, Chen et al. (2023a) explored the idea of “LLM cascades”, where a question is always first answered by a weaker LLM, and then optionally routed to a stronger LLM when the the weaker LLM’s answer is not accepted (Figure 1). To decide this routing, this work suggested fine-tuning a smaller LLM to score each question along with its answer produced by the weaker LLM. While this approach could work for some tasks, in practice, we observed that it did not yield satisfying performance for intricate reasoning tasks. Intuitively, it is very challenging to evaluate the difficulty and the answer correctness of a reasoning question solely based on its literal expression, even with a large enough LLM, since the errors could be nuanced despite the reasoning paths appearing promising (Madaan et al., 2023).

¹<https://openai.com/pricing>.

In this work, we proposed to devise this routing decision-maker from a different angle, i.e., the “answer consistency” of the weaker LLM (Wang et al., 2023). This is inspired by the observation that answers from the weaker LLM tend to be consistent in multiple sampling paths when the question is easy, but inconsistent when the question is hard. To implement this idea, we proposed two types of methods, a vote-based method that examines if the agreement of multiple answer samples on the majority-voted answer surpasses a pre-defined confidence threshold, and a verification-based method that checks if the majority-voted answers sampled from different prompts are consistent. To realize the two methods in reasoning tasks, we further investigated multiple strategies for answer sampling, including sampling from a single set versus two sets of task demonstrations. In particular, we proposed to leverage a “mixture of thought (MoT) representations”, which samples answers from both Chain-of-Thought (Wei et al., 2022, CoT) and Program-of-Thought (Chen et al., 2022; Gao et al., 2023, PoT) prompts, emulating how experts can provide diverse perspectives to the same question. This follows the same spirit of ensembling (Rokach, 2010), but is applied to developing LLM cascades for the first time. By pairing different sampling strategies with the two answer consistency checking methods (i.e., vote and verification), we end up with ten approaches to implementing the LLM cascade.

To evaluate the proposed approaches, we conducted experiments on six reasoning datasets, covering mathematical, symbolic, and causal reasoning tasks, using GPT-3.5-turbo as the weaker LLM and GPT-4 as the stronger one. The experimental results demonstrated the effectiveness of LLM cascades based on answering consistency. That is, different approaches that we proposed for LLM cascades can generally achieve performance comparable to or even better than fully using the stronger LLM, while they require only half or less relative cost to the latter. In particular, our approaches based on a mixture of thought representations achieved comparable task performance with only 40% of the cost of GPT-4. Our results also underscored the effectiveness of sampling answers from diverse prompt settings, such as sampling from different task demonstrations or different thought representations. Our further analysis revealed that different prompt settings can often provide different opinions for the more complex questions while tending to be consistent for easier ones, which allows us to distinguish questions at different difficulty levels more accurately for cascade decision-making. Finally, we also compared our consistency-based approaches with fine-tuned smaller LLMs (Chen et al., 2023a) as well as other variants that make the routing decisions based on the literal expressions of the question and its answer. Our approaches exhibited strong advantages over all of them.

2 LLM CASCADES FOR COST-EFFICIENT REASONING

2.1 OVERVIEW OF LLM CASCADES

We leverage a cascade of LLMs to save the cost of in-context LLM reasoning, as illustrated in Figure 1. Specifically, we assume two LLMs. The weaker LLM (denoted as LLM^w) yields relatively worse performance but is less costly, whereas the stronger LLM (denoted as LLM^s) enjoys better task performance but is more expensive. Given a question Q , the LLM cascade first employs the weaker LLM to obtain an initial answer A^w . This answer, along with other metadata produced by the weaker LLM, will then be fed to a cascade decision maker to decide whether the answer can be accepted as the final one. If the answer is rejected, the stronger LLM should be invoked to provide a more reliable answer A^s . As a consequence, the total cost of answering the question becomes

$$C = C^w + C^d + \mathbb{1}_{reject} C^s, \quad (1)$$

where C^w and C^s indicate the costs from calling the weaker and the stronger LLMs, respectively, C^d denotes any cost involved in the LLM cascade decision-making process, and $\mathbb{1}_{reject} = 1$ holds if and only if the decision maker rejects the answer.

Both LLMs solve the question via few-shot in-context learning, e.g., for the weaker LLM, an answer A^w is produced by sampling from $P_{LLM^w}(A^w | E_1 || E_2 || \dots || E_M || Q)$, where $E_1 || E_2 || \dots || E_M || Q$ denotes a concatenation of M task demonstrations and the input question Q , forming the “prompt input” to the LLM (Brown et al., 2020). As M task examples are used to demonstrate the task,




Figure 1: Illustration of LLM cascade chaining a weaker but cheaper LLM with a stronger but more costly one.

it indicates a “ M -shot in-context learning” of LLM. For reasoning tasks, in practice, a LLM will typically be prompted to elaborate on its reasoning process via “thought representations”, such as Chain-of-Thought (Wei et al., 2022, CoT) and Program-of-Thought (Chen et al., 2022; Gao et al., 2023, PoT), where the reasoning process is described step by step via natural language and programming language, respectively. The answer (e.g., a numerical result of the mathematical calculation) can then be extracted from the texts (for CoT) or obtained by executing the code (for PoT).

2.2 ANSWER CONSISTENCY-BASED CASCADE DECISION-MAKING

The core of our LLM cascade is the decision maker, which takes in the output from the weaker LLM, and then decides whether to route to the stronger LLM or not. An ideal cascade decision maker should call the stronger LLM only when the answer by the weaker LLM is wrong, such that the total cost C can be minimized without degrading the overall task performance (compared with using the stronger LLM all the time). To this end, we propose two methodologies based on the “answer consistency” of the weaker LLM, which we elaborate on below.

Answer Consistency and Sources of Sampling Answer consistency has been found helpful for improving the LLM performance in reasoning tasks (Wang et al., 2023). Instead of greedily decoding one answer for each question, Wang et al. (2023) sampled a diverse set of reasoning paths (or thought processes) and then selected the most consistent answer by marginalizing out the sampled paths. Drawing inspiration from the prior work, we make the following hypothesis: When the weaker LLM samples highly consistent answers for a given question, it reveals a high “confidence” in solving this question and its most consistent answer is likely to be correct; in this case, there is thus no need to invoke the stronger LLM.

To realize this intuition, we generalize from Wang et al. (2023) and consider three sources of sampling consistency:

- **In-distribution sampling:** As Wang et al. (2023), we consider sampling multiple answers given the same prompt input to the weaker LLM. In practice, this can be achieved by setting a non-zero temperature for the weaker LLM.
- **Sampling from different in-context demonstrations:** We further consider sampling answers from two sets of task demonstrations under the same thought representation. For example, to demonstrate the CoT process in mathematical reasoning tasks, Wei et al. (2022) annotated eight math examples as the demonstrations and performed 8-shot in-context learning. We additionally annotated another eight examples as the second set of demonstrations, which allowed us to further diversify the sources of answer sampling.
- **Sampling from different thought representations:** While existing literature typically investigated either CoT or PoT independently, in this work, we propose to leverage the synergy of both thought representations in a single task. We hypothesize that an LLM obtains truly high confidence in its problem-solving, only when it is able to produce a consistent answer agnostic to how the intermediate steps are represented. Therefore, we propose to sample the weaker LLM answers from a “mixture of thought (MoT) representations”, which includes both CoT and PoT prompts.

Below, we introduce our methodologies for consistency checking based on the answer samples.

Method 1: Vote-based decision-making The first method calculates the consistency of the weaker LLM’s answer samples by voting. Formally, for a single prompt, we denote the set of answers produced by the weaker LLM for each question Q as $(A_1^w, A_2^w, \dots, A_K^w)$, where K is the pre-defined number of samples. When sampling from two different prompts, we denote $(A_{11}^w, A_{12}^w, \dots, A_{1K_1}^w)$ and $(A_{21}^w, A_{22}^w, \dots, A_{2K_2}^w)$ as the answer samples produced by each of them, where K_1 and K_2 represent the pre-defined sample size for each prompt setting, respectively. Note that for this method, we do not distinguish answers sampled with a single prompt or multiple prompts (e.g., samples from different prompts have exactly equal weights when voting). The most consistent answer can then be selected as the one that most samples agree with, and this answer will also be regarded as the final answer A^w by the weaker LLM. The decision maker measures the weaker LLM’s consistency via the agreement score

$$s = \frac{\sum_{i=1}^K \mathbb{1}_{A_i^w=A^w}}{K} \text{ or } s = \frac{\sum_{i=1}^{K_1} \mathbb{1}_{A_{1i}^w=A^w} + \sum_{i=1}^{K_2} \mathbb{1}_{A_{2i}^w=A^w}}{K_1 + K_2}. \quad (2)$$

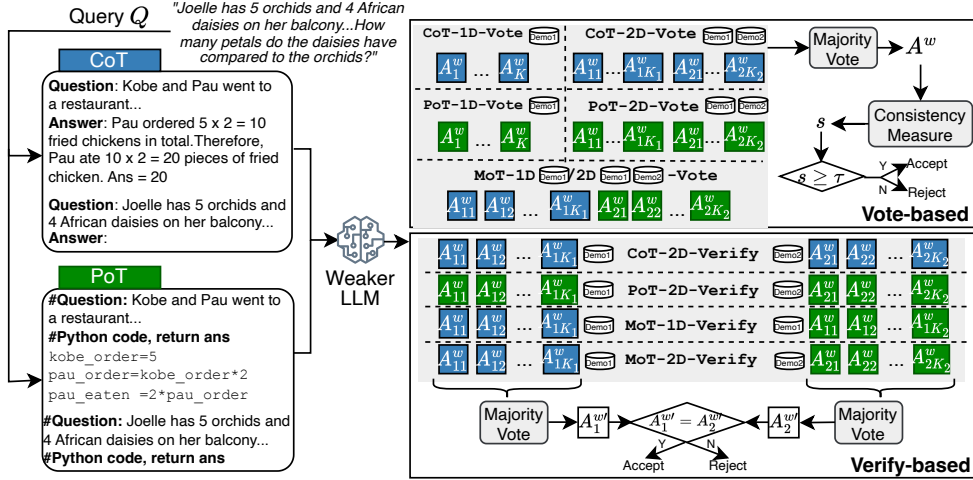


Figure 2: An overview of our approaches (6 vote-based and 4 verification-based). We use ■ to represent the answers from PoT and ■ to represent the answers from CoT. $Demo_i$ is the i -th set of demonstrations.

The larger the s , the more consistent the weaker LLM’s answer samples. In conjunction with a pre-defined threshold value τ , the decision maker accepts the weaker LLM’s most consistent answer A^w when $s \geq \tau$ and rejects it otherwise. As a result, the total cost of answering a question (Eq 1) can vary depending on the threshold.

Method 2: Verification-based decision-making In the case of producing samples from two different prompt settings (i.e., different demonstrations or thought representations), we propose the second method, which compares the most consistent answers produced by each prompt as the answer verification. As previously mentioned, we could obtain two sets of answers from distinct prompts: $(A_{11}^w, A_{12}^w, \dots, A_{1K_1}^w)$ and $(A_{21}^w, A_{22}^w, \dots, A_{2K_2}^w)$. Our method then verifies the most consistent answers within each prompt, denoted as $A_1^{w'}$ and $A_2^{w'}$ respectively, as follows:

$$s = \mathbb{1}_{A_1^{w'} = A_2^{w'}}. \quad (3)$$

Only when s equals 1, i.e., when the two answers are the same, the weaker LLM’s answer will be accepted by the decision maker. In this case, the final answer of the weaker LLM will be the same as the two most consistent answers, i.e., $A^w = A_1^{w'} = A_2^{w'}$.

In comparison, these two methods have different applicable scenarios. The vote-based method is well-suited for scenarios with pre-defined cost constraints. In such cases, we have the flexibility to tune the threshold to ensure it aligns with the constraint. On the other hand, the verification-based method is capable of producing relatively optimal results without the need for threshold tuning, although it lacks flexibility. We will systematically compare the two methods in experiments.

2.3 LLM CASCADES FOR REASONING TASKS

We instantiate the proposed two methods in LLM reasoning tasks with different sampling sources, resulting in 10 approaches, as summarized in Figure 2. Specifically, 6 approaches adopt vote-based decision-making: **CoT-1D-Vote** collects K answers sampled from prompting the weaker LLM with the CoT representation, and then calculates the answer consistency for decision-making, following Eq 2. Similarly, **PoT-1D-Vote** bases its decision-making on answers sampled from a PoT prompt. To diversify the sources of the answers, for each thought representation, we further consider sampling from two sets of CoT or PoT demonstrations, resulting in **CoT-2D-Vote** and **PoT-2D-Vote**, respectively. Finally, the vote-based approaches also include two variants leveraging a mixture of thought (MoT) representations. For **MoT-1D-Vote**, K_1 answers are sampled from the CoT prompt and another K_2 from the PoT prompt, and a union set of their answers are then used to compute the consistency score s . For **MoT-2D-Vote**, the procedure is similar, except that the CoT and the PoT prompts are annotated from two sets of demonstration examples.

The verification-based approaches assume answer samples from two different prompts. We instantiate 4 variants, including **CoT-2D-Verify**, where we prompt the weaker LLM with

two sets of CoT demonstrations, resulting in two answer sets for decision-making (Eq 3); PoT-2D-Verify, where we similarly prompt the weaker LLM with two sets of PoT demonstrations; MoT-1D-Verify, where we consider two sets of answers from two thought representation prompts (but on the same set of task demonstration examples); and MoT-2D-Verify, which additionally employs different sets of demonstrations when prompting the weaker LLM with different thought representations.

Cost-Comparable Sample Size Configuration To fairly compare all approaches in terms of their effectiveness in identifying easy (or correct) vs. hard (or correct) questions, we aim to configure their costs to be comparable. Since we use the same prompt to the stronger LLM, C^s is agnostic to the specific approach. However, different approaches may require varying costs in calling the weaker LLM (i.e., C^w). For example, even with the same total sampling size (i.e., $K = K_1 + K_2$), CoT-1D-Vote and CoT-2D-Vote result in different input token usages because the former prompt the LLM once but the latter prompt it twice, doubling the token usages. Therefore, we unify the C^w by a coarse-grained cost alignment. Our analysis suggests different configurations of K for different approaches. We refer readers to Appendix A for more details. Since our configuration can only yield comparable costs due to certain simplifications, rather than exact costs, we still report the actual token cost for each approach.

3 EXPERIMENT

3.1 EXPERIMENTAL SETTING

We evaluate our LLM cascade approaches on six datasets, covering (1) mathematical reasoning, including GSM8k (Cobbe et al., 2021), ASDIV (Ling et al., 2017), and TabMWP (Lu et al., 2023); (2) symbolic reasoning from BIG-Bench Hard (bench authors, 2023), including DATE and Navigate; and (3) causal reasoning, including CREPE (Zhang et al., 2023). In our pipeline, we leverage the GPT-3.5-turbo (4k context) as the weaker LLM and the GPT-4 (8k context) with CoT self-consistency (Wang et al., 2023, SC) as the stronger LLM. Throughout our experiments, we set the number of task demonstrations as $M = 8$. We use the same demonstration examples as prior work (Chen et al., 2022; Wei et al., 2022). When additional task demonstrations are needed (e.g., for 2D approaches), we randomly sample examples from the training data and manually annotate them with thought representations. We set the number of sampling paths as $K = 20$ for GPT-3.5-turbo and $K = 3$ for GPT-4. The sampling temperature by default is 0.4 for both LLMs. The metrics we use are the task accuracy and the relative cost compared with the cost of GPT-4 with CoT SC (denoted as GPT-4-CoT-SC). A lower relative cost and higher accuracy indicate better performance. We also compared our approaches with baselines using only the weaker LLM (i.e., GPT-3.5-CoT-SC, GPT-3.5-PoT-SC) or only the strong LLM in different ways (i.e., GPT-4-CoT-Greedy, GPT-4-PoT-Greedy, GPT-4-CoT-SC, GPT-4-PoT-SC). Our experimental details can be found in Appendix B and reproducible prompts in Appendix L.

3.2 MAIN RESULTS

Figure 3 illustrates the performance of our proposed approaches. For Vote-based approaches, we draw curves by changing the pre-defined threshold τ varying from 0.4 to 1. A high value of threshold signifies a more rigorous criterion for trusting the answers from the weaker LLM, making more examples transferred to the stronger LLM. Our observations are as follows:

Our pipeline achieves comparable task performance with significantly reduced costs. On average, all of our cascade variants (Vote or Verify) demonstrate significant cost efficiency. In particular, as shown in the average plot, the four MoT variants achieve comparable task performance (~ 0.929 accuracy) to GPT-4-CoT-SC (0.931) while demanding only 40% of its cost. On CREPE, MoT variants even outperform GPT-4-CoT-SC (0.885 vs. 0.871) at 47% of its cost. Our approaches based on CoT and PoT also exhibit the capability to save costs while maintaining the overall task performance. For example, CoT-2D-Vote achieved 0.924 task accuracy on average but demanded only 57% relative cost. These observations suggest the effectiveness of our cascade decision maker via checking the answer consistency of the weaker LLM.

Sampling from diverse prompt settings helps cascade decision-making. Our results show that variants involving diverse sources of sampling, such as CoT/PoT-2D-Vote and MoT-1D/2D-Vote, can more precisely distinguish between easy and hard reasoning questions,

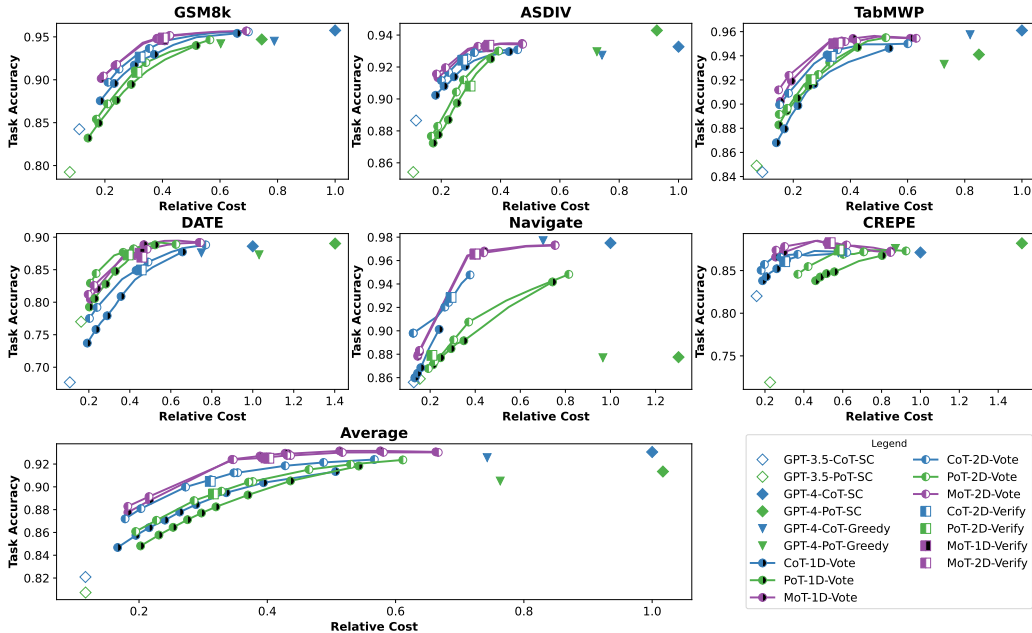


Figure 3: Main experiment results over six reasoning datasets. The bottom figure represents the average performance. The exact numerical results are included in Appendix C.

compared with their counterparts sampling from single sources, i.e., CoT/PoT-1D-Vote. For example, between CoT-2D-Vote and CoT-1D-Vote, the former outperforms the latter by 1.4% absolute accuracy under the same relative cost of 0.4 on average.

Mixing thought representations is particularly effective. Furthermore, we find that mixing the two thought representations (i.e., MoT-1D/2D-Vote) outperforms decision-making using either of them (i.e., CoT-1D/2D-vote and PoT-1D/2D-vote). This is illustrated by the gap in the average plot and is consistent on most datasets except DATE, where many test questions are very similar to the demonstration examples. Intuitively, this is because different thought representations can bring in more diverse “opinions” of the weaker LLM on the same input question, resembling how a group of experts with diverse perspectives could contribute to more effective results in collaborative work. It can also be viewed as “ensembling” LLMs, which utilizes the intuition that variants of the same model typically share few mistakes (Rokach, 2010). We provide a further investigation of this effect in Section 3.3. We also note that when using MoT, no obvious difference is perceived between using one set (i.e., MoT-1D-Vote) or two sets (i.e., MoT-2D-Vote) of task demonstrations. This result reveals that tuning the thought representations is more helpful for measuring an LLM’s (un)certainly on its answer than tuning the task demonstrations.

Increasing the threshold yields marginal benefits for MoT-1D/2D-Vote. As costs increase, the curves of MoT-1D/2D-Vote flatten out, showing that pushing the threshold to exceedingly high is unnecessary. This is because even for an easy question, the weaker LLM may still have hallucinations in a small set of answers. Setting the threshold too high can lead to the decision-making being influenced by them. It may result in easy questions being sent to the stronger LLM incorrectly, thus driving up the overall cost. In our practice, a threshold that can balance cost and accuracy typically falls between 0.5 and 0.6. When it comes to MoT-1D/2D-Verify, we can strike a balance between cost and accuracy as the verification method can tolerate a few hallucinations.

3.3 ANALYSIS ON MIXTURE OF THOUGHT REPRESENTATIONS

To understand the effect of MoT, we analyze three vote-based approaches. We first group questions into “easy” and “hard” based on whether the weaker LLM can answer them correctly (i.e., whether the majority-voted answer A^w is correct or not). For all easy/hard questions, we then calculate the average consistency score following Eq 2. We present the average consistency scores for each approach, categorizing easy and hard questions based on whether the weaker LLM can answer them correctly (i.e., A^w being correct). A higher average consistency score indicates our decision-maker places greater trust in the answer from the weaker LLM. Hence, an effective cascade decision-

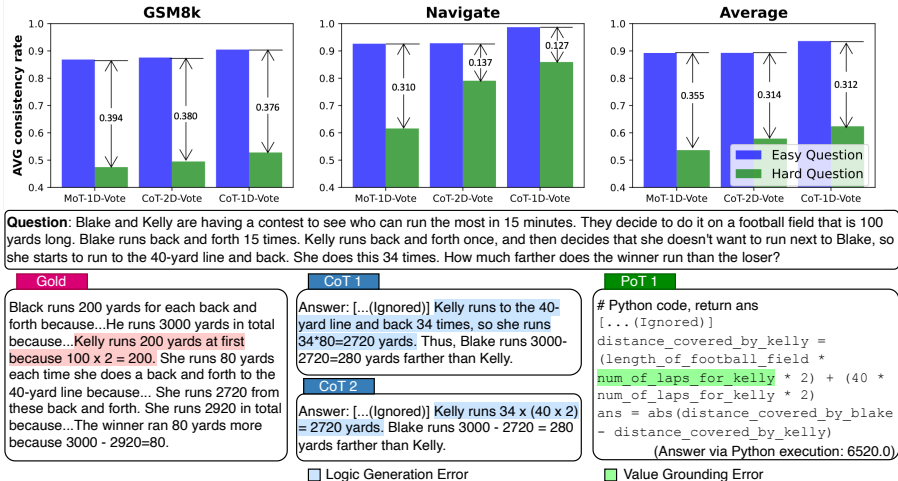


Figure 4: Average consistency scores (Top) and case studies (Bottom; simplified for presentation purposes) of various approaches showing the advantage of MoT. Shades highlight similar mistakes by CoT even based on different task demonstrations, and dissimilar mistakes by PoT.

maker should reveal relatively higher consistency scores for easy questions and lower ones for hard questions, leading to a larger "gap" between them.

We show the results in Figure 4 (Top). It is observed that all approaches lead to higher consistency scores on the easy questions than those on the hard questions, which explains their overall effectiveness for cascade decision-making. Because of involving two different thought representations, MoT-1D-Vote tends to have a lower consistency score compared with the two CoT approaches. However, it still ends up with a larger "gap" in the consistency scores for the easy and the hard questions, which is particularly prominent on the Navigate dataset where it gives the best performance gain (Figure 3). In contrast, CoT-1D-Vote results in the smallest score gap, indicating its weakness in distinguishing between the easy and the hard questions, particularly for Navigate. This weakness is mitigated by diversifying the prompting with two sets of task demonstrations (i.e., CoT-2D-Vote), but it still underperforms mixing the thought representations.

Finally, in Figure 4 (Bottom) we show that when CoT cannot answer a hard question, prompting the weaker LLM with another set of task demonstrations (but still under the CoT representation) often yields the same mistaken answer, which thus results in a high consistency score. On the contrary, PoT tends to make mistakes in a very different way and result in a different mistaken answer, which explains the low consistency score of MoT-1D-Vote. More cases are shown in Appendix D.

3.4 ROBUSTNESS EVALUATION

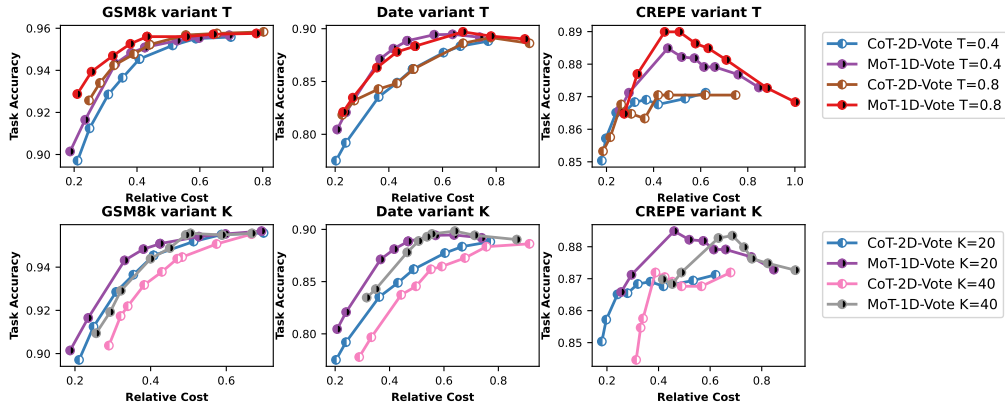


Figure 5: Robustness analysis with varying temperature T and sample size K .

We further analyze if our results are sensitive to the change of the sampling temperature T or the sample size K . We select datasets for each type of reasoning task and conduct experiments with `CoT-2D-Vote` and `MoT-1D-Vote`. Our results are shown in Figure 5. We first look into the effect when increasing the sampling temperature T from 0.4 (our default setting) to 0.8. For both approaches, increasing their temperature yields comparable or slightly better performance. This is owing to the increased answer diversity when the temperature gets higher. However, in any case, `MoT-1D-Vote` consistently outperforms `CoT-2D-Vote`. Increasing the sample size K from 20 to 40, on the other hand, leads to a rightward shift of the curves, implying that it requires higher cost to achieve the same task accuracy. This can be explained by the higher cost of the larger sample size, whereas increasing the sample size does not contribute to the detection of easy vs. hard questions. Like in the results of varying temperatures, `MoT` robustly outperforms `CoT` in any case.

3.5 COMPARISON TO EXTERNAL TEXT-BASED VERIFIERS




Figure 6: Comparison with external verifiers showing the advantages of our approaches based on answer consistency. We do not report finetuned verifiers for DATE as it does not have a training set.

As mentioned, prior work implemented the LLM cascade decision-maker by training an external verifier, which scores a question and its answer (from the weaker LLM) based on their literal descriptions (Chen et al., 2023a; Šakota et al., 2023). Related to the above work, Chen et al. (2023b) and Madaan et al. (2023) also showed the promise of prompting LLMs to evaluate their own responses. To perform a general comparison with such external verifiers, we conducted experiments on GSM8k, DATE, and CREPE with the following baselines: `Finetuned-Q`, which is a RoBERTa-base model (Liu et al., 2020) fine-tuned to decide whether a question should be routed to the stronger LLM based on its description; `Finetuned-QA`, which works similarly as `Finetuned-Q` but additionally takes the majority-voted answer from the weaker LLM (GPT-3.5-CoT-SC) as input; `LLM-Q`, where we instead prompt GPT-3.5-turbo as the verifier to judge based on the question description; and `LLM-QA`, which similarly employs GPT-3.5-turbo to decide upon the question and the weaker LLM’s majority-voted answer. We leave details of the baselines in Appendix E.

The results in Figure 6 show that incorporating the external verifiers cannot achieve comparable accuracy with `GPT-4-CoT-SC`. For example, on the GSM8k dataset, the highest accuracy with the external verifiers is 0.892, which is way lower than the accuracy (0.958) of `GPT-4-CoT-SC` and the accuracy (0.951) of our approaches. They also show lower task accuracies than our approaches under the same cost. It indicates that the external verifiers cannot yield satisfying results in complex reasoning tasks, which can be due to the intrinsic challenge of deciding question difficulty and answer correctness solely based on their textual descriptions. For the calibration analysis, please refer to the appendix I.

3.6 ADDITIONAL STUDIES

How weak can the weaker LLM be? We evaluate our approaches when adopting LLAMA2-13B (Touvron et al., 2023) as the weaker LLM and GPT-4 as the stronger LLM. The results are shown in the Appendix F. On the DATE dataset, our approaches still works. However, LLAMA2-13B as the weaker LLM doesn’t yield ideal results on GSM8k and CREPE. That is because most of the questions in GSM8k and CREPE are excessively complex for LLAMA2-13B. Therefore, LLAMA2-13B often fails to answer the questions consistently across multiple samplings. Hence, the choice of a weaker LLM should be contingent on the task’s level of difficulty. When the current weaker LLM struggles with the task, it is advisable to consider switching to a more powerful LLM.

Can the stronger LLM benefit from the weaker LLM hints? In our LLM cascade, all questions will obtain answers from the weaker LLM, no matter if they will be sent to the stronger LLM

or not. Therefore, an interesting question is whether the answer produced by the weaker LLM (correct or incorrect) can provide “hints” to enhance the stronger LLM. To answer this question, we experimented with MoT-1D-Verify, where we additionally passed the two inconsistent answers $A_1^{w'}$ and $A_2^{w'}$ as hints to the stronger LLM following the format of prior work (Zheng et al., 2023). We observed that the hints can only yield slight improvement on DATE, but greatly hurt the model performance on GSM8k and CREPE. Therefore, we conclude that hints from the weaker LLM, when it is uncertain about the question, do not help the stronger LLM. Details are in Appendix G.

Can our method generalize to factual-based tasks? We also explored whether our method can be generalized to factual-based reasoning tasks in Appendix J.

4 RELATED WORK

Earlier research has delved into techniques to enhance the cost efficiency of LLMs, including quantization, pruning, and decoding methods (Bai et al., 2022; Kurtic et al., 2023; Leviathan et al., 2022). While many LLMs are closed-source, certain studies have concentrated on how to utilize the API efficiently. Chen et al. (2023a) proposed an LLM cascade, which sends a query to a list of LLMs sequentially if the answers provided by the prior LLMs are regarded unacceptable. A concurrent effort is made by Šakota et al. (2023), where the proposed approach determines which LLM to invoke prior to sending the query to any LLM. Both of these approaches employed external verifiers fine-tuned from smaller LMs and therefore demanded substantial training data and additional model learning. In our experiments, we showed the ineffectiveness of these approaches when they were used to save costs for complex reasoning tasks, which echos the discovery of the shortcomings of LLM fine-tuning in recent work (Ni et al., 2023). In contrast, we proposed a novel solution for LLM cascades based on the answer consistency of the weaker LLM, which is training-free and is not limited to a specific weaker LLM.

Reasoning with LLMs LLMs have demonstrated strong capability in solving reasoning tasks (Rae et al., 2021; Lewkowycz et al., 2022; Zhong et al., 2023). The standard few-shot in-context learning for LLMs is to feed the question to the LLM and let it generate the answer directly (Brown et al., 2020). Wei et al. (2022) then proposed the Chain-of-Thought approach, showing that prompting LLMs to think step by step can significantly improve their performance in solving reasoning problems, as it activates the built-in reasoning capabilities of LLMs (Madaan & Yazdanbakhsh, 2023). To further facilitate the mathematical steps in a reasoning task, Chen et al. (2022) and Gao et al. (2023) proposed the Program-of-Thought prompting, using code as the intermediate step and executing it with an external interpreter to entangle reasoning and numerical calculation. While most existing works studied CoT and PoT independently as alternative approaches, several other works have delved into the synergy between them, e.g., integrating code into CoT (Cheng et al., 2023b; He-Yueya et al., 2023), ensembling algebraic and code representations (Imani et al., 2023; Yue et al., 2023), or performing automatic selection between CoT and PoT (Zhao et al., 2023). Similar to these works, our approaches employ CoT and PoT jointly; however, unlike them, we leverage this synergy for the novel application of cost-saving for LLM reasoning. On the other hand, like our LLM cascades, there are also recent works building pipelines chaining multiple LLMs (Varshney & Baral, 2022; Xiong et al., 2023a; Li et al., 2023; Lin et al., 2023a), but these works have a different focus on improving the task accuracy, rather than reducing the LLM costs.

Uncertainty of LLMs Our approach involves measuring the uncertainty of the results of LLMs. Multiple concurrent works have investigated using the voting score to evaluate the answer’s uncertainty (Xiong et al., 2023b; Si et al., 2023; Cai et al., 2023). However, they dismissed the potential application for different representations of intermediate steps in the evaluation.

5 CONCLUSION

We introduce a simple yet efficient and universally applicable economical pipeline to dynamically decide the LLMs in reasoning tasks, so as to save the token costs. Our approach based on checking the answer consistency of the weaker LLM is novel and effective. Our discoveries emphasize that leveraging prompts with a mixture of thought representations in weaker LLM achieves the best performance as it introduces diverse answers. Compared with fully employing the stronger LLM, our pipeline requires approximately 40% of expenses to achieve a comparable result. Future works are listed in Appendix K.

6 ETHICS STATEMENT

Our research endeavors to build a cost-efficient pipeline in solving reasoning problems in the LLM era. The significance of this work holds the potential to benefit a diverse range of organizations, particularly those with limited financial resources such as local businesses, educational institutions, and non-profit organizations. By devising cost-effective solutions, we empower these resource-constrained entities to harness the reasoning ability of LLMs cheaply, thereby fostering fairness and inclusivity within the NLP community. This aligns with the broader goal of ensuring that advancements in NLP do not remain exclusive to a select few, but are accessible to a wider audience, regardless of their financial constraints.

Moreover, for these industry giants, the cost savings facilitated by our findings can also be substantial, particularly when facing extremely large throughputs. By optimizing cost-efficiency, our research contributes to not only economic savings but also environmental sustainability by mitigating the carbon footprint associated with running large-scale computations.

7 REPRODUCIBILITY STATEMENT

Our pipeline is simple to implement and reproducible. We have documented all the experimental details, both in the main text and the appendix sections. While we cannot include the complete text of every prompt due to their excessive length, we do provide examples of each prompt in Appendix L, facilitating the reader’s comprehension of the style employed. All of our implementations (including the complete prompt scripts, the code for training external verifiers, the code for approach evaluation, etc.) are publicly available at https://github.com/MurongYue/LLM_MoT_cascade.

8 ACKNOWLEDGEMENT

This project was supported by Microsoft. It was also partially supported by resources provided by the Office of Research Computing at George Mason University (URL: <https://orc.gmu.edu>) and funded in part by grants from the National Science Foundation (Awards Number 1625039 and 2018631). Murong and Ziyu appreciate the funding support from George Mason College of Engineering and Computing.

REFERENCES

- Haoli Bai, Lu Hou, Lifeng Shang, Xin Jiang, Irwin King, and Michael Lyu. Towards efficient post-training quantization of pre-trained language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=tvDRmAxGIjw>.
- BIG bench authors. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=uyTL5Bvosj>.
- 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. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Zefan Cai, Baobao Chang, and Wenjuan Han. Human-in-the-loop through chain-of-thought. *arXiv preprint arXiv:2306.07932*, 2023.
- Lingjiao Chen, Matei Zaharia, and James Zou. Frugalgpt: How to use large language models while reducing cost and improving performance. *arXiv preprint arXiv:2305.05176*, 2023a.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*, 2022.
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128*, 2023b.

- Zhoujun Cheng, Jungo Kasai, and Tao Yu. Batch prompting: Efficient inference with large language model apis. *arXiv preprint arXiv:2301.08721*, 2023a.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. Binding language models in symbolic languages. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=lH1PV42cbF>.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. PAL: Program-aided language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 10764–10799. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/gao23f.html>.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. *Transactions of the Association for Computational Linguistics (TACL)*, 2021.
- Tanmay Gupta and Aniruddha Kembhavi. Visual programming: Compositional visual reasoning without training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14953–14962, June 2023.
- Joy He-Yueya, Gabriel Poesia, Rose E Wang, and Noah D Goodman. Solving math word problems by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*, 2023.
- Shima Imani, Liang Du, and Harsh Shrivastava. Mathprompter: Mathematical reasoning using large language models. *arXiv preprint arXiv:2303.05398*, 2023.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. Language models (mostly) know what they know, 2022.
- Eldar Kurtic, Elias Frantar, and Dan Alistarh. Ziplm: Hardware-aware structured pruning of language models. *arXiv preprint arXiv:2302.04089*, 2023.
- Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *International Conference on Machine Learning*, 2022. URL <https://api.semanticscholar.org/CorpusID:254096365>.
- Aitor Lewkowycz, Anders Johan Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Venkatesh Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=IFXTZERXdm7>.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12286–12312, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.687. URL <https://aclanthology.org/2023.acl-long.687>.

- Bill Yuchen Lin, Yicheng Fu, Karina Yang, Faeze Brahman, Shiyu Huang, Chandra Bhagavathula, Prithviraj Ammanabrolu, Yejin Choi, and Xiang Ren. Swiftsage: A generative agent with fast and slow thinking for complex interactive tasks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023a. URL <https://openreview.net/forum?id=Rzk3GP1HN7>.
- Jianzhe Lin, Maurice Diesendruck, Liang Du, and Robin Abraham. Batchprompt: Accomplish more with less, 2023b.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 158–167, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1015. URL <https://aclanthology.org/P17-1015>.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Ro{bert}a: A robustly optimized {bert} pre-training approach, 2020. URL <https://openreview.net/forum?id=SyxS0T4tvS>.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. In *International Conference on Learning Representations (ICLR)*, 2023.
- Aman Madaan and Amir Yazdanbakhsh. Text and patterns: For effective chain of thought it takes two to tango, 2023. URL <https://openreview.net/forum?id=z9fXRC5XdT>.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.
- Ansong Ni, Jeevana Priya Inala, Chenglong Wang, Alex Polozov, Christopher Meek, Dragomir Radev, and Jianfeng Gao. Learning math reasoning from self-sampled correct and partially-correct solutions. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=4D4TSJE6-K>.
- OpenAI. Gpt-4 technical report, 2023.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021.
- Lior Rokach. Ensemble-based classifiers. *Artificial intelligence review*, 33:1–39, 2010.
- Marija Šakota, Maxime Peyrard, and Robert West. Fly-swat or cannon? cost-effective language model choice via meta-modeling. *arXiv preprint arXiv:2308.06077*, 2023.
- Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Lee Boyd-Graber, and Lijuan Wang. Prompting GPT-3 to be reliable. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=98p5x51L5af>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Neeraj Varshney and Chitta Baral. Model cascading: Towards jointly improving efficiency and accuracy of NLP systems. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11007–11021, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.756. URL <https://aclanthology.org/2022.emnlp-main.756>.

- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=1PL1NIMMrw>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=_VjQlMeSB_J.
- Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. Diving into the inter-consistency of large language models: An insightful analysis through debate. *arXiv preprint arXiv:2305.11595*, 2023a.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*, 2023b.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. Mammoth: Building math generalist models through hybrid instruction tuning. *arXiv preprint arXiv:2309.05653*, 2023.
- Li Zhang, Hainiu Xu, Yue Yang, Shuyan Zhou, Weiqiu You, Manni Arora, and Chris Callison-burch. Causal reasoning of entities and events in procedural texts. In *Findings of the Association for Computational Linguistics: EACL 2023*, pp. 415–431, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.findings-eacl.31>.
- Xu Zhao, Yuxi Xie, Kenji Kawaguchi, Junxian He, and Qizhe Xie. Automatic model selection with large language models for reasoning. *arXiv preprint arXiv:2305.14333*, 2023.
- Chuanyang Zheng, Zhengying Liu, Enze Xie, Zhenguo Li, and Yu Li. Progressive-hint prompting improves reasoning in large language models. *arXiv preprint arXiv:2304.09797*, 2023.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models. *arXiv preprint arXiv:2304.06364*, 2023.

A COST ANALYSIS OF LLM CASCADE APPROACHES

In Section 2.3, we introduced ten approaches to implement the LLM cascade decision maker. As we are interested in saving the cost of LLM usage, we first analyze the token usage of each approach. While different thought representations of the same task example could induce different token usages, the difference is very hard to quantify and also depends on the specific reasoning task. For example, for GSM8k (Cobbe et al., 2021), the PoT representation of the same example is typically shorter than its CoT counterpart, while a reversed comparison is observed for DATE and Navigate (bench authors, 2023). In addition, different demonstration examples may also consist of different numbers of tokens, which is hard to quantify as well. To provide a unified analysis, we assume that every demonstration example for any task, regardless of the thought representation in use, consumes N_{tok} tokens.

Based on our hypothesis, we conducted a cost analysis, with results summarized in Table 1. To give an example, consider comparing the token costs of CoT-1D-Vote and CoT-2D-Vote. For CoT-1D-Vote, the cost of calling weaker LLM is $(MC_i^w + KC_o^w) \times N_{tok}$; for CoT-2D-Vote, it is $(2MC_i^w + K_1C_o^w + K_2C_o^w) \times N_{tok}$. Here, C_i^w and C_o^w are the input and output per-token cost for the weaker LLM, respectively, and M is the number of demonstrations in each prompt. In practice, we set $K_1 = K_2 = K_{2D}/2$, where K_{2D} is the total sample size of two sets of demonstration examples, and the total cost for CoT-2D-Vote can then be rewritten to $(2MC_i^w + K_{2D}C_o^w)$. To keep the costs of CoT-1D-Vote and CoT-2D-Vote consistent, we set $K_{2D} = K_{1D} - M \times \frac{C_i^w}{C_o^w}$. Similarly, we can get the total sample size of different representations K_{MoT} . The analysis guides us in configuring each approach to be ‘‘cost-comparable’’ (i.e., leading to similar C^w ’s). Finally, we note that we do not change the number of task demonstrations (i.e., M) because LLMs are very sensitive to this configuration.

Method	Certainty Value (s) Definition	Cost Estimation
CoT-1D-Vote	Agreement of K CoT samples on the majority-voted answer A^w , $s = \sum_{i=1}^K \mathbb{1}_{A_i^w=A^w}/K$	$(MC_i^w + K_{1D}C_o^w) \times N_{tok}$
PoT-1D-Vote	Agreement of K PoT samples on the majority-voted answer A^w , $s = \sum_{i=1}^K \mathbb{1}_{A_i^w=A^w}/K$	$(MC_i^w + K_{1D}C_o^w) \times N_{tok}$
MoT-1D-Vote	Agreement of $K_{MoT}/2$ CoT and $K_{MoT}/2$ PoT samples on the majority-voted answer A^w , $s = \sum_{i=1}^{K_{MoT}} \mathbb{1}_{A_i^w=A^w}/K_{MoT}$	$2 \times (MC_i^w + \frac{K_{MoT}}{2}C_o^w) \times N_{tok}$
CoT-2D-Vote	Agreement of $K_{2D}/2$ CoT samples over 2 sets of task demonstrations on the majority-voted answer A^w , $s = \sum_{i=1}^{K_{2D}} \mathbb{1}_{A_i^w=A^w}/K_{2D}$	$2 \times (MC_i^w + \frac{K_{2D}}{2}C_o^w) \times N_{tok}$
PoT-2D-Vote	Agreement of $K_{2D}/2$ PoT samples over 2 sets of task demonstrations on the majority-voted answer A^w , $s = \sum_{i=1}^{K_{2D}} \mathbb{1}_{A_i^w=A^w}/K_{2D}$	$2 \times (MC_i^w + \frac{K_{2D}}{2}C_o^w) \times N_{tok}$
MoT-2D-Vote	Agreement of $K_{MoT}/2$ CoT samples on one set of task demos and $K_{MoT}/2$ PoT samples on another set of task demos on their majority-voted answer A^* , $s = \sum_{i=1}^{K_{MoT}} \mathbb{1}_{A_i^w=A^w}/K_{MoT}$	$2 \times (MC_i^w + \frac{K_{MoT}}{2}C_o^w) \times N_{tok}$
CoT-2D-Verify	Agreement of the two majority-voted answers $A_1^{w'}$ and $A_2^{w'}$ over 2 sets of CoT task demonstrations with K_{2D} examples, $s = \mathbb{1}_{A_1^{w'}=A_2^{w'}}$	$2 \times (MC_i^w + \frac{K_{2D}}{2}C_o^w) \times N_{tok}$
PoT-2D-Verify	Agreement of the two majority-voted answers $A_1^{w'}$ and $A_2^{w'}$ over 2 sets of PoT task demonstrations with K_{2D} examples, $s = \mathbb{1}_{A_1^{w'}=A_2^{w'}}$	$2 \times (MC_i^w + \frac{K_{2D}}{2}C_o^w) \times N_{tok}$
MoT-2D-Verify	Agreement of the two majority-voted answers $A_1^{w'}$ and $A_2^{w'}$ over one set of CoT and another set of PoT task demonstrations with $K_{MoT}/2$ examples, $s = \mathbb{1}_{A_1^{w'}=A_2^{w'}}$	$2 \times (MC_i^w + \frac{K_{MoT}}{2}C_o^w) \times N_{tok}$
MoT-1D-Verify	Agreement of the two majority-voted answers $A_1^{w'}$ and $A_2^{w'}$ for CoT and PoT over the same set of task demonstrations with $K_{MoT}/2$ examples, $s = \mathbb{1}_{A_1^{w'}=A_2^{w'}}$	$2 \times (MC_i^w + \frac{K_{MoT}}{2}C_o^w) \times N_{tok}$

Table 1: Investigated approaches and cost calculations. For simplicity, we assume M -shot demonstration prompting and that every demonstration example has N_{tok} tokens. C_i^w is the input price per token and C_o^w is the output price per token. To keep consistent costs of using the weaker LLM, we set $K_{MoT} = K_{2D} = K_{1D} - M \times \frac{C_i^w}{C_o^w}$.

B MAIN EXPERIMENTAL DETAILS

Implementation Details We run each approach two times to reduce variance and report the average results. Since DATE and Navigate do not have a training set, before experiments we sampled 8 shots of examples randomly and annotated them to be the second set of task demonstrations (for 2D

approaches). The remaining examples are used consistently across all experiments as the test set. We adopt “*Python code, return ans*” in PoT prompt to let the LLM generate the Python code (Chen et al., 2022). The interpreter we used for executing PoT is Python 3.10, with some packages, such as `DateTime`² to facilitate the execution of the generated code.

Metrics Details We evaluate the methods based on task accuracy and cost efficiency, with a lower cost and higher accuracy indicating better performance. For accuracy, when the answer is a string, we use exact matching. When the answer is a number, we relax the evaluation criteria due to potential variations in the exact computations carried out by the external interpreter. Following prior work (Chen et al., 2022), we adopt the tolerance to 0.001. For cost efficiency, we calculate the actual token number based on the `tiktoken`³ and the total cost⁴ for each method and get the relative cost by comparing it with the total cost of `GPT-4-CoT-SC`.

C MAIN RESULT TABLES

In this section, we show the exact numerical results that are reported in Figure 3. The GSM8k result is in Table 2. The ASDIV result is in Table 3. The TabMWP result is in Table 4. The DATE result is in Table 5. The Navigation result is in Table 6. The CREPE result is in Table 7. The average results are in Table 8.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.875	0.184	0.832	0.141	0.897	0.211	0.854	0.170	0.901	0.187	0.904	0.195
0.50	0.896	0.233	0.850	0.179	0.912	0.249	0.872	0.209	0.917	0.235	0.918	0.241
0.55	0.907	0.265	0.861	0.206	0.929	0.309	0.894	0.262	0.943	0.331	0.942	0.334
0.60	0.917	0.302	0.876	0.240	0.937	0.354	0.910	0.305	0.948	0.381	0.948	0.381
0.65	0.924	0.336	0.886	0.265	0.945	0.408	0.920	0.342	0.951	0.425	0.952	0.425
0.70	0.929	0.372	0.895	0.292	0.945	0.408	0.920	0.342	0.951	0.425	0.952	0.425
0.80	0.941	0.438	0.910	0.346	0.952	0.514	0.933	0.427	0.954	0.528	0.955	0.530
0.90	0.950	0.518	0.924	0.410	0.955	0.587	0.938	0.485	0.955	0.598	0.956	0.600
1.00	0.954	0.660	0.940	0.517	0.956	0.699	0.947	0.565	0.957	0.693	0.957	0.691
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.926	0.322	0.909	0.311	0.947	0.401	0.949	0.403				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.842	0.111	0.792	0.078	0.945	0.789	0.942	0.603	0.958	1.000	0.947	0.745

Table 2: Exact numerical results on GSM8k.

²<https://pypi.org/project/DATETIME/>

³<https://platform.openai.com/tokenizer>

⁴<https://openai.com/pricing>

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.902	0.181	0.872	0.173	0.911	0.200	0.877	0.167	0.913	0.189	0.916	0.184
0.50	0.908	0.209	0.878	0.190	0.917	0.225	0.883	0.188	0.918	0.218	0.920	0.214
0.55	0.910	0.226	0.882	0.210	0.924	0.262	0.894	0.222	0.930	0.291	0.931	0.290
0.60	0.914	0.243	0.887	0.224	0.926	0.288	0.904	0.251	0.932	0.324	0.933	0.324
0.65	0.917	0.261	0.893	0.239	0.929	0.312	0.912	0.273	0.933	0.346	0.933	0.348
0.70	0.920	0.281	0.897	0.254	0.929	0.312	0.912	0.275	0.933	0.347	0.933	0.348
0.80	0.924	0.311	0.906	0.283	0.930	0.365	0.923	0.328	0.935	0.396	0.934	0.400
0.90	0.927	0.348	0.916	0.318	0.930	0.402	0.927	0.360	0.935	0.430	0.934	0.433
1.00	0.930	0.429	0.925	0.367	0.931	0.458	0.930	0.393	0.935	0.473	0.934	0.474
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.924	0.272	0.908	0.297	0.933	0.357	0.934	0.361				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.887	0.115	0.854	0.106	0.927	0.742	0.930	0.725	0.933	1.000	0.943	0.927

Table 3: Exact numerical results on ASDIV.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.868	0.141	0.883	0.150	0.900	0.152	0.892	0.151	0.902	0.155	0.912	0.149
0.50	0.880	0.169	0.894	0.177	0.909	0.183	0.897	0.179	0.919	0.193	0.924	0.185
0.55	0.890	0.191	0.899	0.195	0.934	0.273	0.916	0.238	0.951	0.324	0.952	0.329
0.60	0.899	0.216	0.905	0.214	0.941	0.318	0.925	0.288	0.952	0.370	0.952	0.381
0.65	0.910	0.245	0.910	0.236	0.945	0.355	0.935	0.327	0.954	0.409	0.952	0.421
0.70	0.917	0.273	0.915	0.254	0.945	0.355	0.935	0.327	0.954	0.409	0.952	0.421
0.80	0.927	0.327	0.924	0.292	0.950	0.448	0.946	0.411	0.957	0.484	0.956	0.496
0.90	0.935	0.393	0.934	0.345	0.950	0.513	0.952	0.460	0.955	0.531	0.956	0.548
1.00	0.946	0.535	0.947	0.425	0.950	0.601	0.955	0.524	0.955	0.610	0.955	0.630
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.940	0.333	0.920	0.262	0.950	0.342	0.951	0.357				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.844	0.092	0.849	0.072	0.958	0.819	0.933	0.728	0.961	1.000	0.941	0.849

Table 4: Exact numerical results on TabMWP.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.737	0.192	0.793	0.203	0.775	0.203	0.829	0.208	0.805	0.207	0.812	0.197
0.50	0.758	0.235	0.806	0.229	0.792	0.240	0.844	0.235	0.821	0.240	0.826	0.225
0.55	0.769	0.259	0.818	0.256	0.835	0.362	0.872	0.324	0.871	0.367	0.868	0.367
0.60	0.779	0.289	0.828	0.283	0.849	0.431	0.877	0.367	0.881	0.417	0.872	0.425
0.65	0.795	0.328	0.837	0.306	0.862	0.489	0.882	0.419	0.889	0.469	0.881	0.484
0.70	0.809	0.357	0.848	0.329	0.862	0.489	0.882	0.419	0.889	0.469	0.881	0.484
0.80	0.835	0.425	0.860	0.378	0.877	0.604	0.890	0.502	0.894	0.569	0.888	0.583
0.90	0.854	0.511	0.877	0.425	0.883	0.666	0.892	0.552	0.895	0.637	0.891	0.654
1.00	0.877	0.657	0.888	0.523	0.888	0.770	0.889	0.625	0.892	0.739	0.892	0.744
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.850	0.456	0.873	0.392	0.875	0.449	0.869	0.459				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.677	0.107	0.770	0.162	0.876	0.750	0.873	1.031	0.886	1.000	0.890	1.402

Table 5: Exact numerical results on DATE.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.860	0.132	0.871	0.213	0.898	0.125	0.868	0.192	0.878	0.145	0.880	0.146
0.50	0.860	0.134	0.877	0.246	0.898	0.125	0.874	0.219	0.882	0.153	0.883	0.152
0.55	0.862	0.140	0.881	0.270	0.913	0.233	0.882	0.261	0.964	0.368	0.964	0.365
0.60	0.864	0.144	0.885	0.293	0.920	0.265	0.892	0.305	0.966	0.402	0.966	0.400
0.65	0.866	0.149	0.889	0.317	0.924	0.282	0.908	0.371	0.968	0.440	0.967	0.437
0.70	0.869	0.159	0.892	0.349	0.924	0.282	0.908	0.371	0.968	0.440	0.967	0.437
0.80	0.876	0.175	0.903	0.432	0.934	0.318	0.926	0.537	0.971	0.545	0.970	0.543
0.90	0.884	0.193	0.920	0.548	0.941	0.339	0.935	0.649	0.973	0.626	0.972	0.624
1.00	0.901	0.240	0.942	0.743	0.948	0.377	0.948	0.815	0.973	0.754	0.973	0.754
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.929	0.293	0.879	0.208	0.966	0.400	0.966	0.398				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.856	0.128	0.859	0.154	0.977	0.701	0.877	0.966	0.975	1.000	0.878	1.302

Table 6: Exact numerical results on Navigate.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.838	0.188	0.838	0.461	0.850	0.180	0.846	0.369	0.866	0.255	0.874	0.257
0.50	0.843	0.208	0.842	0.485	0.857	0.197	0.855	0.422	0.871	0.295	0.878	0.301
0.55	0.847	0.227	0.846	0.500	0.865	0.241	0.870	0.552	0.885	0.460	0.885	0.471
0.60	0.852	0.261	0.846	0.517	0.865	0.280	0.869	0.603	0.882	0.519	0.882	0.525
0.65	0.856	0.284	0.847	0.537	0.868	0.320	0.869	0.658	0.882	0.573	0.881	0.575
0.70	0.863	0.312	0.849	0.557	0.869	0.369	0.872	0.708	0.879	0.614	0.880	0.620
0.80	0.867	0.374	0.853	0.603	0.868	0.419	0.874	0.754	0.879	0.660	0.878	0.663
0.90	0.873	0.455	0.861	0.683	0.869	0.534	0.874	0.851	0.877	0.760	0.873	0.760
1.00	0.872	0.594	0.868	0.803	0.871	0.621	0.873	0.925	0.873	0.847	0.871	0.846
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.860	0.302	0.874	0.581	0.882	0.534	0.883	0.536				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.820	0.157	0.719	0.226	0.871	0.577	0.876	0.871	0.871	1.000	0.882	1.522

Table 7: Exact numerical results on CREPE.

Voting Method												
Threshold	CoT-1D-Vote		PoT-1D-Vote		CoT-2D-Vote		PoT-2D-Vote		MoT-1D-Vote		MoT-2d-Vote	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
0.40	0.847	0.167	0.848	0.202	0.872	0.178	0.861	0.195	0.878	0.183	0.883	0.183
0.50	0.857	0.195	0.858	0.231	0.881	0.203	0.871	0.227	0.888	0.216	0.891	0.216
0.55	0.864	0.216	0.864	0.253	0.900	0.273	0.888	0.286	0.924	0.345	0.924	0.347
0.60	0.871	0.240	0.871	0.276	0.906	0.312	0.896	0.329	0.927	0.389	0.925	0.391
0.65	0.878	0.263	0.877	0.297	0.912	0.348	0.904	0.371	0.929	0.428	0.928	0.431
0.70	0.884	0.289	0.882	0.320	0.912	0.354	0.905	0.377	0.929	0.433	0.927	0.436
0.80	0.895	0.336	0.893	0.370	0.918	0.428	0.915	0.465	0.932	0.513	0.930	0.517
0.90	0.904	0.394	0.905	0.436	0.921	0.488	0.920	0.530	0.932	0.576	0.930	0.581
1.00	0.913	0.506	0.918	0.543	0.924	0.567	0.924	0.611	0.931	0.663	0.930	0.667
Verify Method												
	CoT-2D-Verify		PoT-2D-Verify		MoT-1D-Verify		MoT-2D-Verify					
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost				
	0.905	0.311	0.894	0.316	0.925	0.398	0.925	0.402				
Without Cascade												
	GPT-3-CoT-SC		GPT-3-PoT-SC		GPT-4 CoT-Greedy		GPT-4 PoT-Greedy		GPT-4 CoT-SC		GPT-4 PoT-SC	
	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost	Acc	Cost
	0.821	0.116	0.807	0.117	0.926	0.743	0.905	0.763	0.931	1.000	0.913	1.017

Table 8: Exact numerical results on average across all datasets.

D CASE STUDY

In this section, we perform a case study to learn why MoT can achieve better performance. When CoT makes a mistake, prompting the LLM with another set of CoT task demonstrations often leads to the same mistake. In contrast, PoT more frequently either directly gives a correct answer or makes a different mistake than CoT’s. Among those cases, 43% fall into the former category that PoT can solve the question, while the remaining 57% fall into the latter that PoT can make a different mistake. Here we use CoT1, CoT2, and PoT1 to represent the different prompts. CoT1 and CoT2 share the same representation but different demonstrations while CoT1 and PoT1 share the same demonstration but different representations.

D.1 CASES POT CAN SOLVE

Here we show cases satisfying that PoT can solve the question. The first case shows that the PoT can provide different reasoning paths to avoid the same mistake with CoT and the second case shows that the PoT can avoid the computation errors.

D.1.1 CASE 1

Question: Juan and his brother Carlos are selling lemonade. For each gallon they make it costs 3 for lemons and 2 for sugar. They sell each glass for 0.50 and get 20 glasses per gallon. If they made 25 in profit, how much did they spend on lemons?

Gold Answer: Each gallon sells for a total of 10 because $20 \times 0.5 = 10$
 Each gallon costs 5 to make because $3 + 2 = 5$
 They make 5 off each gallon because $10 - 5 = 5$
 They made and sold 5 gallons because $25/5 = 5$
 They spent 15 on lemons because $5 \times 3 = 15$

CoT1: Answer: They make 10 on each gallon of lemonade because $20 \times 0.5 = 10$
 They made 2.5 gallons of lemonade because $25/10 = 2.5$
 They spent 7.5 on lemons because $2.5 \times 3 = 7.5$
 They spent 5 on sugar because $2.5 \times 2 = 5$
 They spent 12.5 in total because $7.5 + 5 = 12.5$
 ans = 7.5

CoT2: Answer: They make 10 per gallon because $20 \times 0.5 = 10$
 They sold 2.5 gallons because $25/10 = 2.5$
 They spent 7.5 on sugar because $2.5 \times 2 \times 1.5 = 7.5$
 They spent 7.5 on lemons because $2.5 \times 3 = 7.5$
 ans = 7.5

PoT1:

```
# Python code, return ans
cost_per_gallon = 3 + 2
num_of_glasses_per_gallon = 20
profit_per_gallon = num_of_glasses_per_gallon * 0.5 - \
cost_per_gallon
total_profit = 25
total_gallons = total_profit / profit_per_gallon
cost_of_lemons = total_gallons * 3
ans = cost_of_lemons
#ans = 15.0(This line is not generated by LLM.)
```

D.1.2 CASE 2

Question: Robert wants to buy a bookshelf and a sofa and has received two offers. Furniture United’s offer includes a 1,350 advance payment and 6 monthly installments of 350 each. Cozy Homes’ offer includes a 1,100 advance payment and 9 monthly installments of 250 each. What is the cost difference between the two offers?

Gold Answer: Furniture United’s offer is $1,350 + (6 \times 350) = 3,450$ Cozy Homes’ offer is $1,100 + (9 \times 250) = 3,350$. The difference is $3,450 - 3,350 = 100$

CoT1: Answer: Furniture United’s total cost would be $1350 + 6 \times 350 = 3150$. Cozy Homes’ total cost would be $1100 + 9 \times 250 = 3350$. The cost difference between the two offers is $3350 - 3150 = 200$. ans = 200

CoT2: Answer: The total cost of Furniture United’s offer is $1350 + (6 \times 350) = 3,150$. The total cost of Cozy Homes’ offer is $1100 + (9 \times 250) = 3,350$. The cost difference between the two offers is $3350 - 3150 = 200$. ans = 200.

PoT1:

```
furniture_united_advance = 1350
furniture_united_monthly_installment = 350
cozy_homes_advance = 1100
cozy_homes_monthly_installment = 250
furniture_united_total_cost = furniture_united_advance + \
    (6 * furniture_united_monthly_installment)
cozy_homes_total_cost = cozy_homes_advance + \
    (9 * cozy_homes_monthly_installment)
ans = furniture_united_total_cost - cozy_homes_total_cost
# ans = 100.0(This line is not generated by LLM.)
```

D.2 CASES PoT CANNOT SOLVE

In this section, we show some cases that PoT cannot solve but will make a different mistake. In case 1, the CoT forgot to calculate the first 40-yard line but the PoT calculated the first 40-yard but with an incorrect grounding value. In case 2, the CoT forgot to remove the lowest score. The PoT removed the lowest score but grounded an incorrect number of total tests. We can see that although both the CoT and PoT make mistakes, the reasons why they make mistakes are different, making the final results different.

D.2.1 CASE 1

Question: Blake and Kelly are having a contest to see who can run the most in 15 minutes. They decide to do it on a football field that is 100 yards long. Blake runs back and forth 15 times. Kelly runs back and forth once, and then decides that she doesn’t want to run next to Blake, so she starts to run to the 40-yard line and back. She does this 34 times. How much farther does the winner run than the loser?

Gold Answer: Blake runs 200 yards for each back and forth because $100 \times 2 = 200$ He runs 3000 yards in total because $200 \times 15 = 3000$ Kelly runs 200 yards at first because $100 \times 2 = 200$ She runs 80 yards each time she does a back and forth to the 40-yard line because $40 \times 2 = 80$ She runs 2720 from these back and forth because $34 \times 80 = 2720$ She runs 2920 in total because $2720 + 200 = 2920$ Blake won because $3000 > 2920$ The winner ran 80 yards more than the loser because $3000 - 2920 = 80$

CoT1: Answer: Blake runs back and forth 15 times, so she runs $15 \times 100 \times 2 = 3000$ yards. Kelly runs to the 40-yard line and back 34 times, so she runs $34 \times 80 = 2720$ yards. Thus, Blake runs $3000 - 2720 = 280$ yards farther than Kelly. ans = 280

CoT2: Answer: Blake runs a total of $15 \times 100 \times 2 = 3000$ yards. Kelly runs $34 \times (40 \times 2) = 2720$ yards. Blake runs $3000 - 2720 = 280$ yards farther than Kelly. ans = 280.

PoT1:

```
# Python code, return ans
length_of_football_field = 100
num_of_laps_for_blake = 15
num_of_laps_for_kelly = 34
distance_covered_by_blake = length_of_football_field * \
    num_of_laps_for_blake * 2
distance_covered_by_kelly = (length_of_football_field * \
```

```
num_of_laps_for_kelly * 2) + (40 * num_of_laps_for_kelly * 2)
ans = abs(distance_covered_by_blake - distance_covered_by_kelly)
# ans = 6520.0(This line is not generated by LLM.)
```

D.2.2 CASE 2

Question: Brinley is in Mr. Bert’s math class. Mr. Bert gives six tests each semester. He allows the students to remove the lowest score from the average each semester. If Brinley has scores of 89, 71, 92, 100 and 86 on the first five tests, what score does she need on the sixth test to get an average of 93?

Gold Answer: First, identify the lowest score, which is 71, and remove it. Since Brinley wants a score of 93, we know her total score needs to be 93×5 (tests) = 465 points So far, Brinley has earned $89 + 92 + 100 + 86 = 367$ points The final test must be $465 - 367 = 98$ points.

CoT1: First, we need to find Brinley’s current average by adding up her five test scores and dividing by 5: $(89 + 71 + 92 + 100 + 86) / 5 = 87.6$ We know that Brinley wants her average to be 93, which means she needs to have a total of 6 test scores (since one will be dropped). The total of these 6 scores needs to be $93 \times 6 = 558$ To find out what score Brinley needs on the sixth test, we can subtract the total of her first five scores from the desired total: $558 - (89 + 71 + 92 + 100 + 86) = 120$ So Brinley needs to get a score of 120 on her sixth test to have an average of 93 after dropping her lowest score. ans = 120

CoT2: Answer: Brinley’s first five test scores add up to 438. To have an average of 93, Brinley needs a total of $93 \times 6 = 558$ points for all six tests. To get 558 points after five tests, Brinley needs to score $558 - 438 = 120$ points on the sixth test. ans = 120

PoT1:

```
# Python code, return ans
num_of_tests = 6
total_of_first_five_tests_scores = 89 + 71 + 92 + 100 + 86
desired_average = 93
lowest_score = min(89, 71, 92, 100, 86)
sum_of_best_five_scores = total_of_first_five_tests_scores - \
lowest_score
ans = (desired_average * num_of_tests) - sum_of_best_five_scores
# ans = 191.0(This line is not generated by LLM.)
```

E IMPLEMENTATION DETAILS OF EXTERNAL VERIFIERS

Throughout these experiments, our choice of the weaker LLM remains GPT-3.5-turbo. We only perform our experiment with CoT representation because the finetuned model we utilize is not suitable for understanding Python code.

The input to the verifier can be either the question alone or a concatenation of both the question and the answer produced by the weaker LLM. The question texts include the context, e.g., the goal and context in CREPE. The answer texts contain both intermediate thoughts and the final answer (e.g., an integer in mathematical reasoning). We have multiple generation results from the weaker LLM. To avoid excessive input length, we only employ the majority-voted answer and its intermediate thoughts in the test time. Specifically, in the test time, we extract the majority-voted answer from two CoT prompts with different demonstrations as this setting shows better performance in our method, and then randomly select one collection of intermediate thoughts that result in this majority-voted answer.

We still use task accuracy and relative cost as the metric. Note that only the question as input doesn’t need the answers from the weaker LLM necessarily. Therefore, the total cost C (Eq 1) for only the question as input becomes:

$$C = C^d + (1 - \mathbb{1}_{reject})C^w + \mathbb{1}_{reject} C^s. \quad (4)$$

To use an LLM as the verifier, we have designed two distinct prompts: one for “question difficulty prediction”, which solely relies on the questions (LLM-Q), and another for “answer evalua-

tion”, which incorporates both questions and answers (LLM-QA). The difficulty prediction (LLM-Q) prompt entails the LLM’s determination of whether a question necessitates redirection to the stronger LLM, based solely on the question’s complexity. If the LLM deems a question challenging for difficulty prediction, it triggers forwarding to the stronger LLM. The answer evaluation (LLM-QA) prompt is to give the LLM both the question and the answer. We then let the LLM rethink the solving steps and generate feedback to evaluate whether the solution is correct. If not, the question will be transferred to the stronger LLM. Both prompting methods are implemented with $M = 8$ shots of demonstrations, and we keep the number of “easy” and “hard” questions or the number of “correct” and “incorrect” answers balanced. We set the sampling path as $K = 20$ and temperature as $T = 0.4$ to sample the major answer and then use the self-consistency (SC) to decide the final “difficulty” or “correctness of the answers”. The reproducible prompts are in the Appendix L. Notably, the cost for decision C^d in Eq. 1 is no longer 0 with an external verifier. For the LLM as a verifier, we use GPT-3.5-turbo as it’s much cheaper than GPT-4. We count the number of tokens used for the approaches LLM-Q and LLM-QA.

For finetuning a smaller LLM (RoBERTa-base of Liu et al. (2020)) as the verifier, we similarly tested two variants, when the input contains only the question (Finetuned-Q) or when it includes both the question and the answers (Finetuned-QA), in which case we use a separation token to split them. The training data for fine-tuning the models come from the original training split of each benchmark dataset. Since DATE does not have a training split, we only performed experiments on GSM8k and CREPE. In each training instance, we have $K = 20$ generated responses and we label all of them. We label the generated responses where GPT-3.5-CoT-SC correctly predicts the answer as positive. Notably, the number of total training cases is the number of sampling paths $K = 20$ multiplied by the number of questions. To address the data imbalance issue (i.e., most of the training cases are positive), we perform simple under-sampling over the GSM8k. For CREPE, the resulting training set is relatively smaller, so we did not perform the under-sampling.

In the training process, we fine-tune RoBERTa-base with the learning rate 1×10^{-5} and batch size 32. We randomly split 10% of the training data to be the development set. The performance of the best accuracy over the development set is shown in Table 9. One notable approach, predicting “all positive”, is to assign a positive label to every example in the development set. We could learn that the external verifier could achieve better performance than “all positive” but not significant enough.

Dataset	Method	Dev
GSM8k	All positive	0.5
	Based on Question	0.592
	Based on Question and Answer	0.615
CREPE	All positive	0.715
	Based on Question	0.749
	Based on Question and Answer	0.812

Table 9: Performance of finetuning RoBERTa as verifier over the dev set

Table 10 shows the exact numerical results of those presented in Figure 6. We can learn that the external verifier could boost the performance but not that significant. Comparing LLM-Q with LLM-QA, an interesting observation is that the latter costs less than the former. This may seem counter-intuitive at first glance because LLM-QA includes a much longer input (a concatenation of question and answer) than LLM-Q (only question). However, we found out that for both approaches, their final costs are mainly determined by the number of cases transferred to the stronger LLM. That is, when including both the question and the answer as input, LLM-QA tends to trust the weaker LLM more frequently, leading to fewer cases being routed to the stronger LLM and hence the lower cost. However, both approaches do not yield accuracy as high as our MoT-based approaches.

F LLAMA2-13B AS THE WEAKER LLM

Although GPT-3.5-turbo is considered to be weaker compared to GPT-4, it is still more powerful and expensive than most of the open-source LLMs. In this section, we tested the performance of

Dataset	Method	Accuracy	Cost
GSM8k	GPT-3.5-turbo CoT-SC	0.842	0.111
	LLM-Q	0.884	0.253
	LLM-QA	0.862	0.135
	Finetuned-Q	0.892	0.340
	Finetuned-QA	0.882	0.262
	CoT-2D-Vote with threshold=0.7	0.945	0.408
	MoT-1D-Vote with threshold=0.5	0.917	0.235
	MoT-1D-Vote with threshold=0.7	0.951	0.425
	MoT-1D-Verify	0.947	0.401
GPT-4 CoT-SC	0.958	1.000	
DATE	GPT-3.5-turbo CoT-SC	0.676	0.107
	LLM-Q	0.750	0.367
	LLM-QA	0.707	0.142
	CoT-2D-Vote with threshold=0.65	0.862	0.489
	MoT-1D-Vote with threshold=0.65	0.889	0.469
	MoT-1D-Verify	0.875	0.449
	GPT-4 CoT-SC	0.886	1.000
CREPE	GPT-3.5-turbo CoT-SC	0.820	0.157
	LLM-Q	0.864	0.406
	LLM-QA	0.846	0.169
	Finetuned-Q	0.854	0.323
	Finetuned-QA	0.847	0.322
	CoT-2D-Vote with threshold=0.55	0.865	0.241
	MoT-1D-Vote with threshold=0.55	0.885	0.460
	MoT-1D-Verify	0.882	0.534
	GPT-4 CoT-SC	0.871	1.000

Table 10: Exact numerical results of various external verifiers. Specifically, because DATE does not have a training set, we do not report finetuned models’ performance on it.

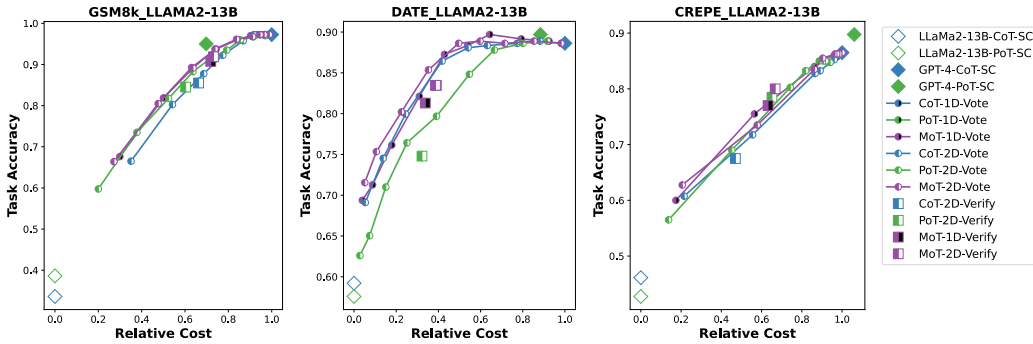


Figure 7: The performance of using LLAMA2-13B as weaker LLM.

LLAMA2-13B (Touvron et al., 2023), an open-source and easy-to-deploy LLM, as the weaker LLM and GPT-4 as the stronger LLM.

The result is shown in the Figure 7. Through experimentation, we observe that choosing LLAMA2-13B as a weaker LLM and applying our strategy on GSM8k and CREPE did not yield ideal results, but it performed well on the DATE dataset. On GSM8k and CREPE datasets, accuracy and cost change approximately linearly and we cannot achieve a comparable result with a lower cost, but on the DATE dataset, we obtain a curve similar to the curve using GPT-3.5-turbo as a weaker LLM. To analyze the cause, we have selected cases in MoT-1D-Vote whose agreement score s (Eq 2) is greater than the 0.6, 0.7, and 0.8. We report their task accuracy and proportion among all cases. A lower accuracy in those cases is indicative of reduced overall task performance, while a diminished proportion implies an increased cost.

The results are displayed in Table 11. Comparing LLAMA2-13B to GPT-3.5-turbo, we observe that LLAMA2-13B exhibits a significantly lower proportion. This suggests that LLAMA2-13B is less likely to provide consistent answers across multiple samplings. For example, only 16.3% of the cases have an agreement score greater than the 0.6 threshold on GSM8k in LLAMA2-13B but 77.3% in GPT-3.5-turbo. However, the difference in accuracy is slight in GSM8k and DATE. This outcome aligns with our hypothesis that the LLM can consistently answer easy questions but is hesitant to the hard questions. The lower proportions show that the number of easy questions for LLAMA2-13B is

far fewer than that in GPT-3.5-turbo as LLAMA2-13B cannot solve these complex reasoning tasks well. Therefore, LLAMA2-13B is not a suitable choice as a weaker LLM for particular tasks.

Threshold	Dataset	Weak LLM	Accuracy	Proportion
0.6	GSM8k	GPT-3.5-turbo	0.964	0.773
		LLAMA2-13B	0.923	0.163
	DATE	GPT-3.5-turbo	0.887	0.721
		LLAMA2-13B	0.824	0.461
	CREPE	GPT-3.5-turbo	0.916	0.722
		LLAMA2-13B	0.639	0.243
0.7	GSM8k	GPT-3.5-turbo	0.981	0.707
		LLAMA2-13B	0.946	0.093
	DATE	GPT-3.5-turbo	0.893	0.656
		LLAMA2-13B	0.893	0.328
	CREPE	GPT-3.5-turbo	0.934	0.655
		LLAMA2-13B	0.658	0.095
0.8	GSM8k	GPT-3.5-turbo	0.985	0.639
		LLAMA2-13B	1.000	0.070
	DATE	GPT-3.5-turbo	0.900	0.593
		LLAMA2-13B	0.933	0.203
	CREPE	GPT-3.5-turbo	0.956	0.588
		LLAMA2-13B	0.600	0.013

Table 11: The performance of cases where the percentage of the major answer under MoT-1D-Vote be greater than the threshold with LLAMA2-13B as the weaker LLM.

G CAN THE STRONGER LLM BENEFIT FROM THE WEAKER LLM’S HINTS?

Dataset	Method	Accuracy
GSM8k	W/o Hints	0.891
	W/ Hints	0.867
DATE	W/o Hints	0.892
	W/ Hints	0.910
CREPE	W/o Hints	0.774
	W/ Hints	0.727

Table 12: The GPT-4-CoT-SC accuracy over the instances that $A_1^{w'} \neq A_2^{w'}$ in the MoT-1D-Verify approach

We conducted an extended experiment to investigate whether progressive hints can further enhance our pipeline. Previous works have proven that progressive hints in multiple iterations can enhance the LLM’s reasoning performance (Zheng et al., 2023). For example, if the answer is 13 in the first iteration and 14 in the second one, the hint for the third iteration can be “*Hint: The answers may be close to 13, 14*”. In MoT-1D/2D-Verify, before the decision maker transfers a question to a stronger LLM, we already have two different answers generated by the weaker LLM. However, the answers from the weaker LLM will be discarded when using the stronger LLM. To incorporate the answers in the leveraging process of the stronger LLM, we follow the prior work (Zheng et al., 2023) to provide these two inconsistent answers as hints. For example, if the main answer is 5 with CoT and 12 with PoT, the hint for the stronger LLM can be “*Hint: The answers may be close to 5, 12*”. We perform the exploration of whether using these answers as hints can improve the stronger LLM’s performance. Prompts can be found in Appendix L. In the prompts, we acknowledge that the hints may not always be correct. In our exploration, we choose the data that will be sent to the stronger LLM by the MoT-1D-Verify approach. We then assess the GPT-4-CoT-SC accuracy of these cases, both with and without the use of hints, respectively. We conduct our experiments on GSM8k, DATE, and CREPE datasets.

The results in Table 12 reveal that progressive hints may not necessarily further enhance our pipeline. Adding progressive hints only yields a slight improvement on the DATE dataset, with diminishing returns on GSM8k and GREPE. The reason is that GPT-4 tends to be misled by the incorrect answer, especially when both answers from GPT-3.5-turbo are incorrect.

	accuracy	cost
w/o Batch Prompting	0.947	0.401
w Batch Prompting	0.924	0.325

Table 13: The GSM8k results with and without batch prompting

H CAN BATCHPROMPT FURTHER REDUCE THE COST?

Batch Prompting is a method to reduce costs by inputting multiple questions into LLM at once (Cheng et al., 2023a; Lin et al., 2023b). In this section, we explore whether our method can be used in conjunction with batch prompting. Our experiments are based on the basic batch prompting setup (Cheng et al., 2023a). We use `MoT-1D-Verify` for decision making. Specifically, we grouped a batch of 4 test questions into each API call of the weaker LLM, in addition to the original 8-shot demonstrations. The prompt is shown in Figure 39. Like in our previous experiments, we obtain multiple samples from running the weaker LLM, and the verification-based method (Eq 3) can then be leveraged independently for each test question. For example, if two of four test instances are rejected by the decision maker, we only feed the rejected two to the stronger LLM. We perform the experiment on the GSM8k dataset.

We show the result in Table 13. The results show that adding batch prompting can indeed reduce costs more. However, using batch prompting slightly affects the accuracy of both the weaker LLM and the stronger LLM. The result indicates that our method is orthogonal to the batch prompting approach.

I CALIBRATION ANALYSIS

In this section, we performed a calibration analysis comparing our approaches with baselines. For our voting-based method, we take n/K as the confidence score. For LLM as an external verifier, we followed the prior works’ prompting method (Kadavath et al., 2022) to let the LLM make the decision according to the textual contents.⁵ Following their setting, we consider two temperature setups, i.e., $T = 1$ and $T = 2$, and denote the variants as `LLM-QA T=1` and `LLM-QA T=2`, respectively. The prompts are shown in Table 31. We also note that, in the original method of Kadavath et al. (2022), the authors used the token probability of an LLM in predicting the answer to be “True” as its confidence score. Since we cannot obtain the token probability of GPT-3.5-turbo, the confidence score is obtained by sampling the True/False prediction K times and reporting the frequency of “True” (n/K). This will induce significantly more costs than our approach, but we adopted the strategy here for this calibration analysis. We perform the calibration analysis in the GSM8k dataset.

The calibration results on GSM8k are illustrated in Figure 8(Left). Our major observations are: (1) All decision-making methods yield a monotone calibration curve, implying that when they have higher confidence in a certain answer, the answer is generally more likely to be true. For our three vote-based approaches, this upward trend also supports our initial hypothesis that a question is easy if the weaker LLM exhibits consistency in its responses across multiple sampling paths; (2) `LLM-QA T=2` with a larger temperature can lead to the least calibrated decision maker; (3) However, there is no a significant difference among other approaches in terms of their calibration degree (though all our variants showed to be better than `LLM-QA`).

We wanted to note that achieving perfect calibration with n/K as the confidence score is not necessary for our task. To give an extreme example, let us consider the case that whenever an answer has $n/K > 0.5$, the answer from the weaker LLM is correct (accuracy 1.0). In this case, we would have a very poor calibration curve, but our pipeline would actually perform very well, because any majority-voted answer it accepts in this case is perfectly true. Another factor to note is the “size” of each bin in the calibration plot (e.g., the number of instances that have confidence scores ranging between 0.6 and 0.7). While this factor will also dramatically impact the LLM cascade performance in practice, it cannot be shown in the calibration plot.

⁵In our experiments in Section 3.5, the `LLM-QA` model is additionally prompted to explain its decision making, but we ignored this explanation following Kadavath et al. (2022) for the purpose of calibration analysis.




Figure 8: The left figure is for the calibration analysis of different approaches. The number of each method is the expected calibration error (ECE). The right figure plots the accuracy of the subset of answers satisfying n/K greater than the confidence score.

To give a more direct comparison among different decision-making approaches, we plot Figure 8(Right), where for each approach we report the accuracy of a subset of answers satisfying n/K greater than the confidence score, following the exact way how the approach will be applied in the LLM cascade pipeline. As such, curves on this plot have taken into account the “bin sizes” and can reflect their actual effectiveness. A similar analysis was also performed by Kadavath et al. (2022). We first observed that as the confidence score grows, there is a consistent upward trend in the subset accuracy for all approaches, which echoes the monotone curves in the calibration analysis. The higher position of MoT compared with other variants and baselines also explains its superior performance in the main experimental results.

J CAN OUR METHOD GENERALIZE TO FACTUALITY-BASED TASKS?

In prior studies, the efficacy of our approach in mathematical, symbolic, and causal reasoning has been evidenced. This section explores its potential to be generalized to factual reasoning tasks. Our investigation utilized StrategyQA (Geva et al., 2021), a factoid question answering dataset that requires implicit reasoning over facts. We employed an open-domain question-answering format, presenting only questions without context to the LLM. Our few-shot prompts are shown in Figure 27 and Figure 28.

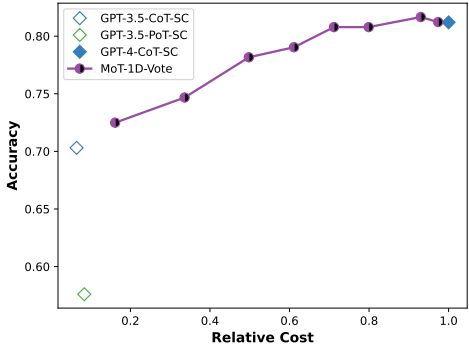


Figure 9: Performance of MoT-1D-Vote and baselines on the dev set of strategyQA.

The result is shown in Figure 9. It shows that our method is still effective on factual reasoning tasks. That is again owing to the fact that using different prompt representations could trigger different reasoning paths, which often results in more trustworthy answers when the two representations agree with each other.

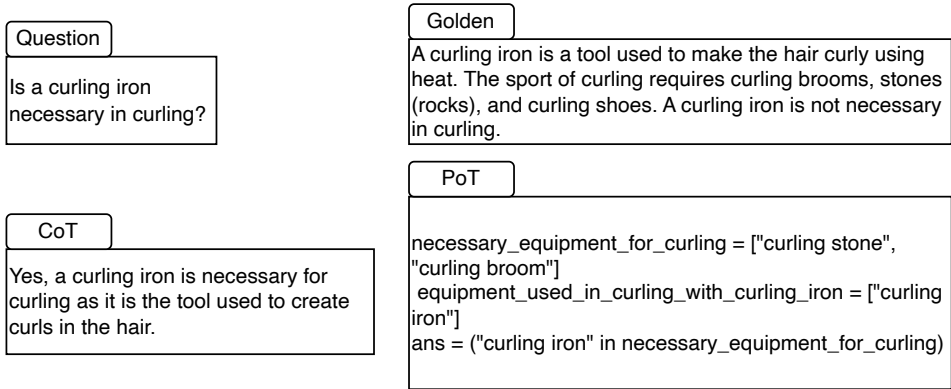


Figure 10: An example showing that MoT helps in factual reasoning tasks such as strategyQA.

An example is shown in Figure 10. For the question “*Is a curling iron necessary in curling?*”, the golden answer is “*No, curling is an ice sport and doesn’t need a curling iron*”. However, most of the CoT answers are “yes” with hallucinations about the concept “curling”. In contrast, most of the PoT answers are “No”. The PoT processes typically list the necessary equipment for curling, such as “curling stone” and “broom”, and then check if “curling iron” is on the list. By checking the consistency between CoT and PoT, MoT-1D-Vote is thus able to identify the incorrect or untrustworthy answer. It is worth noting that, judging from the results of Figure 9, PoT is not better than CoT, but the combination of CoT and PoT generates diverse thoughts and answers, instead of leaning towards one kind of thinking, thus reducing errors in factual reasoning. Therefore, we can still leverage consistency checking across MoT prompts in decision-making to check if the answer from the weaker LLM is trustworthy in factual-based reasoning tasks.

K LIMITATIONS AND FUTURE WORKS

Our approach is still subject to some limitations. Firstly, its applicability is confined to tasks where questions have clear answers, such as mathematical reasoning. Secondly, the utilization of our pipeline might lead to increased latency in cases demanding a stronger LLM. Thirdly, our method is based on the assumption that the intermediate steps could be expressed via different representations. In addition, if weaker LLM has overconfidence in some incorrect beliefs, our method will fail in these cases.

Considering these limitations, we identify some potential avenues for future research. In tackling the initial challenge, we could integrate other metrics, such as semantic similarity, to evaluate the consistency of the general textual generation tasks. To address the latency issue, it is essential to establish a framework for making decisions early for some obviously easy or hard questions. For the third one, we believe that how to represent the intermediate steps in a specific task relies on human designation. Recent research has shown some predefined programming interfaces effectively manage complex tasks, even those that appear unrelated to traditional code generation tasks, such as image understanding (Gupta & Kembhavi, 2023). If a task’s intermediate steps cannot be expressed directly via code, we can also introduce diverse answers by leveraging some programming interfaces/tools.

Additionally, due to CoT and PoT being suitable for different tasks, the answer’s correctness with different representations varies across tasks, e.g., the CoT answers are more reliable in GSM8k but the PoT answers are more reliable in DATE. The incorporation of learning algorithms to perform weighted voting for specific tasks yields a promising improvement.

L FULL PROMPTS

We show the prompts used for this paper according to Table 14. We show one example prompt in each case.

Prompt Type	Dataset	Prompt Table
CoT	GSM8k	prompt 15
	ASDIV	prompt 17
	TabMWP	prompt 19
	Date	prompt 21
	Navigate	prompt 23
	CREPE	prompt 25
PoT	GSM8k	prompt 16
	ASDIV	prompt 18
	TabMWP	prompt ??
	Date	prompt 22
	Navigate	prompt 24
	CREPE	prompt 26
LLM as External Verifier	GSM8k w/ Question	prompt 29
	GSM8k w/ Question and Answer	prompt 30
	DATE w/ Question	prompt 32
	DATE w/ Question and Answer	prompt 33
	CREPE w/ Question	prompt 34
	CREPE w/ Question and Answer	prompt 35
Learning from hints	GSM8k	prompt 36
	Date	prompt 37
	CREPE	prompt 38

Table 14: The index of each prompt.

GSM8k CoT

Complete the text, start with 'Answer' and the last line starts with
 \rightarrow 'ans = '.

Question: Manny had 3 birthday cookie pies to share with his 24
 \rightarrow classmates and his teacher, Mr. Keith. If each of the cookie pies
 \rightarrow were cut into 10 slices and Manny, his classmates, and Mr. Keith
 \rightarrow all had 1 piece, how many slices are left?

Answer: There is a total of $3 \times 10 = \ll 3 \times 10 = 30 \gg 30$ cookie slices. There
 \rightarrow are $24 + 1 + 1 = \ll 24 + 1 + 1 = 26 \gg 26$ people who ate the cookie pieces.
 \rightarrow There is $30 - 26 = \ll 30 - 26 = 4 \gg 4$ cookie slices left. ans = 4
 ...

Table 15: GSM8k CoT task demonstrations

GSM8k PoT

```

# Question: Manny had 3 birthday cookie pies to share with his 24
  ↳ classmates and his teacher, Mr. Keith. If each of the cookie pies
  ↳ were cut into 10 slices and Manny, his classmates, and Mr. Keith
  ↳ all had 1 piece, how many slices are left?
# Python code, return ans
num_cookie_pies = 3
num_slices_per_cookie_pie = 10
total_people = 24 + 1 + 1
total_slices_needed = total_people * 1
total_slices = num_cookie_pies * num_slices_per_cookie_pie
ans = total_slices - total_slices_needed
...

```

Table 16: GSM8k PoT task demonstrations

ASDIV CoT

Please answer the math question and return a number as the answer.

Question: Isabella’s hair is 18 inches long. By the end of the year her
 ↳ hair is 24 inches long. How much hair did she grow?

Answer: Isabella’s hair is 18 inches long and 24 inches long.
 So the length she grow in one year is $24 - 18 = 6$
 ans = 6
 ...

Table 17: ASDIV CoT task demonstrations

ASDIV PoT

```

# Question: Isabella’s hair is 18 inches long. By the end of the year
  ↳ her hair is 24 inches long. How much hair did she grow?
# Python code, return ans
Isabella_hair_before = 18
Isabella_hair_after_year = 24
hair_growth = Isabella_hair_after_year - Isabella_hair_before
ans = hair_growth
...

```

Table 18: ASDIV PoT task demonstrations

TabMWP CoT

Read the following table regarding "Coins" and then answer a question:

Name | Number of coins

Braden | 76

Camilla | 94

Rick | 86

Mary | 84

Hector | 80

Devin | 83

Emily | 82

Avery | 87

Question: Some friends discussed the sizes of their coin collections.

↪ What is the mean of the numbers?

Explain: Let's think step by step.

The numbers of coins of each one are in [76, 94, 86, 84, 80, 83, 82, 87].

So the mean of the numbers is $(76+94+86+84+80+83+82+87)/8 = 648/8 = 81$

Answer: 81

...

Table 19: TabMWP CoT task demonstrations

StrategyQA PoT

Question: Are more people today related to Genghis Khan than Julius

↪ Caesar?

Python code, return ans

children_num_of_Julius_Caesar = 3

children_num_of_Genghis_Khan = 6

ans = (children_num_of_Genghis_Khan > children_num_of_Julius_Caesar)

Table 20: StrategyQA PoT task demonstrations

DATE CoT

Q: Today is Christmas Eve of 1937. What is the date tomorrow in

↪ MM/DD/YYYY?

Explain: Today is the Christmas Eve of 1937, so today is 12/24/1937.

Today is 12/24/1937, the date tomorrow is 12/25/1937.

A: 12/25/1937

...

Table 21: DATE CoT task demonstrations

DATE PoT

```

# Write Python Code to solve the following questions .
from datetime import date , timedelta
from dateutil.relativedelta import relativedelta

# Q: Today is Christmas Eve of 1937. What is the date tomorrow in
    ↪ MM/DD/YYYY?
# today is Christmas Eve of 1937, then today is 12/24/1937
today = date(1937, 12, 24)
# tomorrow
date_tomorrow = today + relativedelta(days=1)
# The answer formatted with %m/%d/%Y is
ans = date_tomorrow.strftime('%m/%d/%Y')
...

```

Table 22: DATE PoT task demonstrations

Navigate CoT

Following these instructions , if we return to the starting point , return
 ↪ 'yes'; else return 'no'.

Instruction: Take 1 step. Take 2 steps. Take 3 steps. Turn around. Take
 ↪ 6 steps. Turn left.

Explain: start_position = [0, 0], assume the start face is to x-positive

1. Take 1 step. The current position is [1, 0]
2. Take 2 steps. The current position is [3, 0]
3. Take 3 steps. The current position is [6, 0]
4. Turn around. The face is to x-negative in the following steps.
5. Take 6 steps. The current position is [0, 0]
6. Turn left. The face is y-negative in the following steps.

After all the steps , the position is [0, 0], the same as the starting
 ↪ point.

Answer: yes
 ...

Table 23: Navigate CoT task demonstrations

Navigate PoT

Following these instructions, if we return to the starting point, return
 ↪ 'yes'; else return 'no'.

```
def left_rotate(face_direct):
    face_direct = tuple(face_direct)
    mapping_dict = {(1, 0): (0, 1), (0, 1): (-1, 0), (-1, 0): (0, -1),
                    ↪ (0, -1): (1, 0)}
    return list(mapping_dict[face_direct])

def right_rotate(face_direct):
    face_direct = tuple(face_direct)
    mapping_dict = {(1, 0): (0, -1), (0, 1): (1, 0), (-1, 0): (0, 1),
                    ↪ (0, -1): (-1, 0)}
    return list(mapping_dict[face_direct])

def around_rotate(face_direct):
    face_direct = tuple(face_direct)
    mapping_dict = {(1, 0): (-1, 0), (0, 1): (0, -1), (-1, 0): (1, 0),
                    ↪ (0, -1): (0, 1)}
    return list(mapping_dict[face_direct])

def move_steps(current_position, face_direct, step, step_direct=''):
    new_list = []
    if step_direct == 'left':
        face_direct = left_rotate(face_direct)
    elif step_direct == 'right':
        face_direct = right_rotate(face_direct)
    elif step_direct == 'backward':
        face_direct = around_rotate(face_direct)
    for i in range(len(current_position)):
        new_list.append(current_position[i] + face_direct[i] * step)
    return new_list

# Instruction: Take 1 step. Take 2 steps. Take 3 steps. Turn around.
↪ Take 6 steps. Turn left.
# Python code, return ans
start_position = [0, 0]
current_position = start_position
# assume the start face is to x-positive
face_direct = [1, 0]
# Take 1 step.
current_position = move_steps(current_position, face_direct, 1)
# Take 2 steps.
current_position = move_steps(current_position, face_direct, 2)
# Take 3 steps.
current_position = move_steps(current_position, face_direct, 3)
# Turn around.
face_direct = around_rotate(face_direct)
# Take 6 steps.
current_position = move_steps(current_position, face_direct, 6)
# Turn left. Now the face is to y-positive
face_direct = left_rotate(face_direct)
ans = 'yes' if current_position == start_position else 'no'
...
```

Table 24: Navigate PoT task demonstrations

CREPE CoT

Answer the question. The final answers are "more likely", "equally
↪ likely" and "less likely".

Goal: Wash sneakers

Current_Context: Brush off dirt from the surface of the sneakers. Remove
↪ shoelaces. Rinse the shoelaces in soapy water and air dry. Apply
↪ mild detergent to the sneakers and rub gently.

Question: What's the likelihood that The sneakers are damp?

Explain:

For step 1, after brushing off dirt, the sneakers are damp is "equally
↪ likely" because there is no change for the dampness of the shoes.

For step 2, after removing shoelaces, the sneakers are damp is "equally
↪ likely" because there is no change for the dampness of the shoes.

For step 3, after rinsing the shoelaces in soapy water and air dry, the
↪ sneakers are damp is "equally likely" because we removed the
↪ shoelace from the sneaker so we don't change the dampness of
↪ sneaker.

For step 4, after applying mild detergent to the sneakers and rubbing
↪ gently, the sneakers are damp is "more likely" because we put mild
↪ detergent to the sneakers.

Therefore, in the final step, the event is "more likely".

Answer: more likely

...

Table 25: CREPE CoT task demonstrations

CREPE PoT

```

# Answer the question. The answers are "more likely", "equally likely"
  ↳ and "less likely".
# Goal: Wash sneakers
# Current_Context: Brush off dirt from the surface of the sneakers.
  ↳ Remove shoelaces. Rinse the shoelaces in soapy water and air dry.
  ↳ Apply mild detergent to the sneakers and rub gently.
# Question: What's the likelihood that The sneakers are damp?
# Python code
class Wash_Sneakers():
    # Init from Current_Context
    # Brush off dirt from the surface of the sneakers
    # Remove shoelaces
    # Rinse the shoelaces in soapy water and air dry
    # Apply mild detergent to the sneakers and rub gently.
    def __init__(self):
        self.event0 = None # event0 is the likelihood that The sneakers
          ↳ are damp.
    def brush_off_dirt(self):
        # After brushing off dirt, event0 becomes "equally likely"
        self.event0 = "equally likely"
    def remove_shoelaces(self):
        # After removing shoelaces, event0 becomes "equally likely"
        self.event0 = "equally likely"
    def rinse_shoelace_in_water(self):
        # After rinsing in water, event0 becomes "equally likely"
        self.event0 = "equally likely"
    def apply_detergent_sneaker(self):
        # After rinsing in water, event0 becomes "more likely"
        self.event0 = "more likely"
    def call_all_funcs_in_order(self):
        self.brush_off_dirt()
        self.remove_shoelaces()
        self.rinse_shoelace_in_water()
        self.apply_detergent_sneaker()
c = Wash_Sneakers()
c.call_all_funcs_in_order()
ans = c.event0
...

```

Table 26: CREPE PoT task demonstrations

StrategyQA CoT

Question: Are more people today related to Genghis Khan than Julius
 ↳ Caesar?
 Explain: Julius Caesar had three children.
 Genghis Khan had sixteen children.
 The children number of Genghis Khan is larger than Julius Caesar.
 Answer: Yes

Table 27: StrategyQA CoT task demonstrations

StrategyQA PoT

Question: Are more people today related to Genghis Khan than Julius
 ↳ Caesar?
 # Python code, return ans
 children_num_of_Julius_Caesar = 3
 children_num_of_Genghis_Khan = 6
 ans = (children_num_of_Genghis_Khan > children_num_of_Julius_Caesar)

Table 28: StrategyQA PoT task demonstrations

GSM8k Question Difficulty Prediction

Predict the hardness level of the questions.
 Question: Shawna’s workout goal is 30 situps. On Monday, Shawna was only
 ↳ able to do 12 situps, so she decided that she would make up for
 ↳ the rest on Tuesday. However, she was only able to do 19 situps on
 ↳ Tuesday. How many situps would Shawna have to do on Wednesday to
 ↳ meet her minimum goal and make up for the ones she didn’t do?
 Level: Hard
 ...

Table 29: GSM8k question difficulty prediction task demonstrations

GSM8k Question and Answer Difficulty Prediction

Generate feedback and predict if the generated answer is trustful.
 Question: Shawna’s workout goal is 30 situps. On Monday, Shawna was only
 ↳ able to do 12 situps, so she decided that she would make up for
 ↳ the rest on Tuesday. However, she was only able to do 19 situps on
 ↳ Tuesday. How many situps would Shawna have to do on Wednesday to
 ↳ meet her minimum goal and make up for the ones she didn’t do?
 Answer: Shawna’s goal is 30 situps, and she has already done $12 + 19 =$
 ↳ $\ll 12+19=31 \gg 31$ situps. \n To meet her goal, she needs to do $30 - 31$
 ↳ $= \ll 30-31=-1 \gg -1$ situps. \n Since she can’t do negative situps, she
 ↳ doesn’t need to do any more situps to meet her goal. \n ans = 0
 Feedback: The answer is incorrect. The goal is 30 situps everyday. So
 ↳ the total is $3*30 = 90$ rather than 30.
 Trustful: No
 ...

Table 30: GSM8k question and answer difficulty prediction task demonstrations

Calibration

Question: Shawna's workout goal is 30 situps. On Monday, Shawna was only
→ able to do 12 situps, so she decided that she would make up for
→ the rest on Tuesday. However, she was only able to do 19 situps on
→ Tuesday. How many situps would Shawna have to do on Wednesday to
→ meet her minimum goal and make up for the ones she didn't do?
Proposed Answer: Shawna's goal is 30 situps, and she has already done 12
→ + 19 = $\ll 12+19=31 \gg 31$ situps. \n To meet her goal, she needs to do
→ $30 - 31 = \ll 30-31=-1 \gg -1$ situps. \n Since she can't do negative
→ situps, she doesn't need to do any more situps to meet her
→ goal. \n ans = 0
Is the proposed answer:
(A) True
(B) False
The proposed answer is: B

Table 31: Calibration task demonstrations

DATE Question Difficulty Prediction

Predict the hardness level of the questions.
Question: Today is Christmas Eve of 1937. What is the date tomorrow in
→ MM/DD/YYYY?
Level: Easy
...

Table 32: DATE question difficulty prediction task demonstrations

DATE Question and Answer Difficulty Prediction

Generate feedback and predict if the generated answer is trustful.
Question: Today is Christmas Eve of 1937. What is the date tomorrow in
→ MM/DD/YYYY?
Explain: Today is the Christmas Eve of 1937, so today is 12/24/1937.
Today is 12/24/1937, the date tomorrow is 12/25/1937.
A: 12/25/1937
Feedback: The answer is correct.
Trustful: Yes
...

Table 33: DATE question and answer difficulty prediction task demonstrations

CREPE Question Difficulty Prediction

Predict the hardness level of the questions.

Goal: Making Matcha Green Tea the Traditional Way

Current_Context: Boil 3/4 cup (180 ml) of water and pour it into your

- ↪ matcha bowl. Whisk the hot water with a chasen. Discard the hot
- ↪ water from the bowl. Scoop 2 tsp (1.5 heaping teaspoons or 10 g)
- ↪ of matcha into a fine mesh strainer. Sift the matcha into your
- ↪ empty, dry bowl. Pour boiling water into a teacup. Add a small
- ↪ amount of hot water into the matcha bowl and whisk it.

Question: What's the likelihood that I drink the matcha?

Level: Easy

...

Table 34: CREPE question difficulty prediction task demonstrations

CREPE Question and Answer Difficulty Prediction

Generate feedback and predict if the generated answer is trustful.

Goal: Boating in lake

Current_Context: Rent a boat. Find a lake which allows boating. Drive to

- ↪ the lake and set the boat in lake.

Question: What's the likelihood that the boat is in lake?

Explain: For step 1, after renting a boat, that the boat is in the lake

- ↪ is "less likely" because the boat is not yet in the lake.

For step 2, after finding a lake which allows boating, that the boat is

- ↪ in the lake is "equally likely" because we have found a lake but
- ↪ the boat is not yet in the lake.

For step 3, after driving to the lake and setting the boat in the lake,

- ↪ that the boat is in the lake is "more likely" because we have set
- ↪ the boat in the lake.

Therefore, in the final step, the event is "more likely".

Feedback: The answer is correct.

Trustful: Yes

...

Table 35: CREPE question and answer difficulty prediction task demonstrations

GSM8k Hints

Given the hints(may not be correct), answer the question.

Question: Tara has a shoebox that is 4 inches tall and 6 inches wide.

- ↪ She puts a square block inside that is 4 inches per side. How many
- ↪ square inches of the box are left uncovered?

Hints: The answer may be near to 24 or 8.

Answer: The shoebox is 24 square inches because $4 \times 6 = \ll 4*6=24 \gg 24$

The block is 16 square inches because $4 \times 4 = \ll 4*4=16 \gg 16$

There are 8 square inches uncovered because $24 - 16 = \ll 24-16=8 \gg 8$

ans = 8

...

Table 36: GSM8k with hints task demonstrations

DATE Hints

Given the hints (may not be correct), answer the question.
Question: Jane was born on the last day of February in 2000. Today is
 ↔ her 16-year-old birthday. What is the date 24 hours later in
 ↔ MM/DD/YYYY?
Hints: The answer may be near 02/29/2016 or 03/01/2016.
Answer: Jane was born on the last day of February in 2000, so she born
 ↔ on 02/29/2000.
Today is her 16-year-old birthday. So today is 02/29/2016.
The date 24 hours later than today is 03/01/2016.
A: 03/01/2016
...

Table 37: DATE with hints task demonstrations

CREPE Hints

Given the hints (which may not be correct), answer the question with
 ↔ "more likely", "equally likely" or "less likely".
Goal: Buy a flight ticket from SF to Hawaii.
Context: Open google flights and check the flights availability during
 ↔ the dates wanted. Choose non-stop flights and check if the time
 ↔ slots work. After deciding right time slots and price, proceed to
 ↔ buy tickets. Got an email confirmation of reservation.
Question: Jane was born on the last day of February in 2000. Today is
 ↔ her 16-year-old birthday. What is the date 24 hours later in
 ↔ MM/DD/YYYY?
Hints: The answer may be near "more likely" or "less likely".
Explain: For step 1, after checking flights availability, that I can
 ↔ refund the ticket if it is refundable is "equally likely" because
 ↔ I haven't purchased the ticket yet.
For step 2, after choosing non-stop flights and checking if the time
 ↔ slots work, that I can refund the ticket if it is refundable is
 ↔ "equally likely" because I haven't purchased the ticket yet.
For step 3, after deciding on the right time slots and price and
 ↔ proceeding to buy tickets, that I can refund the ticket if it is
 ↔ refundable is "less likely" because I have already purchased the
 ↔ ticket.
For step 4, after getting an email confirmation of reservation, that I
 ↔ can refund the ticket if it is refundable is "more likely" because
 ↔ I can check the terms and conditions of the ticket and see if it
 ↔ is refundable or not.
Therefore, in the final step, the event is "more likely".
Answer: more likely
...

Table 38: CREPE with hints task demonstrations

Batch Prompting

```

{
Question1: A robe takes 2 bolts of blue fiber and half that much white
  ↳ fiber. How many bolts in total does it take?
Question2: Josh decides to try flipping a house. He buys a house for
  ↳ $80,000 and then puts in $50,000 in repairs. This increased the
  ↳ value of the house by 150%. How much profit did he make?
Question3: Every day, Wendi feeds each of her chickens three cups of
  ↳ mixed chicken feed, containing seeds, mealworms and vegetables to
  ↳ help keep them healthy. She gives the chickens their feed in
  ↳ three separate meals. In the morning, she gives her flock of
  ↳ chickens 15 cups of feed. In the afternoon, she gives her
  ↳ chickens another 25 cups of feed. How many cups of feed does she
  ↳ need to give her chickens in the final meal of the day if the size
  ↳ of Wendi's flock is 20 chickens?
Question4: Kylar went to the store to buy glasses for his new apartment.
  ↳ One glass costs $5, but every second glass costs only 60% of the
  ↳ price. Kylar wants to buy 16 glasses. How much does he need to pay
  ↳ for them?

Answer1: It takes  $2/2=1$  bolt of white fiber\nSo the total
  ↳ amount of fabric is  $2+1=3$  bolts of fabric
ans = 3
Answer2: The cost of the house and repairs came out to
  ↳  $80,000+50,000=\$130,000$ \nHe increased the
  ↳ value of the house by  $80,000*1.5=120,000$ \nSo
  ↳ the new value of the house is
  ↳  $120,000+80,000=\$200,000$ \nSo he made a
  ↳ profit of  $200,000-130,000=\$70,000$ 
ans = 70000
Answer3: If each chicken eats 3 cups of feed per day, then for 20
  ↳ chickens they would need  $3*20=60$  cups of feed per
  ↳ day.\nIf she feeds the flock 15 cups of feed in the morning, and
  ↳ 25 cups in the afternoon, then the final meal would require
  ↳  $60-15-25=20$  cups of chicken feed.
ans = 20
Answer4: The discount price of one glass is  $60/100 * 5 =$ 
  ↳  $\$3$ .\nIf every second glass is cheaper, that means
  ↳ Kylar is going to buy  $16 / 2 = 8$  cheaper glasses.\nSo
  ↳ for the cheaper glasses, Kylar is going to pay  $8 * 3 =$ 
  ↳  $\$24$ .\nAnd for the regular-priced glasses, Kylar will pay
  ↳  $8 * 5 = \$40$ .\nSo in total Kylar needs to pay  $24 + 40 =$ 
  ↳  $\$64$  for the glasses he wants to buy.
ans = 64
}
...

```

Table 39: Batch Prompting task demonstrations