From Deferral to Learning:
Online In-Context Knowledge Distillation for LLM Cascades

Abstract

Standard *LLM Cascades* improve efficiency by *deferring* difficult queries from weak to strong models. However, these systems are typically static: when faced with repeated or semantically similar queries, they redundantly consult the expensive model, failing to adapt during inference. To address this, we propose *Inter-Cascade*, an online, interactive framework that transforms the strong model from a temporary helper into a long-term teacher. In our approach, when the strong model resolves a deferred query, it generates a generalized, reusable problem-solving *strategy*. These strategies are stored in a dynamic repository and retrieved via similarity matching to augment the weak model's context for future queries. This enables the weak model to "learn" on the job without expensive parameter fine-tuning. We theoretically show that this mechanism improves the weak model's confidence calibration. Empirically, Inter-Cascade outperforms standard cascades on multiple benchmarks, improving weak model and overall system accuracy by up to 33.06% and 6.35%, while reducing strong model calls by up to 48.05% and saving fee by up to 49.63%. Inter-Cascade demonstrates the effective in-context knowledge transfer between LLMs, and provides a general, scalable framework applicable to both open-source and API-based LLMs.

1 Introduction

Large Language Models (LLMs) demonstrate remarkable performance across a wide range of generation and reasoning tasks. Generally, performance scales with model size (Kaplan et al., 2020), creating a fundamental trade-off: larger models are more capable but significantly more expensive and slower. To address this, the *LLM Cascade* paradigm has emerged as a standard solution, where weaker (cheaper) models handle routine queries and *defer* only uncertain or complex cases to stronger (expensive) models (Chen et al., 2024b).

However, current cascade systems suffer from a critical "memoryless" limitation. Standard *deferral functions* are optimized offline and remain static during deployment (Shen et al., 2024; Jung et al., 2025). When the strong model resolves a difficult query, its expensive reasoning process is discarded immediately after the answer is returned. Consequently, the system fails to learn from its own previous operations. As noted in recent position paper from NVIDIA (Belcak et al., 2025), real-world query streams often exhibit a "similarity phenomenon," containing repeated or semantically similar tasks (e.g., slight variations of math problems in GSM-Plus (Li et al., 2024).) Faced with these recurring patterns, static LLM cascades redundantly consult the strong model for every instance, leading to a substantial waste of computation and tokens. While fine-tuning the weak model could theoretically solve this, it is often prohibitively expensive, slow, or impossible for API-based models.

To bridge this gap, we propose moving from static deferral to adaptive learning. We argue that the strong model should not merely serve as a temporary "backup" but as a long-term "teacher." To achieve this, we introduce *Inter-Cascade*, an online framework that transforms the interaction between LLMs. Unlike simple caching which only memorizes specific answers, Inter-Cascade extracts generalized problem-solving *strategies* from the strong model's reasoning. By retrieving and injecting these strategies into the weak model's context via similarity matching, we realize online in-context knowledge distillation. This enables the weak model to

"learn" on the job, dynamically improving its local success rate for future similar queries without parameter updates.

Our approach advances the concepts of In-Context Learning (ICL) (Dong et al., 2024), few-shot prompting (Parnami & Lee, 2022) and Retrieval-Augmented Generation (RAG) (Lewis et al., 2020). While traditional ICL relies on fixed demonstrations and RAG typically queries static, human-curated databases, Inter-Cascade builds a self-evolving strategy repository autonomously. The "corpus" is generated by the strong model and curated by the system's own interaction history, requiring no human intervention. This creates a closed-loop system where the weak model continuously distills "wisdom" from the strong model to handle increasingly complex tasks locally.

Primary contributions. Our contributions are as follows: (1) We propose Inter-Cascade, a general and modular framework for online interactive LLM Cascades. It allows the strong model to "teach" the weak model via a similarity-based strategy repository, effectively implementing widely applicable in-context knowledge transfer for both open-source and API-based models. Inter-Cascade is designed as a universal booster that also works for existing LLM Cascade methods. (2) We provide a theoretical framework proving that integrating strong-model strategies improves the weak model's confidence calibration. We show that this mechanism allows the weak model to more accurately assess its own competence, thereby reducing unnecessary deferrals while maintaining safety bounds. (3) We show empirically that compared to state-of-the-art cascades (Jung et al., 2025), Inter-Cascade improves the weak model's accuracy by up to 33.06% and overall system accuracy by up to 6.35%. Crucially, it reduces calls to the strong model by up to 48.05%, translating to significant cost savings (up to 49.63%) while strictly adhering to risk tolerance guarantees.

2 Related Work

LLM Cascades and learning to defer. LLM cascades route queries across models of different cost and capability using confidence-based deferral policies, aiming to balance quality and compute (Chen et al., 2024a). Recent work explores token-level deferral and post-hoc routing functions (Shen et al., 2024; Rayan & Tewari, 2025), learned routers that decide *before* invoking a stronger model (Ong et al., 2025), and cost-aware extensions such as early discarding or rational tuning (Zellinger et al., 2025; Zellinger & Thomson, 2025). These lines build on learning-with-reject frameworks (Chow, 1957; 1970; Madras et al., 2018; Mozannar & Sontag, 2020; Wu et al., 2025). However, most deployed cascades remain largely static after training: similar hard queries can repeatedly trigger strong-model calls without transferring knowledge to the weak model.

Distillation and retrieval-augmented generation. Knowledge distillation transfers capabilities from a strong teacher to a weaker student, typically via (re)training with soft targets or intermediate supervision (Hinton et al., 2015; Romero et al., 2015). RAG methods instead augment generation with non-parametric memory, usually retrieving from a fixed external corpus (Lewis et al., 2020) or from human-chatbot interaction histories for personalization (Zhang et al., 2025; Mo et al., 2025). Inter-Cascade connects these directions: when the weak model defers, the strong model produces reusable strategies that are stored and later retrieved to guide future weak-model attempts, yielding an *online, in-context distillation* mechanism at inference time that is complementary to classical distillation and RAG, without parameter updates and human involvement.

Other related topics There are also a weak model and strong model in *Speculative decoding* (Leviathan et al., 2023; Narasimhan et al., 2025), where the weak model works as a answer draft while the strong model works as a verifier to speed up the generation compared to only using strong model. However, in Inter-Cascade, Strong LLM is called only when the Weak LLM is unable to handle current query. *CombLM* (Ormazabal et al., 2023) and *LLM Debate* (Irving et al., 2018; Du et al., 2024; Estornell & Liu, 2024; Khan et al., 2024; Zhou et al., 2025) are other branches of works that also involve interaction between LLMs. CombLM integrates the logit distribution of two LLMs while LLM Debate requires different LLMs to argue and refine their initial answers and eventually reach consensus through multiple rounds of interaction. The key difference between Inter-Cascade and them is that Inter-Cascade let the Strong LLM and Weak LLM work in a sequential order can conduct early stop to save tokens.

A more comprehensive discussion of related works is in Appendix A.

3 Improving the LLM Cascade

Figure 1: (a) Pipeline of standard LLM Cascade systems. (b) Pipeline of Inter-Cascade. The unique components in Inter-Cascade are painted in orange. For the sake of clarity and readability, we only present the case of a two-LLM Inter-Cascade system and the scalable parts beyond two LLMs are rendered in a lighter color.

In this Section, We first describe the standard LLM Cascade (Chen et al., 2024a) and revisit the accuracy bound and calibration method for the deferral threshold Jung et al. (2025). We then introduce our proposed method Inter-Cascade and provide a theoretical framework to show when a weak model will be improved by a strong model's strategies. After that, we conduct diverse experiments to evaluate the performance of Inter-Cascade in Section 4.

3.1 Standard LLM Cascade

Figure 1(a) shows the general N-LLM Cascade system (Chen et al., 2024a). Each LLM $M_i : i \in [N]$ contains two key components. One is the *generation function* $g_i \colon \mathcal{Q} \to \mathcal{A}$, where $\mathcal{Q}$ is the space of queries and $\mathcal{A}$ is the space of answers. The other is *deferral function* $d_i \colon \mathcal{Q} \to \{0, 1\}$, which determines whether the i-th LLM will answer the query by itself ($d_i(q) = 0$) or defer it to the $(i+1)$-th LLM ($d_i(q) = 1$). Processing by the LLMs proceeds sequentially from M_1 to M_N. We define a partial order $\preccurlyeq_{\mathrm{wbc}}$ ("weaker but cheaper") to compare models (see Appendix B) and assume that in the cascade, $M_1 \preccurlyeq_{\mathrm{wbc}} M_2 \preccurlyeq_{\mathrm{wbc}} \dots \preccurlyeq_{\mathrm{wbc}} M_N$. For each query $q \in \mathcal{Q}$, the first LLM M_1 takes the query q and gives a final answer $g_1(q)$ if deferral function $d_1(q) = 0$, otherwise M_1 defers this query to the next LLM M_2 if $d_1(q) = 1$. If M_2 takes the query from M_1, it repeats the same process and so do the other LLMs except the last model M_N. As M_N doesn't have another LLM to offload the query, M_N discards this query if $d_N(q) = 1$. Recent studies propose different deferral functions d_i to meet the demands in different scenarios. We focus on the two-LLM case in the rest of this paper, as shown in Figure 1(b). We call M_1 the *Weak LLM* and M_2 the *Strong LLM*. One common

choice of deferral function is:

$$d_i(q) = \begin{cases} 0, & \text{if } c(q) \geq \lambda, \\ 1, & \text{otherwise,} \end{cases} \tag{1}$$

where $c : \mathcal{Q} \to [0,1]$ is a pre-defined or pre-trained "confidence" metric (usually defined in terms of the probability of output tokens) and λ is a confidence threshold, which is a hyperparameter that controls the trade-off between the system performance and cost.

Accuracy Guaranteed LLM Cascade. It is well known that LLMs suffer from systematic bias (Wang et al., 2024b; Thakur et al., 2025) and over-confidence (Xiong et al., 2024). To address this, Jung et al. (2025) propose a post-hoc calibration algorithm, which provably guarantees that with the derived λ,

$$P\left(g_i(q) = a_{\text{true}} \mid c(q) \geq \lambda\right) \geq 1 - \alpha \tag{2}$$

with probability at least $1 - \delta$, as proved in Theorem 1 of their work, where a_{true} is the ground-truth answer to query q. The risk tolerance α and error level δ are hyperparameters corresponding to the applications and users' demands. To instantiate this guarantee, they first used fixed-sequence testing (Bauer, 1991) procedure to find the largest threshold λ from a calibration set, such that $\Pr\left(g_i(q) = a_{\text{true}} \mid c(q) \geq \lambda\right)$ is exactly and tightly bounded. The procedure is summarized in Algorithm 1. They also extend the single-model guarantees to the full cascade; see Section 2 and Appendix A.2 in Jung et al. (2025)'s paper for details.

Algorithm 1 Calibrating Deferral Threshold λ (Jung et al., 2025)

Input: Calibration set $(q,a) \in D_{\text{cal}}$, confidence metric $c(\cdot)$, risk tolerance α, error level δ
Output: Threshold λ
 1: Initialize $\Lambda = \{0.999, 0.998, \ldots\}$ in decreasing order
 2: **for** $\lambda \in \Lambda$ **do**
 3:      $n(\lambda) \leftarrow \sum_{(q,a) \in D_{\text{cal}}} \mathbf{1}\{c(q) \geq \lambda\}$
 4:      $\hat{R}(\lambda) \leftarrow \frac{1}{n(\lambda)} \sum_{(q,a) \in D_{\text{cal}}} \mathbf{1}\{g_i(q) \neq a_{\text{true}} \wedge c(q) \geq \lambda\}$
 5:      $\hat{R}^+(\lambda) \leftarrow \sup\{R : \Pr[\text{Bin}(n(\lambda), R) \leq n(\lambda)\hat{R}(\lambda)] \geq \delta\}$
 6:      **if** $\hat{R}^+(\lambda) \leq \alpha$ **then return** λ

The general pipeline of LLM Cascade is shown in Figure 1(a). By using this LLM cascade diagram, the deferral function can keep "confident" queries on Weak LLMs and only send "uncertain" queries to Strong LLMs, dramatically reducing at most 82.5% usage of the strongest LLM as shown by Jung et al. (2025) while ensuring the error rate is bounded by α with probability at least $1 - \delta$.

3.2 Interactive LLM Cascade

LLM Cascade methods can be efficient and reliable although they still incur some waste in terms of tokens and latency as noted in Section 1. In particular, for workloads in which the Weak LLM is fed a similar or repeated queries for which it chooses to defer, the Strong LLM is called repeatedly to generate the same tokens. To address this issue, we propose *Inter-Cascade*. In Inter-Cascade, for both Weak LLM and Strong LLM, besides deferral function and generation function, we add the following components: *strategy generator* and *strategy repository*. In Strong LLM, we set up a *strategy generator* $h : \mathcal{Q} \to \mathcal{S}$, where $\mathcal{S}$ is the space of strategies. The strategy $s \in \mathcal{S}$ is defined as a sequence of tokens that contains the query and the answer of Strong LLM, together with a generalized ideas or tips to solve logically similar problems. To store those strategies, we construct a Strategy Repository called Repo. The Repo is accompanied by a *strategy matching function* $f : \mathcal{Q} \times \mathcal{Q}^N \to \mathcal{S}^k$, where N is the size of current Repo and k is a predefined hyperparameter that determines the number of strategies retrieved.

Strategy Repository. The Strategy Repository Repo is formally defined as a collection of query-strategy pairs: $\text{Repo} = (q_j, s_j)_{j=1}^N$ where $q_j \in \mathcal{Q}$ are previously solved queries and $s_j \in \mathcal{S}$ are their corresponding strategies generated by Strong LLM. The strategy matching f operates through multiple stages. The

repository is initialized as an empty set and dynamically updated: when the Strong LLM generates a strategy $s = h(q)$ for a new query q, the pair (q, s) is added to Repo, enabling future reuse through the matching function f.

For a query $q \in \mathcal{Q}$ that is sent to the Weak LLM, let $\mathrm{sim} \colon \mathcal{Q} \times \mathcal{Q} \to [0, 1]$ be a ranking function. Let the Top-k indices (sorted by decreasing similarity) be

$$\mathrm{TopIndex}(q) \triangleq (t_1, t_2, \ldots, t_k), \tag{3}$$

where each $t_i \in \{1, \ldots, N\}$ indexes an item in Repo and $\mathrm{sim}(q, q_{t_1}) \geq \cdots \geq \mathrm{sim}(q, q_{t_k}) \geq \mathrm{sim}(q, q_{\mathrm{else}})$. After ranking, these strategies with Top-k indexes are chosen to help the Weak LLM. Then the output of strategy matching function is $f(q, \mathrm{Repo}) \triangleq \{s^{t_i} \mid t_i \in \mathrm{TopIndex}(q)\}$.

Remark 3.1. Compared with finetuning or paying for Strong LLM, the cost of maintaining a Repo and running similarity-based matching algorithms are negligible. According to the estimate formula suggested by Johnson et al. (2021), conducting retrieval and Top-2 ranking on 1 million query embeddings, which are 384 dimensional vectors (the same size we used in experiments), only requires $0.2 - 0.8$ ms with $70 - 80$ MB GPU VRAM and $80 - 100$ MB RAM for long term storage. The demand can be easily fulfill on any PC or even phone, and imperceptible to human users.

Algorithm 2 Inter-Cascade Inference Pipeline

Input: Test set $\mathcal{T} = \{q_1, \ldots, q_I\} \subseteq \mathcal{Q}$; Weak LLM with deferral function d_1, generation function g_1, strategy repository $\mathrm{Repo} = \emptyset$; strategy matching function f; Strong LLM with deferral d_2, generator g_2, and strategy generator h.

    Deferral convention:
1:      $0 = $ handle locally, $1 = $ defer/forward.
2: **for** $i \leftarrow 1$ **to** I **do**
3:      $[s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}] \leftarrow f(q_i, \mathrm{Repo})$ ▷ **Retrieval**
4:      $q_i' \leftarrow [q_i, s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}]$ ▷ Concatenate strategies
5:      **if** $d_1(q_i') = 0$ **then**   ▷ **Weak LLM decision**
6:          $a_i \leftarrow g_1(q_i')$ ▷ Answer locally
7:      **else**
8:         **if** $d_2(q_i) = 0$ **then**   ▷ **Strong LLM decision**
9:           $s_{\mathrm{new}} \leftarrow h(q_i)$ ▷ **Strategy generation**
10:          $\mathrm{Repo} \leftarrow \mathrm{Repo} \cup \{(q_i, s_{\mathrm{new}})\}$ ▷ Send back strategy to Weak LLM and store
11:         generate answer $a_i \leftarrow g_2(x_i)$ ▷ Answer at Strong LLM
12:        **else**
13:         Discard current query q_i ▷ None of LLMs are confident to answer the query

---

Inter-Cascade Pipeline. The overall pipeline of Inter-Cascade is presented in Algorithm 2 and in Figure 1(b). For each query q, the Weak LLM first uses the strategy matching function $f(q, \mathrm{Repo})$ to find the most related strategies. The query and these strategies are then sent to deferral function. The augmented input is the prompt concatenation of query and strategies: $q' = [q, s^{t_1}, s^{t_2}, ..., s^{t_k}]$. If the Weak LLM's deferral function $d_1(q') = 0$, then final answer a for current query is $g_1(q')$. If $d_1(q') = 1$, the query q' is deferred to Strong LLM. Each time the query is sent to the Strong LLM, the deferral function in Strong LLM is called. If $d_2(q) = 0$, this query is discarded (since Strong LLM is the last model in two LLMs Cascade), otherwise $g_2(q)$ produces the answer and further, a new strategy is produced by $h(q)$. Then, the strategy will be stored into Repo. Given α and δ, we can derive the λ from Algorithm 1 and determine deferral function d_1 and d_2 as defined by Equation (1). Our algorithm can be extended to multi-LLM cases, the corresponding Algorithm 3 is shown in Appendix C.

Strategies Provide Improved Calibration. The Repo we build during the usage of the combination of LLMs collects the strategies of the Strong LLM and provides strategies to help the Weak LLM answer queries. With the help of strategies, the Weak LLM is able to solve the more challenging problems that appear frequently and be more aware of its correctness of answering the queries, leading better confidence.

However, it is not clear that how this increment in the accuracy and the quality of confidence could be preserved in the queries after the filtration. After all, all the queries, even to which the Weak LLM answers correctly would be deferred if the Weak LLM's confidence can not pass the threshold. Therefore, we present the following theories to estimate such an increment that would remain in the filtered queries.

To be specific, we first assume that, after adding strategies, under the same confidence threshold λ, the number of queries that pass the confidence threshold increases from $n(\lambda)$ to $bn \coloneqq n'(\lambda)$, $b \in [1, \infty)$, where $n(\lambda)$ is first defined in Algorithm 1. The number of wrongly answered queries before and after the help of strategies are denoted by x and ϵx, respectively, where $\epsilon \in (0, 1)$. We want to understand the potential benefit in terms of the reduction in risk α under the same error level δ. We do not change the threshold λ, which is the case when the strategy repository is enlarged during the running process of the Inter-Cascade. Theorem 3.2 states our main result. For the convenience of the statement, we define $\alpha(\epsilon, b)$ as the value of risk tolerance α when total number of queries that pass threshold is bn and incorrectly answered queries is ϵx.

Theorem 3.2. *Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ. Fix $\delta \in (0, 1)$ and an integer $n \geq 1$. For $x \in \{0, 1, \dots, n\}$, $\epsilon \in (0, 1]$, and $b \in [1, \infty)$. Then:*

(a) Decrease in value. $\alpha(\epsilon, b) \leq \alpha(1, 1)$ *when $\epsilon \in (0, 1]$ and $b \in [1, \infty)$.*

(b) Normal approximation for the amount of decrease. *Let $z \coloneqq \Phi^{-1}(1 - \delta)$, where Φ is the standard Normal cumulative distribution function. In the large–sample interior regime $n \to \infty$ with $x/n \to r \in (0, 1)$ and ϵ, b, δ held fixed, the decrease of the risk under the same level of tolerance satisfies*

$$\alpha(1, 1) - \alpha(\epsilon, b) = \left(\frac{x + 1}{n + 1} - \frac{\epsilon x + 1}{bn + 1} \right) + z \left[\sqrt{\frac{(x + 1)(n - x)}{(n + 1)^2(n + 2)}} - \sqrt{\frac{(\epsilon x + 1)(bn - \epsilon x)}{(bn + 1)^2(bn + 2)}} \right] + O\left(\frac{1}{n} \right). \tag{4}$$

The proof of this theorem is in Appendix E. Theorem 3.2 states that, when the δ and confidence threshold λ do not change, if more queries can pass the threshold, after combining with strategies and under certain conditions, we can ensure a smaller risk tolerance α in the guarantee of this inequality (2). That is, Inter-Cascade yields a higher success rate for Weak LLM.

Other than the case that λ remains unchanged, which is analyzed above, another case may be that when the users want the same number of queries to be covered by the Weak LLM during two rounds of queries (before and after adding strategies). This case considers the influence of a better Weak LLM on our pipeline. In this case, we instead assume that $n(\lambda) = n(\lambda')$, which ensures the same coverage of Weak LLM. We also show that we can ensure a smaller risk tolerance α when threshold becomes λ' while δ and number of queries that pass threshold remain unchanged. And the reduction in tolerance level $\alpha(1, 1) - \alpha(\epsilon, 1)$ is approximately linear to $1 - \epsilon$. The full statement of Theorem F.1 and the proof are shown in Appendix F.

4 Experiments

4.1 Benchmarks

We conduct a comprehensive evaluation on a suite of eight diverse benchmarks. To provide a focused analysis in the main text, we select four representative datasets spanning two primary categories: reasoning-intensive tasks (*GSM-Symbolic* (Mirzadeh et al., 2025), *GSM-Plus* (Li et al., 2024), *MetaMath* (Yu et al., 2024)) and factual knowledge tasks (*NASA-History-MCQ* (Fleith, 2025)). While the reasoning datasets evaluate Inter-Cascade's ability to handle structural variations, NASA-History-MCQ is featured specifically for its lack of explicit sample variants. This benchmark serves as a robustness test, allowing us to evaluate whether Inter-Cascade can still enhance efficiency and calibration in general scenarios where the "similarity phenomenon" is less pronounced. These selections highlight our method's adaptability across different difficulty levels. Full results for the remaining four benchmarks, including standard baselines (*GSM8K* (Cobbe et al., 2021), *BigBench Hard* (Suzgun et al., 2022)) and domain-specific tasks like legal benchmark *BarExamQA* (Zheng

et al., 2025) and medicine benchmark *MedMCQA* (Pal et al., 2022), are detailed in Appendix J, further demonstrating the framework's generalizability across broader scenarios without explicit query variants. The detailed descriptions of selected benchmarks are in Appendix H. The prompt template and an example problem for each benchmark are provided in Appendix L.

4.2 Experimental Settings

Inter-Cascade. On all benchmarks, *Gemini-2.0-flash* consistently outperforms *GPT-3.5-turbo* (see Table 1), and is therefore designated as the Strong LLM in our two-LLM Inter-Cascade, with *GPT-3.5-turbo* as the Weak LLM. We extract the normalized token probability from the LLM's output as confidence score $c(q)$ in following experiments. In preparation phase, with given risk tolerance α and error level δ, we derive desired confidence threshold λ from calibration set by following Algo. 1. Then deploy corresponding deferral functions d_i according to equation (1).

Our similarity-based strategy matching process on Repo works as follows. Given a new query, it is encoded into a vector and used to retrieve the top-k semantically similar queries from Repo. We employ the *all-MiniLM-L6-v2* transformer (Reimers & Gurevych, 2019) to produce 384-dimensional sentence embeddings and use the FAISS library (Douze et al., 2025) for efficient approximate nearest-neighbor search. FAISS returns the top-k vectors that minimize cosine distance, providing the Inter-Cascade with prior Strong LLM responses, including queries, answers and strategies, which can inform the Weak LLM's responses.

Table 1: Accuracies of the base LLMs on four benchmarks.

Benchmark	LLM	Accuracy
GSM-Symbolic	gpt-3.5-turbo	13.36%
	gemini-2.0-flash	69.36%
GSM-Plus	gpt-3.5-turbo	23.00%
	gemini-2.0-flash	73.57%
MetaMath	gpt-3.5-turbo	37.30%
	gemini-2.0-flash	79.70%
NASA-History	gpt-3.5-turbo	65.30%
	gemini-2.0-flash	78.80%

Inter-Cascade with No Strategies. To isolate the impact of strategy on the result in our pipeline, we only integrate the most similar questions and answers without the problem strategies to query.

Inter-Cascade with Random Strategies. To evaluate the impact of similarity-based retrieval on Repo, we randomly select the same number of strategies for each query, instead of choosing the top-k most similar queries.

Jung Proposed LLM Cascade. To evaluate the performance and effectiveness of the Inter-Cascade, we choose Jung et al. (2025)'s *Cascaded Selective Evaluation* as the baseline model. Its method for deriving confidence scores and thresholds provides a provable lower bound on the error risk and achieves state-of-the-art performance compared with other confidence-based LLM cascades.

4.3 Evaluation Metrics

We first define the notations used in our evaluation. Let T and U denote the total number of queries and the number of uncovered queries in a benchmark, respectively. Let N_w and N_s be the number of times the Weak and Strong LLMs are invoked, and let C_w and C_s denote the number of queries correctly answered by these models that also pass the confidence threshold. C_w^{total} denotes the total number of queries answered correctly by the Weak LLM. Let Tok_J and Tok_O be the tokens consumed by Jung's method and our proposed Inter-Cascade pipeline, and Cost_J and Cost_O denote their corresponding costs. The evaluation metrics are summarized in Table 2.

Table 2: Evaluation Metrics

Metric	Formula
Pipeline Accuracy	$(C_w + C_s)/(T - U)$
Strong LLM Call Rate	N_s/T
Weak LLM Accuracy	$C_w^{\text{total}}/(T - U)$
Weak Correct Accepted	$C_w/(T - U)$
Coverage Rate	$(T - U)/T$
Token Reduction	$(\text{Tok}_J - \text{Tok}_O)/\text{Tok}_J$
Cost Reduction	$(\text{Cost}_J - \text{Cost}_O)/\text{Cost}_J$

Table 3: Results across datasets using different pipelines. "Jung et al." denotes Jung et al.'s LLM-Cascade and "Our (Retrieval)" denotes the Inter-Cascade with similarity-based retrieval. The number of strategies is fixed at $k = 2$ for both Inter-Cascade settings. Metrics reported are Pipeline Accuracy (Pipeline Acc.), Strong LLM Call Rate (Strong Call), and Coverage Rate (Cov.). (a) GSM-Symbolic: For the Strong LLM, $\alpha_s = 0.2, \delta_s = 0.8, \lambda_s = 0.47$. For the Weak LLM, $\alpha_w = 0.6, \delta_w = 0.6, \lambda_w = 0.45$. (b) GSM-Plus: For the Strong LLM, $\alpha_s = 0.2, \delta_s = 0.8, \lambda_s = 0.51$. For the Weak LLM, $\alpha_w = 0.6, \delta_w = 0.6, \lambda_w = 0.48$. (c) MetaMath: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.4, \delta_w = 0.6, \lambda_w = 0.61$. (d) NASA-History: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.2, \delta_w = 0.7, \lambda_w = 0.87$.

Benchmark	Pipeline	Pipeline Acc. (%) ↑	Strong Call (%) ↓	Coverage (%)
GSM-Symb.	Jung et al.	66.04	59.37	86.31
	Our	**70.37**	**30.84**	**90.35**
GSM-Plus	Jung et al.	52.78	46.29	93.57
	Our	**58.31**	**32.44**	**94.79**
MetaMath (20K)	Jung et al.	65.21	49.26	100.00
	Our	**71.56**	**23.68**	100.00
NASA-Hist.	Jung et al.	71.88	26.68	100.00
	Our	**72.64**	**22.54**	100.00

4.4 Performance and Cost Analysis

Inter-Cascade vs. Jung's LLM Cascade. We evaluate our *Inter-Cascade* pipeline and Jung's method, as shown in Table 3. Our method outperforms Jung's, with a $4.33\% - 6.35\%$ increase in Pipeline Accuracy and $29.92\% - 51.93\%$ reduction in Strong LLM Call Rate on reasoning benchmarks. Crucially, on the NASA-History benchmark, which lacks structural variants, Inter-Cascade maintains high accuracy ($+0.76\%$) while still successfully reducing strong model calls by 15.5% (relative). These results indicate that *Inter-Cascade* pipeline is beneficial across different categories of tasks and particularly effective for reasoning-intensive tasks. Experiment results on extensive and diverse benchmarks are attached in Appendix J.

Impact of Inter-Cascade on Weak LLM. Having examined the overall pipeline improvements, including Pipeline Accuracy and Strong LLM Call Rate reduction, we now investigate how our proposed *Inter-Cascade* affects the Weak LLM. As shown in Table 4, our Weak LLM outperforms the Weak LLM in the other pipeline across all benchmarks. The improvements are particularly pronounced on reasoning benchmarks, with gains of 23.21%, 16.2%, and 33.06% on MetaMath, GSM-Plus, and GSM-Symbolic, respectively. On NASA-History, while the absolute accuracy gain is modest ($+0.48\%$), the Weak Correct Accepted Rate increases by 3.03% (from 55.37% to 58.40%). Importantly, improvements in the Weak LLM's accuracy contribute to the pipeline's performance only when the correctly answered queries exceed the confidence threshold. This is captured by the *Weak Correct Accepted* metric in Table 4, which represents the proportion of correctly answered queries that surpass the Weak LLM's threshold. The observed increase in Weak Correct Accepted shows Strong LLM's strategies helped the weak model better calibrate its confidence, validating our theoretical claim that

Table 4: Results on Weak LLM across datasets. Reported metrics are Weak LLM Accuracy (Weak Acc.) and Weak Correct Accepted (Weak Corr. Accpt.). Parameter settings are the same as in Table 3.

Benchmark	Pipeline	Weak Acc. (%) ↑	Weak Corr. Accpt. (%) ↑
GSM-Symb.	Jung et al.	15.04	12.34
	Our	**48.10**	**46.09**
GSM-Plus	Jung et al.	22.46	19.13
	Our	**38.66**	**35.73**
MetaMath(20K)	Jung et al.	34.95	28.54
	Our	**58.16**	**54.07**
NASA-Hist.	Jung et al.	66.22	55.37
	Our	**66.70**	**58.40**

strategies help the Weak LLM identify correct answers it would otherwise have deferred unnecessarily. This is a crucial factor in converting local improvements into overall pipeline gains.

Table 5: Token and API cost changes across datasets for Inter-Cascade compared with Jung's pipeline. More detailed analysis with input/output tokens is in Appendix I: Table 9.

Benchmark	Weak LLM Tokens	Strong LLM Tokens	Token Price
	Total	Total	
GSM-Symb.	+147.66%	-47.80%	-49.63%
GSM-Plus	+145.96%	-29.95%	-30.41%
Meta(20K)	+127.90%	-52.18%	-52.15%
NASA-Hist.	+132.58%	-15.47%	-15.75%

Table 6: Processing Latency and Strategy Repository Size across different datasets. Retrieval refers to the time spent on strategies matching and ranking. Generation refers to time spent on generating answer via API.

Benchmark	Tested Samples	Our			Jung et al.	Repository Size
		Total	Retrieval	Generation	Total	
GSM-Symb.	11250	2.19s	0.10s	2.09s	1.83s	15.4 MB
GSM-Plus	9504	1.72s	0.06s	1.66s	1.66s	12.9 MB
Meta(20K)	20000	1.60s	0.06s	1.54s	1.54s	19.6 MB
NASA-Hist.	6469	1.28s	0.07s	1.21s	1.30s	8.8 MB

Effect of Strategies on Accuracy and Confidence Calibration. As mentioned earlier, one notable observation from our experiments is that providing strategies enhances the Weak LLM's ability to assess its own accuracy. To further investigate this observation, we present Figure 2 for the GSM-Symbolic dataset. Analyses for the other three datasets, which exhibit similar patterns, are provided in Appendix G. Figure 2a depicts the accuracy of the Weak LLM as a function of the confidence threshold. For each threshold, only queries with confidence equal to or above the threshold are considered, and accuracy is calculated as the proportion of correct predictions. The figure further demonstrates that our pipeline consistently improves the accuracy of queries that pass the threshold. Figures 2b, 2c, and 2d illustrate the distribution of query confidence. The histogram offers insight into prediction coverage across different confidence thresholds and shows that our method outperforms the baselines in terms of coverage. Together, these figures indicate that our method not only helps the Weak LLM produce correct answers, but also enables it to better calibrate its confidence by being more confident when the answer is correct and less confident when it is incorrect.

(a) Accuracy vs. Confidence Threshold

(b) Base Weak LLM

(c) Inter-Cascade Random Strategies

(d) Inter-Cascade Retrieval Strategies

Figure 2: GSM-Symbolic dataset: (a) Accuracy as a function of the confidence threshold for the base Weak LLM, Inter-Cascade with random strategies, and Inter-Cascade with retrieval strategies, and (b) - (d) their corresponding confidence histograms. Our Inter-Cascade (Retrieval) consistently concentrates probability mass near high confidence ($0.9-1.0$), while the weak and random variants place more mass at low confidence, which explains the accuracy gains observed in (a).

Table 7: Pipeline Accuracy and Strong LLM Call Rate in the ablation study on strategy selection : Our "No strategy" (Our NS) vs. Our "Random" Rand vs. Our "Retrieval" (Ret). Parameter settings are the same as Table 3.

Benchmark	Pipeline	Pipeline Acc. (%) ↑	Strong Call (%) ↓	Cov. (%)
GSM Symb.	Our (NS)	67.55	65.15	83.14
	Our (Rand)	63.61	54.20	87.90
	Our (Ret)	**70.37**	**30.84**	**90.35**
GSM Plus	Our (NS)	58.12	54.81	93.83
	Our (Rand)	53.63	43.64	94.10
	Our (Ret)	**58.31**	**32.44**	**94.79**
MetaMath (20K)	Our (NS)	**74.48**	57.32	100.00
	Our (Rand)	67.85	45.99	100.00
	Our (Ret)	71.56	**23.68**	100.00
NASA-Hist.	Our (NS)	**74.64**	65.12	100.00
	Our (Rand)	71.32	25.09	100.00
	Our (Ret)	72.64	**22.54**	100.00

Token and API Cost Savings. Our pipeline not only improves accuracy but also reduces the number of Strong LLM calls, resulting in substantially lower token consumption on Strong LLM. Table 5 shows the percentage changes in token usage and corresponding API costs compared with Jung's pipeline. Table 6 shows the average processing time per query (including the call of Strong LLM) and the final size of strategies

Table 8: Weak LLM performance in the ablation study on strategy selection: Our "No strategy" (Our NS) vs. Our "Random" Rand vs. Our "Retrieval" (Ret). Parameter settings are the same as Table 3.

Benchmark	Pipeline	Weak Acc. (%) ↑	Weak Corr. Accpt. (%) ↑
GSM-Symb.	Our (NS)	10.23	17.08
	Our (Rand)	17.40	15.27
	Our (Ret)	**48.10**	**46.09**
GSM-Plus	Our (NS)	20.20	17.08
	Our (Rand)	25.51	22.38
	Our (Ret)	**38.66**	**35.73**
MetaMath(20K)	Our (NS)	33.40	28.38
	Our (Rand)	38.64	32.66
	Our (Ret)	**58.16**	**54.07**
NASA-Hist.	Our (NS)	28.21	22.88
	Our (Rand)	65.22	55.56
	Our (Ret)	**66.70**	**58.40**

repository across datasets. The results imply that the time difference is between -0.02s and +0.36s, which won't impact the user experience. The size of repository is at level of 10MB+ when the number of queries is at 10K+ level, which can be easily maintained in resource limited settings like mobile or edge device. More promisingly, accumulated queries and responses can serve as training data for periodic offline fine-tuning the Weak LLM (for example as part of a software update), enabling a self-improving pipeline that dynamically adapts to new data.

Cold Start of the Strategy Repository. To evaluate the effect of cold start in our strategy repository, we analyze the evolution of pipeline accuracy for both Jung's method and our standard Inter-Cascade on GSM-Symbolic (Figure 3). As shown in Figure 3a, at early stages, Inter-Cascade performs comparably to the baseline method of Jung (Jung et al., 2025). As the repository grows with accumulated strategies, its performance steadily improves, eventually surpassing the baseline and converging. Figure 3b shows the evolution of the weak LLM when equipped with strategies derived from the strong LLM, illustrating how its performance changes as the repository grows.

(a) Pipeline accuracy comparison

(b) Weak LLM evolution with strong LLM strategies

Figure 3: Performance trends on GSM-Symbolic. (a) compares pipeline accuracy between Jung et al.'s method and our standard Inter-Cascade. (b) shows the evolution of weak LLM performance when enhanced with strategies from strong LLMs.

Ablation Study on Strategy Selection In order to evaluate the impact of each part when we add strategies to the input of Weak LLM, we conduct ablation experiments for different settings: only adding

similar questions and answers (No strategy), adding randomly selected strategies (Random), and our standard Inter-Cascade pipeline (Retrieval). The results in Table 7 and Table 8, show that the performance of Random Strategy method is between our standard pipeline and Jung's method, while No Strategy is not an acceptable option. Although in benchmarks like NASA-History, the overall accuracy is 2.00% higher than our standard pipeline, the cost is significant: the Strong Call Rate increase by 42.58%, which means only adding similar question and answer to the input of Weak LLM would use 2.89x of the Strong LLM. Moreover, the Weak LLM's accuracy would be dramatically undermined by adding non-strategy information to the input of Weak LLM compared to the accuracy for single Weak LLM in Table 1. Only adding retrieved question and answers without instructive and generalized problem solving strategy to Weak LLM input is harmful: not only lower the accuracy of Weak LLM, but also call more Strong LLM, which is more expensive. Extensive Ablation studies on **effect of the size** of strategies and different **selection of LLM pairs** are attached in Appendix K.

Inter-Cascade Robustness under Automatic Strategies. All strategies and their corresponding answers are generated by the Strong LLM in a streaming manner, and any strategy whose confidence exceeds the threshold λ_s is automatically accepted. This differentiates *Inter-Cascade* from other LLM augmentation methods such as manually selected in-context learning, few-shot prompting, or static retrieval-augmented generation. Consequently, the strategy repository may contain incorrect strategies. Nonetheless, the results in Table 3 and Table 4 demonstrate the effectiveness of λ_s and the robustness of the *Inter-Cascade* pipeline.

5 Conclusion

We propose *Inter-Cascade*, an online interactive framework that enables Weak LLMs to learn from Strong LLMs' prior reasoning without fine-tuning. By transforming the strong model into a teacher, Inter-Cascade significantly improves both the weak model's accuracy and overall system performance while reducing computational costs and reliance on expensive models compared to standard static cascades.

As a general and scalable framework, Inter-Cascade opens several avenues for future research. Immediate improvements could focus on refining strategy generation, optimizing similarity retrieval algorithms, and mitigating context mismatch. Furthermore, the framework is naturally suited for distributed systems, allowing local models to tailor their capabilities by selectively querying Strong LLM. Finally, Inter-Cascade bridges the gap between online and offline learning. The dynamically generated strategy repository not only augments inference in real-time but can also serve as a high-quality dataset for periodic fine-tuning, permanently internalizing the strong model's capabilities. We hope this work inspires further exploration into interactive, teacher-student dynamics within multi-LLM systems.

References

Yu Bai, Yukai Miao, Li Chen, Dawei Wang, Dan Li, Yanyu Ren, Hongtao Xie, Ce Yang, and Xuhui Cai. Pistis-RAG: Enhancing retrieval-augmented generation with human feedback. *arXiv preprint arXiv:2407.00072*, 2024.

Peter Bauer. Multiple testing in clinical trials. *Statistics in medicine*, 10(6):871–890, 1991.

Peter Belcak, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. Small language models are the future of agentic AI. *arXiv preprint arXiv:2506.02153*, 2025.

Lawrence D. Brown, T. Tony Cai, and Anirban DasGupta. Interval estimation for a binomial proportion. *Statistical Science*, 16(2):101–133, 2001.

Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to use large language models while reducing cost and improving performance. *Transactions on Machine Learning Research*, 2024a. ISSN 2835-8856.

Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to Use Large Language Models While Reducing Cost and Improving Performance. *Transactions on Machine Learning Research*, July 2024b. ISSN 2835-8856.

Qinwen Chen, Wenbiao Tao, Zhiwei Zhu, Mingfan Xi, Liangzhong Guo, Yuan Wang, Wei Wang, and Yunshi Lan. ComRAG: Retrieval-augmented generation with dynamic vector stores for real-time community question answering in industry. *arXiv preprint arXiv:2506.21098*, 2025a.

Zhijun Chen, Jingzheng Li, Pengpeng Chen, Zhuoran Li, Kai Sun, Yuankai Luo, Qianren Mao, Ming Li, Likang Xiao, Dingqi Yang, et al. Harnessing multiple large language models: A survey on LLM ensemble. *arXiv preprint arXiv:2502.18036*, 2025b.

C Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on information theory*, 16(1): 41–46, 1970.

C. K. Chow. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, EC-6(4):247–254, December 1957. ISSN 0367-9950. doi: 10.1109/TEC.1957.5222035.

Yu-Neng Chuang, Helen Zhou, Prathusha Kameswara Sarma, Parikshit Gopalan, John Boccio, Sara Bolouki, and Xia Hu. Learning to Route LLMs with Confidence Tokens. In *Proceedings of the Forty-Second International Conference on Machine Learning*. PMLR, 2025.

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*, November 2021.

E. A. Cornish and R. A. Fisher. Moments and cumulants in the specification of distributions. *Revue de l'Institut International de Statistique*, 5(4):307–320, 1938.

Corinna Cortes, Giulia DeSalvo, and Mehryar Mohri. Learning with Rejection. In Ronald Ortner, Hans Ulrich Simon, and Sandra Zilles (eds.), *Algorithmic Learning Theory*, volume 9925, pp. 67–82. Springer International Publishing, Cham, 2016. ISBN 978-3-319-46378-0 978-3-319-46379-7. doi: 10.1007/978-3-319-46379-7_5.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. A survey on in-context learning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 1107–1128, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.64.

Matthijs Douze, Alexandr Guzhva, Chengqi Deng, Jeff Johnson, Gergely Szilvasy, Pierre-Emmanuel Mazaré, Maria Lomeli, Lucas Hosseini, and Hervé Jégou. The Faiss library. *arXiv preprint arXiv:2401.08281*, 2025.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*, 2024.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *arXiv preprint arXiv:2404.16130*, February 2025.

Andrew Estornell and Yang Liu. Multi-LLM Debate: Framework, Principals, and Interventions. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, November 2024.

Patrick Fleith. NASA-history-MCQ. Dataset, Hugging Face, 2025.

Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. From RAG to memory: Non-parametric continual learning for large language models. *arXiv preprint arXiv:2502.14802*, 2025.

Peter Hall. *The Bootstrap and Edgeworth Expansion*. Springer, 1992.

Siwei Han, Peng Xia, Ruiyi Zhang, Tong Sun, Yun Li, Hongtu Zhu, and Huaxiu Yao. MDocAgent: A multi-modal multi-agent framework for document understanding. *arXiv preprint arXiv:2503.13964*, 2025.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. *NeurIPS*, 2021.

Radu Herbei and Marten H. Wegkamp. Classification with Reject Option. *The Canadian Journal of Statistics / La Revue Canadienne de Statistique*, 34(4):709–721, 2006. ISSN 0319-5724.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531*, March 2015.

Geoffrey Irving, Paul Christiano, and Dario Amodei. AI safety via debate. *arXiv preprint arXiv:1805.00899*, October 2018.

Weisen Jiang, Han Shi, Longhui Yu, Zhengying Liu, Yu Zhang, Zhenguo Li, and James Kwok. Forward-backward reasoning in large language models for mathematical verification. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 6647–6661, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.397.

Wittawat Jitkrittum, Neha Gupta, Aditya K. Menon, Harikrishna Narasimhan, Ankit Rawat, and Sanjiv Kumar. When Does Confidence-Based Cascade Deferral Suffice? *Advances in Neural Information Processing Systems*, 36:9891–9906, December 2023.

Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2021. doi: 10.1109/TBDATA.2019.2921572.

Norman L. Johnson, Samuel Kotz, and Narayanaswamy Balakrishnan. *Continuous Univariate Distributions, Volume 2*, volume 2. John Wiley & Sons, 2nd edition, 1995.

Chaitanya K. Joshi, Fayao Liu, Xu Xun, Jie Lin, and Chuan-Sheng Foo. On Representation Knowledge Distillation for Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 35 (4):4656–4667, April 2024. ISSN 2162-237X, 2162-2388. doi: 10.1109/TNNLS.2022.3223018.

Jaehun Jung, Faeze Brahman, and Yejin Choi. Trust or escalate: LLM judges with provable guarantees for human agreement. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Representation Learning*, volume 2025, pp. 3101–3125, 2025.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Akbir Khan, John Hughes, Dan Valentine, Laura Ruis, Kshitij Sachan, Ansh Radhakrishnan, Edward Grefenstette, Samuel R. Bowman, Tim Rocktäschel, and Ethan Perez. Debating with more persuasive LLMs leads to more truthful answers. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Hyoje Lee, Yeachan Park, Hyun Seo, and Myungjoo Kang. Self-knowledge distillation via dropout. *Comput. Vis. Image Underst.*, 233(C), August 2023. ISSN 1077-3142. doi: 10.1016/j.cviu.2023.103720.

Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative decoding. In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in neural information processing systems*, 33:9459–9474, 2020.

Qintong Li, Leyang Cui, Xueliang Zhao, Lingpeng Kong, and Wei Bi. GSM-Plus: A comprehensive benchmark for evaluating the robustness of LLMs as mathematical problem solvers. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2961–2984, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.163.

Yangning Li, Weizhi Zhang, Yuyao Yang, Wei-Chieh Huang, Yaozu Wu, Junyu Luo, Yuanchen Bei, Henry Peng Zou, Xiao Luo, Yusheng Zhao, et al. Towards agentic RAG with deep reasoning: A survey of RAG-reasoning systems in LLMs. *arXiv preprint arXiv:2507.09477*, 2025.

Chengyuan Liu, Fubang Zhao, Kun Kuang, Yangyang Kang, Zhuoren Jiang, Changlong Sun, and Fei Wu. Evolving knowledge distillation with large language models and active learning. In Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 6717–6731, Torino, Italia, May 2024. ELRA and ICCL.

Pei Liu, Xin Liu, Ruoyu Yao, Junming Liu, Siyuan Meng, Ding Wang, and Jun Ma. HM-RAG: Hierarchical multi-agent multimodal retrieval augmented generation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 2781–2790, 2025.

Chang Han Low, Ziyue Wang, Tianyi Zhang, Zhitao Zeng, Zhu Zhuo, Evangelos B Mazomenos, and Yueming Jin. Surgraw: Multi-agent workflow with chain-of-thought reasoning for surgical intelligence. *arXiv preprint arXiv:2503.10265*, 2025.

David Madras, Toniann Pitassi, and Richard Zemel. Predict responsibly: improving fairness and accuracy by learning to defer. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pp. 6150–6160, Red Hook, NY, USA, 2018. Curran Associates Inc.

Anqi Mao, Mehryar Mohri, and Yutao Zhong. Principled Approaches for Learning to Defer with Multiple Experts. In Reneta P. Barneva, Valentin E. Brimkov, Claudio Gentile, and Aldo Pacchiano (eds.), *Artificial Intelligence and Image Analysis*, volume 14494, pp. 107–135. Springer Nature Switzerland, Cham, 2024a. ISBN 978-3-031-63734-6 978-3-031-63735-3. doi: 10.1007/978-3-031-63735-3_7.

Anqi Mao, Mehryar Mohri, and Yutao Zhong. Theoretically Grounded Loss Functions and Algorithms for Score-Based Multi-Class Abstention. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pp. 4753–4761. PMLR, April 2024b.

Katerina Margatina, Timo Schick, Nikolaos Aletras, and Jane Dwivedi-Yu. Active learning principles for in-context learning with large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 5011–5034, December 2023.

Seyed Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. GSM-Symbolic: Understanding the limitations of mathematical reasoning in large language models. In *The Thirteenth International Conference on Learning Representations*, 2025.

Fengran Mo, Chuan Meng, Mohammad Aliannejadi, and Jian-Yun Nie. Conversational search: From fundamentals to frontiers in the LLM era. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 4094–4097, 2025.

Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7076–7087. PMLR, 13–18 Jul 2020.

Harikrishna Narasimhan, Wittawat Jitkrittum, Ankit Singh Rawat, Seungyeon Kim, Neha Gupta, Aditya Krishna Menon, and Sanjiv Kumar. Faster cascades via speculative decoding. In *The Thirteenth International Conference on Learning Representations*, 2025.

Cuong C Nguyen, Thanh-Toan Do, and Gustavo Carneiro. Probabilistic learning to defer: Handling missing expert annotations and controlling workload distribution. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Thang Nguyen, Peter Chin, and Yu-Wing Tai. MA-RAG: Multi-agent retrieval-augmented generation via collaborative Chain-of-Thought reasoning. *arXiv preprint arXiv:2505.20096*, 2025b.

Lunyiu Nie, Zhimin Ding, Erdong Hu, Christopher Jermaine, and Swarat Chaudhuri. Online cascade learning for efficient inference over streams. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, and Ion Stoica. RouteLLM: Learning to route LLMs from preference data. In *The Thirteenth International Conference on Learning Representations*, 2025.

Aitor Ormazabal, Mikel Artetxe, and Eneko Agirre. CombLM: Adapting black-box language models through small fine-tuned models. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 2961–2974, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.180.

Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. MedMCQA: A large-scale multi-subject multi-choice dataset for medical domain question answering. In Gerardo Flores, George H Chen, Tom Pollard, Joyce C Ho, and Tristan Naumann (eds.), *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pp. 248–260. PMLR, 07–08 Apr 2022.

Archit Parnami and Minwoo Lee. Learning from few examples: A summary of approaches to few-shot learning. *arXiv preprint arXiv:2203.04291*, 2022.

Cuong Pham, Van-Anh Nguyen, Trung Le, Dinh Phung, Gustavo Carneiro, and Thanh-Toan Do. Frequency attention for knowledge distillation. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 2266–2275, 2024.

Sahana Rayan and Ambuj Tewari. Learning to Partially Defer for Sequences. *arXiv preprint arXiv:2502.01459*, February 2025.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3982–3992, November 2019.

Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for Thin Deep Nets. *arXiv preprint arXiv:1412.6550*, March 2015.

Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2655–2671, July 2022.

Zejiang Shen, Hunter Lang, Bailin Wang, Yoon Kim, and David Sontag. Learning to decode collaboratively with multiple language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12974–12990, August 2024.

Yunxiao Shi, Xing Zi, Zijing Shi, Haimin Zhang, Qiang Wu, and Min Xu. ERAGent: Enhancing retrieval-augmented language models with improved accuracy, efficiency, and personalization. *arXiv preprint arXiv:2405.06683*, 2024.

Robik Shrestha, Yang Zou, Qiuyu Chen, Zhiheng Li, Yusheng Xie, and Siqi Deng. FairRAG: Fair human generation via fair retrieval augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11996–12005, 2024.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ambrose Slone, Ameet Rahane, Anantharaman S. Iyer, Anders Johan Andreassen, Andrea Madotto, Andrea Santilli, Andreas Stuhlmüller, Andrew M. Dai, Andrew La, Andrew Kyle Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash

Gholamidavoodi, Arfa Tabassum, Arul Menezes, Arun Kirubarajan, Asher Mullokandov, Ashish Sabharwal, Austin Herrick, Avia Efrat, Aykut Erdem, Ayla Karakaş, B. Ryan Roberts, Bao Sheng Loe, Barret Zoph, Bartłomiej Bojanowski, Batuhan Özyurt, Behnam Hedayatnia, Behnam Neyshabur, Benjamin Inden, Benno Stein, Berk Ekmekci, Bill Yuchen Lin, Blake Howald, Bryan Orinion, Cameron Diao, Cameron Dour, Catherine Stinson, Cedrick Argueta, Cesar Ferri, Chandan Singh, Charles Rathkopf, Chenlin Meng, Chitta Baral, Chiyu Wu, Chris Callison-Burch, Christopher Waites, Christian Voigt, Christopher D Manning, Christopher Potts, Cindy Ramirez, Clara E. Rivera, Clemencia Siro, Colin Raffel, Courtney Ashcraft, Cristina Garbacea, Damien Sileo, Dan Garrette, Dan Hendrycks, Dan Kilman, Dan Roth, C. Daniel Freeman, Daniel Khashabi, Daniel Levy, Daniel Moseguí González, Danielle Perszyk, Danny Hernandez, Danqi Chen, Daphne Ippolito, Dar Gilboa, David Dohan, David Drakard, David Jurgens, Debajyoti Datta, Deep Ganguli, Denis Emelin, Denis Kleyko, Deniz Yuret, Derek Chen, Derek Tam, Dieuwke Hupkes, Diganta Misra, Dilyar Buzan, Dimitri Coelho Mollo, Diyi Yang, Dong-Ho Lee, Dylan Schrader, Ekaterina Shutova, Ekin Dogus Cubuk, Elad Segal, Eleanor Hagerman, Elizabeth Barnes, Elizabeth Donoway, Ellie Pavlick, Emanuele Rodolà, Emma Lam, Eric Chu, Eric Tang, Erkut Erdem, Ernie Chang, Ethan A Chi, Ethan Dyer, Ethan Jerzak, Ethan Kim, Eunice Engefu Manyasi, Evgenii Zheltonozhskii, Fanyue Xia, Fatemeh Siar, Fernando Martínez-Plumed, Francesca Happé, Francois Chollet, Frieda Rong, Gaurav Mishra, Genta Indra Winata, Gerard de Melo, Germàn Kruszewski, Giambattista Parascandolo, Giorgio Mariani, Gloria Xinyue Wang, Gonzalo Jaimovitch-Lopez, Gregor Betz, Guy Gur-Ari, Hana Galijasevic, Hannah Kim, Hannah Rashkin, Hannaneh Hajishirzi, Harsh Mehta, Hayden Bogar, Henry Francis Anthony Shevlin, Hinrich Schuetze, Hiromu Yakura, Hongming Zhang, Hugh Mee Wong, Ian Ng, Isaac Noble, Jaap Jumelet, Jack Geissinger, Jackson Kernion, Jacob Hilton, Jaehoon Lee, Jaime Fernández Fisac, James B Simon, James Koppel, James Zheng, James Zou, Jan Kocon, Jana Thompson, Janelle Wingfield, Jared Kaplan, Jarema Radom, Jascha Sohl-Dickstein, Jason Phang, Jason Wei, Jason Yosinski, Jekaterina Novikova, Jelle Bosscher, Jennifer Marsh, Jeremy Kim, Jeroen Taal, Jesse Engel, Jesujoba Alabi, Jiacheng Xu, Jiaming Song, Jillian Tang, Joan Waweru, John Burden, John Miller, John U. Balis, Jonathan Batchelder, Jonathan Berant, Jörg Frohberg, Jos Rozen, Jose Hernandez-Orallo, Joseph Boudeman, Joseph Guerr, Joseph Jones, Joshua B. Tenenbaum, Joshua S. Rule, Joyce Chua, Kamil Kanclerz, Karen Livescu, Karl Krauth, Karthik Gopalakrishnan, Katerina Ignatyeva, Katja Markert, Kaustubh Dhole, Kevin Gimpel, Kevin Omondi, Kory Wallace Mathewson, Kristen Chiafullo, Ksenia Shkaruta, Kumar Shridhar, Kyle McDonell, Kyle Richardson, Laria Reynolds, Leo Gao, Li Zhang, Liam Dugan, Lianhui Qin, Lidia Contreras-Ochando, Louis-Philippe Morency, Luca Moschella, Lucas Lam, Lucy Noble, Ludwig Schmidt, Luheng He, Luis Oliveros-Colón, Luke Metz, Lütfi Kerem Senel, Maarten Bosma, Maarten Sap, Maartje Ter Hoeve, Maheen Farooqi, Manaal Faruqui, Mantas Mazeika, Marco Baturan, Marco Marelli, Marco Maru, Maria Jose Ramirez-Quintana, Marie Tolkiehn, Mario Giulianelli, Martha Lewis, Martin Potthast, Matthew L Leavitt, Matthias Hagen, Mátyás Schubert, Medina Orduna Baitemirova, Melody Arnaud, Melvin McElrath, Michael Andrew Yee, Michael Cohen, Michael Gu, Michael Ivanitskiy, Michael Starritt, Michael Strube, Michał Swędrowski, Michele Bevilacqua, Michihiro Yasunaga, Mihir Kale, Mike Cain, Mimee Xu, Mirac Suzgun, Mitch Walker, Mo Tiwari, Mohit Bansal, Moin Aminnaseri, Mor Geva, Mozhdeh Gheini, Mukund Varma T, Nanyun Peng, Nathan Andrew Chi, Nayeon Lee, Neta Gur-Ari Krakover, Nicholas Cameron, Nicholas Roberts, Nick Doiron, Nicole Martinez, Nikita Nangia, Niklas Deckers, Niklas Muennighoff, Nitish Shirish Keskar, Niveditha S. Iyer, Noah Constant, Noah Fiedel, Nuan Wen, Oliver Zhang, Omar Agha, Omar Elbaghdadi, Omer Levy, Owain Evans, Pablo Antonio Moreno Casares, Parth Doshi, Pascale Fung, Paul Pu Liang, Paul Vicol, Pegah Alipoormolabashi, Peiyuan Liao, Percy Liang, Peter W Chang, Peter Eckersley, Phu Mon Htut, Pinyu Hwang, Piotr Miłkowski, Piyush Patil, Pouya Pezeshkpour, Priti Oli, Qiaozhu Mei, Qing Lyu, Qinlang Chen, Rabin Banjade, Rachel Etta Rudolph, Raefer Gabriel, Rahel Habacker, Ramon Risco, Raphaël Millière, Rhythm Garg, Richard Barnes, Rif A. Saurous, Riku Arakawa, Robbe Raymaekers, Robert Frank, Rohan Sikand, Roman Novak, Roman Sitelew, Ronan Le Bras, Rosanne Liu, Rowan Jacobs, Rui Zhang, Russ Salakhutdinov, Ryan Andrew Chi, Seungjae Ryan Lee, Ryan Stovall, Ryan Teehan, Rylan Yang, Sahib Singh, Saif M. Mohammad, Sajant Anand, Sam Dillavou, Sam Shleifer, Sam Wiseman, Samuel Gruetter, Samuel R. Bowman, Samuel Stern Schoenholz, Sanghyun Han, Sanjeev Kwatra, Sarah A. Rous, Sarik Ghazarian, Sayan Ghosh, Sean Casey, Sebastian Bischoff, Sebastian Gehrmann, Sebastian Schuster, Sepideh Sadeghi, Shadi Hamdan, Sharon Zhou, Shashank Srivastava, Sherry Shi, Shikhar Singh, Shima Asaadi, Shixiang Shane Gu, Shubh Pachchigar, Shubham Toshniwal, Shyam Upadhyay, Shyamolima Shammie Debnath, Siamak Shakeri, Simon Thormeyer,

Simone Melzi, Siva Reddy, Sneha Priscilla Makini, Soo-Hwan Lee, Spencer Torene, Sriharsha Hatwar, Stanislas Dehaene, Stefan Divic, Stefano Ermon, Stella Biderman, Stephanie Lin, Stephen Prasad, Steven Piantadosi, Stuart Shieber, Summer Misherghi, Svetlana Kiritchenko, Swaroop Mishra, Tal Linzen, Tal Schuster, Tao Li, Tao Yu, Tariq Ali, Tatsunori Hashimoto, Te-Lin Wu, Théo Desbordes, Theodore Rothschild, Thomas Phan, Tianle Wang, Tiberius Nkinyili, Timo Schick, Timofei Kornev, Titus Tunduny, Tobias Gerstenberg, Trenton Chang, Trishala Neeraj, Tushar Khot, Tyler Shultz, Uri Shaham, Vedant Misra, Vera Demberg, Victoria Nyamai, Vikas Raunak, Vinay Venkatesh Ramasesh, vinay uday prabhu, Vishakh Padmakumar, Vivek Srikumar, William Fedus, William Saunders, William Zhang, Wout Vossen, Xiang Ren, Xiaoyu Tong, Xinran Zhao, Xinyi Wu, Xudong Shen, Yadollah Yaghoobzadeh, Yair Lakretz, Yangqiu Song, Yasaman Bahri, Yejin Choi, Yichi Yang, Sophie Hao, Yifu Chen, Yonatan Belinkov, Yu Hou, Yufang Hou, Yuntao Bai, Zachary Seid, Zhuoye Zhao, Zijian Wang, Zijie J. Wang, Zirui Wang, and Ziyi Wu. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. Featured Certification.

Joshua Strong, Qianhui Men, and J. Alison Noble. Trustworthy and practical AI for healthcare: a guided deferral system with large language models. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence and Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence and Fifteenth Symposium on Educational Advances in Artificial Intelligence*, AAAI'25/IAAI'25/EAAI'25. AAAI Press, 2025a. ISBN 978-1-57735-897-8. doi: 10.1609/aaai.v39i27.35063.

Joshua Strong, Pramit Saha, Yasin Ibrahim, Cheng Ouyang, and Alison Noble. Expert-agnostic learning to defer. *arXiv preprint arXiv:2502.10533*, 2025b.

Mirac Suzgun, Nathan Scales, Nathanael Scharli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging BIG-Bench tasks and whether Chain-of-Thought can solve them. In *Annual Meeting of the Association for Computational Linguistics*, 2022.

Dharmesh Tailor, Aditya Patra, Rajeev Verma, Putra Manggala, and Eric Nalisnick. Learning to Defer to a Population: A Meta-Learning Approach. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, pp. 3475–3483. PMLR, April 2024.

Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 2464–2469. IEEE, 2016.

Aman Singh Thakur, Kartik Choudhary, Venkat Srinik Ramayapally, Sankaran Vaidyanathan, and Dieuwke Hupkes. Judging the judges: Evaluating alignment and vulnerabilities in LLMs-as-Judges. In *Proceedings of the Fourth Workshop on Generation, Evaluation and Metrics (GEM²)*, pp. 404–430, July 2025.

Rajeev Verma and Eric Nalisnick. Calibrated Learning to Defer with One-vs-All Classifiers. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 22184–22202. PMLR, June 2022.

Rajeev Verma, Daniel Barrejon, and Eric Nalisnick. Learning to Defer to Multiple Experts: Consistent Surrogate Losses, Confidence Calibration, and Conformal Ensembles. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, pp. 11415–11434. PMLR, April 2023.

Futing Wang, Jianhao Yan, Yue Zhang, and Tao Lin. ELICIT: LLM augmentation via external in-context capability. In *The Thirteenth International Conference on Learning Representations*, 2025a.

Haorui Wang, Rongzhi Zhang, Yinghao Li, Lingkai Kong, Yuchen Zhuang, Xiusi Chen, and Chao Zhang. TPD: Enhancing student language model reasoning via principle discovery and guidance. In *First Conference on Language Modeling*, 2024a.

Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. Large language models are not fair evaluators. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pp. 9440–9450, August 2024b.

Ruotong Wang, Xinyi Zhou, Lin Qiu, Joseph Chee Chang, Jonathan Bragg, and Amy X Zhang. Social-RAG: Retrieving from group interactions to socially ground AI generation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–25, 2025b.

Zheng Wang, Shu Teo, Jieer Ouyang, Yongjun Xu, and Wei Shi. M-RAG: Reinforcing large language model performance through retrieval-augmented generation with multiple partitions. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1966–1978, 2024c.

Yu Wu and Anand Sarwate. Learning to help: Training models to assist legacy devices. *arXiv preprint arXiv:2409.16253*, 2024.

Yu Wu, Yansong Li, Zeyu Dong, Nitya Sathyavageeswaran, and Anand D. Sarwate. Learning to help in multi-class settings. In *The Thirteenth International Conference on Learning Representations*, 2025.

Yu Xia, Fang Kong, Tong Yu, Liya Guo, Ryan A. Rossi, Sungchul Kim, and Shuai Li. Which LLM to play? convergence-aware online model selection with time-increasing bandits. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pp. 4059–4070, 2024. doi: 10.1145/3589334.3645420.

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. In *The Twelfth International Conference on Learning Representations*, 2024.

Zhichao Xu, Minheng Wang, Yawei Wang, Wenqian Ye, Yuntao Du, Yunpu Ma, and Yijun Tian. RECON: Reasoning with condensation for efficient retrieval-augmented generation. *arXiv preprint arXiv:2510.10448*, 2025.

Diji Yang, Jinmeng Rao, Kezhen Chen, Xiaoyuan Guo, Yawen Zhang, Jie Yang, and Yi Zhang. Im-RAG: Multi-round retrieval-augmented generation through learning inner monologues. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 730–740, 2024.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng YU, Zhengying Liu, Yu Zhang, James Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. MetaMath: Bootstrap your own mathematical questions for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.

Michael J. Zellinger and Matt Thomson. Rational tuning of LLM cascades via probabilistic modeling. *Transactions on Machine Learning Research*, 2025.

Michael J. Zellinger, Rex Liu, and Matt Thomson. Cost-Saving LLM Cascades with Early Abstention. *arXiv preprint arXiv:2502.09054*, February 2025.

Feiyuan Zhang, Dezhi Zhu, James Ming, Yilun Jin, Di Chai, Liu Yang, Han Tian, Zhaoxin Fan, and Kai Chen. DH-RAG: A dynamic historical context-powered retrieval-augmented generation method for multi-turn dialogue. *arXiv preprint arXiv:2502.13847*, 2025.

Lucia Zheng, Neel Guha, Javokhir Arifov, Sarah Zhang, Michal Skreta, Christopher D. Manning, Peter Henderson, and Daniel E. Ho. A reasoning-focused legal retrieval benchmark. In *Proceedings of the 2025 Symposium on Computer Science and Law*, CSLAW '25, pp. 169–193, New York, NY, USA, 2025. Association for Computing Machinery. ISBN 9798400714214. doi: 10.1145/3709025.3712219.

Xiaofeng Zhou, Heyan Huang, and Lizi Liao. Debate, reflect, and distill: Multi-agent feedback with tree-structured preference optimization for efficient language model enhancement. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 9122–9137, July 2025.

A Extended Related Work

LLM Cascade There are many LLM paradigms that contain collaboration between multiples LLMs in a system (Chen et al., 2025b): a)Ensemble before inference, where router choose one LLM from candidates for inference; b)Ensemble during inference, where LLMs work in parallel; c) Ensemble after inference, where LLMs work in sequence and LLM Cascade belongs to this filed. LLM Cascade is firstly proposed by Chen et al. (2024a) to balance the LLM performance and cost by allocating queries to a weak model or a strong model according to the confidence estimate of the queried question. Shen et al. (2024) propose a latent variable model to let the weak model learn the deferral function at the token-level. Rayan & Tewari (2025) also extend the Learning to Defer (Madras et al., 2018) setting to LLM by training a post-hoc deferral function for each token of the sequence. Ong et al. (2025) train a separate router such that deferral decision can be made before sending the query to weak LLM, saving more tokens. Zellinger et al. (2025) provide extra option to early discard the unsolvable queries in weak model. Xia et al. (2024); Nie et al. (2024) formulate LLM Cascade as online problem to dynamically adjust its deferral policy over time. Zellinger & Thomson (2025) propose a rational tuning pipeline for LLM Cascade via probabilistic modeling. Since the deferral result relies on the confidence score of weak model, there are are literatures focusing on boosting the the measure of confidence of weak model's output (Jitkrittum et al., 2023; Chuang et al., 2025). Together with experimental verification, Jung et al. (2025) conduct fixed sequence testing to provably guarantee the lower bound of accuracy. Therefore, we choose Cascaded Selective Evaluation by Jung et al. (2025) as the baseline of our work. Beside deferring to strong model, Beyond standard LLM Cascade, Strong et al. (2025a) propose a deferral system that weak model also sends its generated intelligent guidance to strong model once deferred, boosting the performance of next level model. However, current LLM Cascades cannot adapt to the query streaming once trained and deployed. And the weak model cannot learn from the previous deferrals and corresponding strategies generated by the strong model, causing the waste of computation, tokens, money and sometimes communication.

Learning With Reject Option The general framework that allows a machine learning model to abstain from making decision was originally propose by Chow (1957; 1970) in the 1950s. After decades, the Learning with reject option was continuously explored in different periods by Herbei & Wegkamp (2006) and Cortes et al. (2016). The more recent works extend the framework to a multi models system where the local model can learn to defer its task to one expert (human or existing model) (Madras et al., 2018; Mozannar & Sontag, 2020; Verma & Nalisnick, 2022; Mao et al., 2024b), multiple experts (Verma et al., 2023; Mao et al., 2024a) or unknown experts (Nguyen et al., 2025a; Strong et al., 2025b; Tailor et al., 2024). There are literature that also explore the case when expert can learn to adaptively help the local model (Wu & Sarwate, 2024; Wu et al., 2025). Adding reject option at the network layer level is another branch of works called early exiting (Teerapittayanon et al., 2016). However, most of the learning with reject option works focus on classical prediction tasks, few of them address the NLP tasks that rely on generative-based model while this work focus on the collaboration between LLMs.

Knowledge Distillation Knowledge distillation (KD) is a machine learning technique for training smaller "student" models by transferring "knowledge" from larger, more powerful "teacher" models. Classical knowledge distillation use soft-labels (Hinton et al., 2015) to let the student model learn the distribution of teacher model. The concept of KD is expanded to more levels: besides mimicking the output of teacher model, the student model can also learn from intermediate features (Romero et al., 2015; Pham et al., 2024), relationships (Joshi et al., 2024), actively chosen sample (Liu et al., 2024), principle discovery (Wang et al., 2024a) and itself (Lee et al., 2023). Our Inter-Cascade also helps the knowledge transfer from the Strong LLM to Weak LLM. However, current knowledge distillation relies on the training or finetuning of the student model and can not continue learning process during inference phase while our method doesn't require the updating of the LLM parameters and continually improves during the inference phase via dynamically matching stored Strong LLM's strategy.

Retrieval-Augmented Generation(RAG) RAG (Lewis et al., 2020) is an approach that combines pre-trained parametric and non-parametric memory for language generation. Given the focus of our work, we group RAG-style approaches into three categories: static RAG, history-aware RAG, and agentic RAG.

Static RAG. Classical RAG assumes a fixed, pre-constructed external corpus and focuses on how to retrieve, re-rank, and fuse evidence to support generation. Works in this line focus on design dense retrieval and re-ranking pipelines over a static collection (Lewis et al., 2020; Edge et al., 2025; Wang et al., 2025a; Rubin et al., 2022; Margatina et al., 2023).In all these methods, the source of knowledge is an offline, human-curated dataset, and the system's adaptivity lies purely in how it accesses this corpus, not what the corpus contains. By contrast, Inter-Cascade does not assume any pre-existing database: the "corpus" is constructed online as the strong LLM generates strategies and reasoning traces that are stored for future reuse by the weak LLM. Thus, our system is closer to an online, LLM-driven knowledge construction mechanism than to classical static RAG.

History-Aware RAG. A second line of work augments RAG with dialogue history and user feedback, dynamically updating a memory store based on past interactions. Conversational RAG frameworks like DH-RAG (Zhang et al., 2025), CHIQ (Mo et al., 2025) maintain short-term and long-term memories of successful dialogue turns, using them to improve future retrieval and personalization. Other methods such as ComRAG (Chen et al., 2025a), ERAGent (Shi et al., 2024), Pistis-RAG (Bai et al., 2024), and Social-RAG (Wang et al., 2025b) update user profiles or QA memories when users provide explicit positive feedback or when high-quality answers are validated by the social community. Despite their dynamism, these systems either take history information for self usage or treat the human user (or user community) as the source of new content. The resulting models are primarily personalized assistants. In Inter-Cascade, the update loop is fundamentally different: the weak LLM decides when to update, and the strong LLM decides what to write, without any human in the loop. The stored content is not user utterances or QA pairs, but LLM-generated strategies and reasoning structures distilled from a stronger model. Rather than personalizing to a single user, Inter-Cascade uses interaction between two models to build a reusable strategic knowledge base for many users and tasks.

Agentic RAG A third, increasingly prominent direction combines RAG with multi-agent or agentic architectures (Li et al., 2025). In these systems, different agents are assigned distinct roles, e.g., planner, retriever, answer generator, or verifier. Those agents collaborate via tool calls and message passing. For centralized systems like MA-RAG (Nguyen et al., 2025b), HM-RAG (Liu et al., 2025), and SurgRaw (Low et al., 2025), the focus is on managing the workflow, such as deciding when to use the retriever to access the existing database. Decentralized methods like M-RAG (Wang et al., 2024c) and MDocAgent (Han et al., 2025) consider retrieval from partitioned databases. There are also works like RECOND (end-to-end generation) (Xu et al., 2025), Hippo (knowledge-graph) (Gutiérrez et al., 2025), IM-RAG (multi step refinement) (Yang et al., 2024) and FAIR-RAG (fair retrieval) (Shrestha et al., 2024) propose algorithms to refine answers from RAG database. However, in all such designs, the RAG component itself remains an external, fixed resource: agents coordinate how to use RAG, but no agent is responsible for constructing a new corpus of knowledge for others. Inter-Cascade differs from these agentic RAG systems in two key aspects. First, there are only two "agents": a weak LLM and a strong LLM, but their interaction is explicitly teacher–student and online knowledge distillation, rather than mere division of labor. Second, the strong LLM actively produces the knowledge store that the weak LLM later retrieves, making the RAG-like database a product of model interaction rather than a static tool.

Across all three categories, existing RAG approaches either (i) operate over a fixed, human-curated external corpus, (ii) update a memory store using human dialogue and feedback, or (iii) update a memory using self history for personalization without knowledge transfer. To our knowledge, Inter-Cascade is the first framework where a weak LLM and a strong LLM jointly and autonomously build a RAG-like corpus under the framework of LLM Cascade, with the weak model deciding when to consult and update it, and the strong model providing the organized knowledge. This yields a new form of online, interaction-driven distillation, particularly suitable for small models without access to large external knowledge bases or the Internet.

Other related topics There are also a weak model and strong model in *s.peculative decoding* (Leviathan et al., 2023; Narasimhan et al., 2025). In speculative decoding, the weak model works as a answer draft while the strong model work as a verifier to speed up the generation compared to only using strong model. However, in Inter-Cascade, Strong LLM is called only when the Weak LLM is unable to handle current query. *CombLM* (Ormazabal et al., 2023) and *LLM Debate* (Irving et al., 2018; Du et al., 2024; Estornell & Liu, 2024; Khan et al., 2024; Zhou et al., 2025) are other branches of works that also involve interaction between LLMs. CombLM integrates the logit distribution of two LLMs while LLM Debate requires different LLMs to

argue and refine their initial answers and eventually reach consensus through multiple rounds of interaction. The key difference between Inter-Cascade and them is that Inter-Cascade let the Strong LLM and Weak LLM work in a sequential order can conduct early stop to save tokens.

Algorithm 3 Inter-Cascade Inference Pipeline

Input: Test set $\mathcal{T} = \{q_1, \ldots, q_I\} \subseteq \mathcal{Q}$; LLM M_n with deferral function d_n, generation function g_n, strategy repository Repo_n and strategy generator h_n.

    Deferral convention: $0 = \text{handle locally}$, $1 = \text{defer/forward}$.

1:   $\text{Repo} = \emptyset$
2: **for** $n \leftarrow 1$ **to** N **do**
3:     **for** $i \leftarrow 1$ **to** I **do**
4:        **if** $n < N$ **then**
5:           **(Strategy matching)**
6:           $[s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}] \leftarrow f_n(q_i, \text{Repo}_n)$ ▷ Find most relevant top-k strategies to q_i
7:           $q_i' \leftarrow [q_i, s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}]$ ▷ concatenate query and strategies
8:        **else**
9:           $q_i' = q_i$ ▷ Last LLM doesn't maintain Repo

10:        **(Deferral Decision)**
11:        **if** $d_n(q_i') = 0$ **then**
12:           generate answer $a_i \leftarrow g_1(q_i')$ ▷ Answer locally at Weaker LLM
13:           $s_{\text{new}} \leftarrow h(q_i)$
14:           $\text{Repo}_{<n} \leftarrow \text{Repo}_{<n} \cup \{s_{\text{new}}\}$ ▷ Add strategy to all the weaker LLMs
15:        **else**
16:           **if** $n < N$ **then**
17:              Pass   ▷ Defer to next level
18:           **else**
19:              Discard current query q_i ▷ None of LLMs are confident to answer the query

---

B   Order of LLMs

To distinguish two LLMs into strong model M_s and weak model M_w, we make following definitions. For a task distribution $\mathcal{D}$, we denote the performance of a model M by $\text{Perf}(M)$, which can be instantiated by measures such as the expected accuracy or negative loss on $\mathcal{D}$. Similarly, we let $\text{Cost}(M)$ represent the expected cost of using M on $\mathcal{D}$, such as the price, latency, or required computation resource. Note that Cost also depends on the task distribution $\mathcal{D}$, for simplicity, we only use the notation $\text{Cost}(M)$. We say that M_w is weaker than M_s if $\text{Perf}(M_w) \leq \text{Perf}(M_s)$, and that it is cheaper if $\text{Cost}(M_w) \leq \text{Cost}(M_s)$. To simplify notation, we introduce the shorthand relation

$$M_w \preccurlyeq_{\text{wbc}} M_s \tag{5}$$

if and only if

$$\text{Perf}(M_w) \leq \text{Perf}(M_s) \quad \text{and} \quad \text{Cost}(M_w) \leq \text{Cost}(M_s), \tag{6}$$

where the term "wbc" represents "weaker but cheaper". Consider a multi-LLM inference/generation system, which contains N LLM models, $\mathcal{M} = \{M_1, M_2, ..., M_N\}$, with different capacities and use costs to a query. WLOG, we assume that $M_1 \preccurlyeq_{\text{wbc}} M_2 \preccurlyeq_{\text{wbc}} ... \preccurlyeq_{\text{wbc}} M_N$.

C   Algorithm for General Inter-Cascade

Since Inter-Cascade is scalable to any number of layers for LLM, the general Inter-Cascade pipeline for N-LLM cascade system is shown in Algo. 3.

D Proof: Clopper–Pearson Upper Bound as a Beta Quantile

In the lemma below, we apply the Clopper–Pearson upper bound to rewrite $\widehat{R}^+(\lambda)$, yielding a clearer form that facilitates computation. This helps the proof of Theorem 3.2 and Theorem F.1.

Lemma D.1 (Clopper–Pearson upper bound as a Beta quantile). *Let $n(\lambda) \in \mathbb{N}$ be the number of evaluated items at threshold λ, let $R(\lambda) \in [0,1]$ denote the unknown risk, and suppose*

$$X \sim \mathrm{Bin}\big(n(\lambda),\, R(\lambda)\big), \tag{7}$$

and $x \in \{0, 1, \ldots, n(\lambda)\}$ is the number of errors observed (a realization of X). For a fixed $\delta \in (0,1)$, define the one-sided $(1 - \delta)$ upper confidence limit by

$$\widehat{R}^+(\lambda) := \sup \Big\{ p \in [0,1] :\ \Pr\big(\mathrm{Bin}(n(\lambda), p) \leq x\big) \geq \delta \Big\}. \tag{8}$$

Then

$$\widehat{R}^+(\lambda) = \mathrm{Beta}^{-1}\big(1 - \delta;\ x + 1,\ n(\lambda) - x\big), \tag{9}$$

where $\mathrm{Beta}^{-1}(q; a, b)$ denotes the q-quantile (inverse CDF in its first argument, the probability level) of the $\mathrm{Beta}(a, b)$ distribution. The case $x = 0$ gives the closed form $\widehat{R}^+(\lambda) = 1 - \delta^{1/n(\lambda)}$, while $x = n(\lambda)$ uses the boundary convention $\widehat{R}^+(\lambda) = 1$ (the parameter $\mathrm{Beta}(n(\lambda) + 1, 0)$ being degenerate).

Proof. For fixed $x < n(\lambda)$ the map $p \mapsto F(p) := \Pr\big(\mathrm{Bin}(n(\lambda), p) \leq x\big)$ is strictly decreasing in p, so the set in the definition of $\widehat{R}^+(\lambda)$ is an interval $[0, p^\star]$ and the supremum $p^\star$ uniquely solves

$$F(p^\star) = \Pr\big(\mathrm{Bin}(n(\lambda), p^\star) \leq x\big) = \delta. \tag{10}$$

Using the standard identity linking the binomial tail to the regularized incomplete beta function, for integers $0 \leq x \leq n(\lambda) - 1$,

$$\Pr\big(\mathrm{Bin}(n(\lambda), p) \leq x\big) = \sum_{k=0}^{x} \binom{n(\lambda)}{k} p^k (1-p)^{n(\lambda)-k} = 1 - I_p\big(x+1,\ n(\lambda) - x\big), \tag{11}$$

where $I_p(a, b)$ is the CDF of $\mathrm{Beta}(a, b)$ at p (the regularized incomplete beta function), so that $\mathrm{Beta}^{-1}(q; a, b)$ is the value p solving $I_p(a, b) = q$, i.e. the inverse of the map $p \mapsto I_p(a, b)$. Plugging this into (10) gives

$$I_{p^\star}\big(x+1,\ n(\lambda) - x\big) = 1 - \delta, \tag{12}$$

so $p^\star$ is the $(1 - \delta)$ quantile of the $\mathrm{Beta}\big(x + 1,\ n(\lambda) - x\big)$ distribution:

$$p^\star = \mathrm{Beta}^{-1}\big(1 - \delta;\ x + 1,\ n(\lambda) - x\big). \tag{13}$$

This equals $\widehat{R}^+(\lambda)$ by definition. For the edge cases: when $x = 0$, $F(p) = (1 - p)^{n(\lambda)}$, so $F(p^\star) = \delta$ gives $p^\star = 1 - \delta^{1/n(\lambda)}$; when $x = n(\lambda)$, $F(p) = \Pr\big(\mathrm{Bin}(n(\lambda), p) \leq n(\lambda)\big) = 1 \geq \delta$ for every p, so the feasible set is $[0, 1]$ and $\widehat{R}^+(\lambda) = 1$. $\qquad\square$

E Proof: Unchanged Threshold

Theorem E.1 (Restatement of Theorem 3.2). *Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ. Fix $\delta \in (0, 1)$ and an integer $n \geq 1$, and let $x \in \{0, 1, \ldots, n\}$, $\epsilon \in (0, 1]$, and $b \in [1, \infty)$. Then:*

(a) Decrease in value. $\alpha(\epsilon, b) \leq \alpha(1, 1)$ when $\epsilon \in (0, 1]$ and $b \in [1, \infty)$.

(b) Normal approximation for the amount of decrease. Let $z := \Phi^{-1}(1 - \delta)$, where Φ is the standard Normal cumulative distribution function. In the large–sample interior regime $n \to \infty$ with $x/n \to r \in (0, 1)$ and ϵ, b, δ held fixed, the decrease of the risk under the same level of tolerance satisfies

$$\alpha(1, 1) - \alpha(\epsilon, b) = \left(\frac{x+1}{n+1} - \frac{\epsilon x + 1}{bn + 1} \right) + z \left[\sqrt{\frac{(x+1)(n-x)}{(n+1)^2(n+2)}} - \sqrt{\frac{(\epsilon x + 1)(bn - \epsilon x)}{(bn + 1)^2(bn + 2)}} \right] + O\left(\frac{1}{n} \right). \tag{14}$$

Proof. We use the Beta quantile function Beta^{-1} to represent the variable $\widehat{R}^+(\lambda)$, which is equivalent to the risk α, when $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ. We then use the Normal approximation to the Beta distribution to evaluate the decrease of α by definition. For the convenience of statement of our theorems, we define $\alpha(\epsilon, b)$ as the value of risk bound α when the obtained λ satisfies $n(\lambda) = bn$ and incorrectly answered queries among $n(\lambda)$ is $x(\lambda) = \epsilon x$, given the δ fixed. (a) Notice that we assume that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ. Let us suppose that λ_0 satisfies that $n(\lambda_0) = bn$ and $x(\lambda_0) = \epsilon x$. By Algorithm 1, this shows that $\widehat{R}^+(\lambda_0) = \alpha(\epsilon, b)$.

From Lemma D.1, we know that

$$\alpha(\epsilon, b) := \mathrm{Beta}^{-1}\big(1 - \delta;\ \epsilon x + 1,\ bn - \epsilon x\big). \tag{15}$$

Let $p_1 = \mathrm{Beta}^{-1}(1 - \delta;\ x + 1,\ n - x)$. Then, by the property of Beta distribution, $\Pr\big(\mathrm{Bin}(n, p_1) \leq x\big) = \delta$. It follows that,

$$\Pr\big(\mathrm{Bin}(bn, p_1) \leq \epsilon x\big)\ \leq\ \Pr\big(\mathrm{Bin}(n, p_1) \leq x\big) = \delta, \tag{16}$$

because lowering the threshold ($\epsilon x \leq x$) and increasing trials ($bn \geq n$) makes the left tail event rarer. Let us assume that $p_2 = \mathrm{Beta}^{-1}(1 - \delta;\ \epsilon x + 1,\ bn - \epsilon x)$. From the proof of Lemma D.1, it is equivalent to that $\Pr\big(\mathrm{Bin}(bn, p_2) \leq \epsilon x\big) = \delta$. It follows that $\Pr\big(\mathrm{Bin}(bn, p_2) \leq \epsilon x\big) = \delta \geq \Pr\big(\mathrm{Bin}(bn, p_1) \leq \epsilon x\big)$, which implies that $p_2 \leq p_1$. Hence the new upper bound $p_2 = \mathrm{Beta}^{-1}(1 - \delta;\ \epsilon x + 1,\ bn - \epsilon x)$ satisfies $p_2 \leq p_1$. This shows the statement (a).

(b) Write

$$\mu_{\epsilon, b} := \frac{\epsilon x + 1}{bn + 1}, \qquad \sigma_{\epsilon, b} := \sqrt{\frac{(\epsilon x + 1)(bn - \epsilon x)}{(bn + 1)^2(bn + 2)}}. \tag{17}$$

In the large–sample interior regime ($n \to \infty$ with $x/n \to r \in (0, 1)$ and ϵ, b, δ fixed), the shape parameters $\epsilon x + 1$ and $bn - \epsilon x$ both diverge while $\frac{\epsilon x + 1}{bn + 1}$ stays bounded away from 0 and 1, so the standardized $\mathrm{Beta}(\epsilon x + 1,\ bn - \epsilon x)$ law converges to $\mathcal{N}(0, 1)$ with skewness $O\big((bn + 1)^{-1/2}\big)$ (Johnson et al., 1995, Ch. 25). By the Cornish–Fisher quantile expansion (Cornish & Fisher, 1938; Hall, 1992),

$$\mathrm{Beta}^{-1}\big(1 - \delta;\ \epsilon x + 1,\ bn - \epsilon x\big)\ =\ \mu_{\epsilon, b}\ +\ z\,\sigma_{\epsilon, b}\ +\ O\!\left(\frac{1}{n}\right). \tag{18}$$

Equivalently, $\widehat{R}^+(\lambda)$ is the upper Clopper–Pearson limit for a binomial proportion, and the display is its standard Normal (Edgeworth) approximation (Brown et al., 2001). Calculate $\alpha(1, 1) - \alpha(\epsilon, b)$ demonstrate the result of theorem. $\square$

F Proof: Unchanged Used-Queries

Other than the case that the threshold remains unchanged, which is analyzed above, another case may be that when the user want the same number of queries to be covered by the Weak LLM during two rounds of queries (before and after adding strategies), one of which has a better Weak LLM. Such a case controls the cost. This case considers the influence of a better Weak LLM to our pipeline. In this case, we instead assume that $n(\lambda) = n(\lambda')$, and abbreviate them as n for simplicity, which ensures the same coverage of Weak LLM. The number of wrongly answered queries before and after getting a better Weak LLM are denoted by x and ϵx, and we still estimate the decrease of α under the same level of tolerance δ. We give an approximation on the change rate of the risk bound with respect to the proportion of decrease of errors. We denote by $\alpha(\epsilon)$ the $\alpha(\epsilon, b = 1)$ for simplicity, and present the analysis in Theorem F.1.

Theorem F.1. *Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ. Fix $\delta \in (0, 1)$ and an integer $n \geq 1$, and let $x \in \{0, 1, \ldots, n\}$ and $\epsilon \in (0, 1]$. Then:*

(a) Exact monotonicity. $\alpha(\epsilon)$ is strictly increasing in ϵ. In particular, for any $\epsilon \in (0, 1)$,

$$\alpha(\epsilon)\ <\ \alpha(1). \tag{19}$$

(a) GSM-Plus   (b) MetaMath   (c) Nasa-History-MCQ

Figure 4: Accuracy as a function of the confidence threshold for the base Weak LLM and for the Weak LLM within the Inter-Cascade using random and retrieval strategies across three benchmarks.

(b) Normal approximation for the amount of decrease. *Let* $z := \Phi^{-1}(1-\delta)$. *In the large–sample interior regime* $n \to \infty$ *with* $x/n \to r \in (0,1)$ *and* δ *fixed, a first–order expansion in* $(1-\epsilon)$ *gives, as* $\epsilon \to 1$,

$$\alpha(1) - \alpha(\epsilon) = (1-\epsilon)\left[\frac{x}{n+1} + \frac{z}{2(n+1)\sqrt{n+2}}\frac{x(n-1-2x)}{\sqrt{(x+1)(n-x)}}\right] + o(1-\epsilon). \tag{20}$$

Hence the decrease is approximately linear in $(1-\epsilon)$ *with the coefficient in brackets; in particular, when* $x \leq n/2$ *the variance term is nonnegative and the decrease is at least* $(1-\epsilon)\,x/(n+1)$ *to first order.*

Proof. (a) Similar to the proof of the statement (a) of Theorem 3.2, increasing x moves mass to the right in the Binomial, so the lower-tail CDF in p decreases and its $(1-\delta)$ quantile increases; with n fixed this is equivalent to $\alpha(\epsilon)$ being strictly increasing in ϵ.

(b) Similar to the proof of the statement (a) of Theorem 3.2, notice that

$$\alpha(\epsilon) := \mathrm{Beta}^{-1}\big(1-\delta;\ \epsilon x + 1,\ n - \epsilon x\big). \tag{21}$$

For $i = \epsilon x + 1$, $j = n - \epsilon x$, the Beta (i,j) mean and variance are $\mu_\epsilon = i/(i+j)$ and $\sigma_\epsilon^2 = ij/[(i+j)^2(i+j+1)]$. By the Cornish–Fisher quantile expansion in the large–sample interior regime (Cornish & Fisher, 1938; Hall, 1992; Johnson et al., 1995) (equivalently, the Normal/Edgeworth approximation to the Clopper–Pearson limit (Brown et al., 2001)), $\alpha(\epsilon) = \mu_\epsilon + z\sigma_\epsilon + O(1/n)$. Differentiate at $\epsilon = 1$ to obtain the first-order change:

$$\frac{d\mu_\epsilon}{d\epsilon}\bigg|_{\epsilon=1} = \frac{x}{n+1} \tag{22}$$

$$\frac{d\sigma_\epsilon}{d\epsilon}\bigg|_{\epsilon=1} = \frac{1}{2(n+1)\sqrt{n+2}} \cdot \frac{(n-1-2x)x}{\sqrt{(x+1)(n-x)}}. \tag{23}$$

A first-order Taylor expansion around $\epsilon = 1$ yields the displayed approximation. $\qquad\square$

G   Confidence Distribution

Figures 4 and 5 present results for the GSM-Plus, MetaMath, and Nasa-History-MCQ datasets, complementing the GSM-Symbolic analyses in the main text.

Figure 4 shows accuracy as a function of the confidence threshold for the base Weak LLM and for the Weak LLM within the Inter-Cascade using random and retrieval strategies. For each threshold, only queries with confidence equal to or above the threshold are considered, and accuracy is calculated as the proportion of correct predictions. Across the reasoning datasets (GSM-Plus and MetaMath), the Inter-Cascade with retrieval strategies consistently improves accuracy over the baseline and random-strategy variants. For the factual non-reasoning dataset (Nasa-History-MCQ), the Inter-Cascade achieves comparable performance.

Figure 5: Confidence histograms for three benchmarks. Columns correspond to (a)(d)(g) the base Weak LLM, (b)(e)(h) the Weak LLM within the Inter-Cascade using random strategies, and (c)(f)(i) the Weak LLM within the Inter-Cascade using retrieval strategies. Across all datasets, the Inter-Cascade with retrieval strategies concentrates probability mass near high confidence (0.9–1.0), while the base and random-strategy variants place more mass at lower confidence levels.

Figure 5 depicts the distribution of query confidence for the three benchmarks. Across all datasets, the Inter-Cascade with retrieval strategies concentrates probability mass near high confidence (0.9–1.0), whereas the base and random-strategy variants place more mass at lower confidence levels. These results further confirm that providing strategies helps the Weak LLM not only produce more accurate predictions but also better calibrate its confidence.

H   Full Description of Benchmarks

GSM-Symbolic. The GSM-Symbolic benchmark, released by Apple's team (Mirzadeh et al., 2025), is a structured variant of GSM8K (Cobbe et al., 2021). Unlike traditional benchmarks such as GSM8K, which present problems in a plain context, GSM-Symbolic reformulates problems into a more structured and abstract format following a symbolic template, providing a more reliable measure of models' reasoning capabilities. The dataset contains $12,500$ grade-school math problems. We randomly sample $1,250$ problems as the calibration set for threshold computation and use the remaining $11,250$ problems as the test set. The prompt template and an example problem are provided in Appendix L.

GSM-Plus. GSM-Plus (Li et al., 2024) is derived from the $1,319$ test questions in GSM8K by introducing eight types of question variations: numerical substitution, digit expansion, integer-decimal-fraction conversion, adding operation, reversing operation, problem understanding, distractor insertion, and critical thinking. GSM-Plus thus comprises a total of $10,552$ question variations. We randomly sample $1,048$ problems as the

calibration set for threshold computation and use the remaining $9,504$ problems as the test set. The prompt template and an example problem are provided in Appendix L.

MetaMath. MetaMath (Yu et al., 2024) is a dataset generated by bootstrapping the mathematical benchmarks GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). The augmentation is performed in both forward and backward directions. In the forward direction, MetaMath contains the original and LLM-rephrased questions, while in the backward direction, it includes self-verification questions and FOBAR questions (Jiang et al., 2024), resulting in a total of $395K$ diverse problems. For our experiments, we randomly select $1,000$ problems as the calibration set for threshold computation and use $20,000$ additional problems as the test set. The prompt template and an example problem are provided in Appendix L.

NASA-History-MCQ. NASA-History-MCQ (Fleith, 2025) is a multiple-choice question benchmark on the history of NASA. It contains $7.47K$ questions, and each question provides four answer choices. We randomly sample $1,000$ problems as the calibration set for threshold computation and use the remaining $6,469$ problems as the test set. The prompt template and an example problem are provided in Appendix L.

BarExamQA. BarExamQA (Zheng et al., 2025) is a legal reasoning benchmark constructed from real U.S. bar examination questions. Each question is posed in a multiple-choice format and requires multi-step legal reasoning over complex legal fact patterns. BarexamQA contains a total of 954 problems, we randomly sample 95 problems as the calibration set for threshold computation and remaining 859 as the test set.

BigBench Hard. BIG-Bench Hard (Suzgun et al., 2022) is a subset of 23 particularly challenging BIG-Bench tasks for which no prior result from (Srivastava et al., 2022) has outperformed the average human-rater score. It is a diverse benchmark designed to test capabilities of language models on a diverse set of crowd-sourced tasks. The benchmark aims to focus on the problems that beyond the capabilities of existing LLMs. We use 5412 problems as test set and 599 problems as calibration set for threshold computation. The calibration set are selected from each tasks with the same proportion.

GSM8K. GSM8K (Cobbe et al., 2021) is a widely used grade-school math word problem benchmark designed to evaluate multi-step numerical reasoning. The dataset contains 7473 training questions and 1719 test questions, with each problem requiring several arithmetic operations and logical reasoning steps to reach the final answer. Following standard practice, we use problems in calibration set for threshold computation and use the remaining problems as the test set.

MedMCQA. MedMCQA (Pal et al., 2022) is a large-scale multiple-choice question benchmark in the medical domain. It covers high-quality AIIMS and NEET PG entrance exam MCQs covering 2400 healthcare topics and 21 medical subjects. It contains over $194,000$ questions, each with four answer choices and a single correct answer. We randomly sample $2,000$ problems as the calibration set for threshold computation and use 8000 additional problems as the test set.

I Full Description of Token and API Cost Analysis

The full analysis on the token consumption, including input tokens and output token for the four benchmarks presented in main text is shown in Table 9.

Table 9: Token and API cost changes across datasets for Inter-Cascade compared with Jung's pipeline.

Benchmark	Weak LLM Tokens			Strong LLM Tokens			Token Price
	Total	Input	Output	Total	Input	Output	
GSM-Symb.	+147.66%	+148.80%	-17.10%	-47.80%	-45.80%	-51.32%	-49.63%
GSM-Plus	+145.96%	+147.11%	-3.56%	-29.95%	-29.51%	-30.90%	-30.41%
Meta.(20K)	+127.90%	+128.66%	-1.38%	-52.18%	-52.20%	-52.12%	-52.15%
NASA-Hist.	+132.58%	+133.40%	0.99%	-15.47%	-15.22%	-16.07%	-15.75%

Table 10: Accuracies of the base LLMs on extensive benchmarks

Dataset	LLM	Accuracy	Dataset	LLM	Accuracy
GSM8K	gpt-3.5-turbo	31.46%	BigBench	gpt-3.5-turbo	49.75%
	gemini-2.0-flash	74.83%		gemini-2.0-flash	78.80%
BarExamQA	gpt-3.5-turbo	48.42%	MedMCQA	gpt-3.5-turbo	62.80%
	gemini-2.0-flash	78.95%		gemini-2.0-flash	83.05%

Table 11: Results across extensive datasets using different pipelines. "Jung et al." denotes Jung et al.'s LLM-Cascade and "Our (Retrieval)" denotes the Inter-Cascade with similarity-based retrieval. The number of strategies is fixed at $k = 2$ for both Inter-Cascade settings. Metrics reported are Pipeline Accuracy (Pipeline Acc.), Strong LLM Call Rate (Strong Call), and Coverage Rate (Cov.). (a) GSM8K: For the Strong LLM, $\alpha_s = 0.2, \delta_s = 0.8, \lambda_s = 0.44$. For the Weak LLM, $\alpha_w = 0.5, \delta_w = 0.5, \lambda_w = 0.49$. (b) BigBench: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.4, \delta_w = 0.6, \lambda_w = 0.61$. (c) BarExamQA: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.5, \delta_w = 0.5, \lambda_w = 0.51$. (d) MedMCQA: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.3, \delta_w = 0.8, \lambda_w = 0.69$.

Data	Pipeline	Pipeline Acc. (%)↑	Strong Call (%)↓	Cov. (%)
GSM8K	Jung et al.	59.02	37.03	95.95
	Our (Retrieval)	**60.62**	**35.46**	**96.05**
BigBench	Jung et al.	64.14	33.04	100.00
	Our (Retrieval)	**64.32**	**23.84**	100.00
BarExamQA	Jung et al.	57.39	23.17	100.00
	Our (Retrieval)	**58.67**	**21.65**	100.00
MedMCQA	Jung et al.	71.69	18.74	100.00
	Our (Retrieval)	**75.65**	**2.60**	100.00

J   Additional Experiments on Other Benchmarks

Although the Inter-Cascade diagram is motivated by the real-world scenarios that contain similar or repeated tasks, we also provide the result of our Inter-Cascade on extensive benchmarks that are more diverse and do not contain explicit sample variants: GSM8K (Cobbe et al., 2021), BigBench Hard (Suzgun et al., 2022), BarExamQA (Zheng et al., 2025) and MedMCQA (Pal et al., 2022). The full description of those benchmarks are in Appendix H. We firstly test the accuracy of each single LLM on those benchmarks and the result is in Table 10.

Inter-Cascade vs. Jung's LLM Cascade. We evaluate our *Inter-Cascade* pipeline and Jung's method, as shown in Table 11. Our method outperforms Jung's, with a $0.18\% - 3.96\%$ increase in Pipeline Accuracy. The Strong LLM Call Rate is reduced on all benchmarks, with reductions ranging from 1.52% to 16.14%. Compared with the results on GSM-Symbolic, GSM-Plus and MetaMath benchmarks, the accuracy

Figure 6: Effect of number of Strategies on pipeline accuracy for GSM-Symbolic Benchmark

Table 12: Results on Weak LLM across extensive datasets. Reported metrics are Weak LLM Accuracy (Weak Acc.) and Weak Correct Accepted (Weak Corr. Accpt.). Parameter settings are the same as in Table 11.

Data	Pipeline	Weak Acc. (%)↑	Weak Corr. Accpt. (%)↑
GSM8K	Jung et al.	37.06	33.38
	Our (Retrieval)	**39.30**	**35.62**
BigBench	Jung et al.	49.02	39.34
	Our (Retrieval)	**49.93**	**46.60**
BarExamQA	Jung et al.	47.50	39.81
	Our (Retrieval)	**51.22**	**43.31**
MedMCQA	Jung et al.	64.95	58.16
	Our (Retrieval)	**74.51**	**73.72**

Table 13: Token and API cost changes across extensive datasets for Inter-Cascade compared with Jung's pipeline.

Benchmark	Weak LLM Tokens			Strong LLM Tokens			Token Price
	Total	Input	Output	Total	Input	Output	
GSM8K	+115.89%	+116.56%	-2.27%	-3.25%	-4.10%	-1.28%	-2.33%
BigBench	+134.53%	+135.32%	-5.47%	-26.37%	-30.90%	-19.67%	-22.70%
BarExamQA	+216.37%	+216.90%	+0.12%	-5.70%	-5.39%	-6.28%	-5.98%
MedMCQA	+129.64%	+130.70%	-0.16%	-84.74%	-85.58%	-83.17%	-83.94%

improvement is not that large, but the more important part is that our Inter-Cascade can still reach a better trade-off between accuracy and cost since our method still remarkably reduce the usage of Strong LLM. These results indicate that *Inter-Cascade* pipeline is also beneficial across different categories of tasks on diverse benchmarks.

Impact of Inter-Cascade on Weak LLM. Having examined the overall pipeline improvements, including Pipeline Accuracy and Strong LLM Call Rate reduction, we now investigate how our proposed *Inter-Cascade* affects the Weak LLM. As shown in Table 12, our Weak LLM still outperforms the Weak LLM in the other pipeline across all benchmarks. The improvements on *Weak Accuracy* are between 0.91% and 9.56% and the improvements on *Weak Correct Accepted* are between 2.24% and 15.56%. The results implies that even though we test our Inter-Cascade on diverse benchmarks, retrieving most similar problems and solution strategies can still help boosting the performance and confidence of Weak LLM.

According to experiment results for extensive benchmarks, it shows that Inter-Cascade not only work for tasks that contain constructive similarity, but also help in more general and diverse cases, since explicit or implicit similarity occurs everywhere and the pipeline in our Inter-Cascade take the advantage of the similarity nature of daily tasks.

Token and API Cost Savings. The results of analysis on cost and latency for extensive benchmarks are attached in Table 13 and Table 14. The tendency is similar: integrating with strategies, the token usages on

Table 14: Processing Latency and Strategy Repository Size across extensive datasets. Retrieval refers to the time spent on strategies matching and ranking. Generation refers to time spent on generating answer via API.

Benchmark	Tested Samples	Our			Jung et al.	Repository Size
		Total	Retrieval	Generation	Total	
GSM8K	7473	1.344s	0.005s	1.339s	1.216s	6.3MB
BigBench	5412	1.456s	0.004s	1.452s	1.227s	3.4MB
BarExamQA	859	1.686s	0.254s	1.432s	1.312s	1.1MB
MedMCQA	8000	0.975s	0.004s	0.971s	0.970s	6.3MB

Table 15: Accuracies of new pair of base LLMs on GSM-Symbolic Benchmark

Dataset	LLM	Accuracy
GSM-Symbolic	gemini-2.0-flash	69.36%
	gemini-2.5-flash	89.28%

Table 16: New LLM Pairs (Weak LLM: Gemini-2.0-flash; Strong LLM: Gemini-2.5-flash) Results on GSM-Symbolic dataset using different pipelines. "Jung et al." denotes Jung et al.'s LLM-Cascade and "Our (Retrieval)" denotes the Inter-Cascade with similarity-based retrieval. The number of strategies is fixed at $k = 2$ for both Inter-Cascade settings. Metrics reported are Pipeline Accuracy (Pipeline Acc.), Strong LLM Call Rate (Strong Call), and Coverage Rate (Cov.). GSM-Symbolic: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w = 0.2, \delta_w = 0.8, \lambda_w = 0.47$.

Data	Pipeline	Pipeline Acc. (%)↑	Strong Call (%)↓	Cov. (%)
GSM-Symbolic	Jung et al.	79.10	19.10	100.00
	Our (Retrieval)	**85.50**	**9.90**	100.00

Weak LLM increase between 115.89% and 216.37%, but since the *Strong Call* decrease on all benchmark, the token usages on Strong LLM decrease between 1.28% and 83.17% and therefore, we can save 2.33%-83.94% money on API price. On the other hand, the average latency change on each query is between $0.005\,\text{s}$ and $0.374\,\text{s}$ on different benchmarks, which is acceptable to the user experience.

K   Additional Ablation Results

To better evaluate the performance and generalization capacity of Inter-Cascade, we set up extra ablation studies in this section.

K.1   Effect of Strategies Number

To evaluate the effect the number of strategies we matched for each queries, we test the pipeline accuracy with different number of strategies that used for integrating with the input of Weak LLM. The result in Figure 6 shows that the trend of pipeline accuracy is increasing first, reaching peak and then decreasing along with the number of strategies. The result makes sense because too few strategies might not retrieve the best strategy in repository, while too many strategies might distract the answer from certain query question, furthermore, there is a chance that the longer contexts may exceed the the maximum limit of the input context window. Both factors might undermine the performance of the pipeline accuracy. In our experiment on GSM-Symbolic benchmark, the empirical best number of strategies k is 2.

K.2   Results on New LLM Pairs

To show that our Inter-Cascade is a framework that work general multiple LLM collaboration systems, we also test the result on different choice of Weak LLM and Strong LLM. We switch our Weak LLM to Gemini-2.0-flash and switch our Strong LLM to Gemini-2.5-flash. The results on single LLM are in Table 15. We also analyze the performance on those metrics: Pipeline Accuracy, Strong Call Rate, Weak Accuracy and Weak Correct Accept in Table 16 and Table 17. The results shows that although we test on different pairs of Weak LLM and Strong LLM, the trend doesn't change: Inter-Cascade would help improve the accuracy of Weak LLM, pipeline accuracy, reduce the usage of Strong LLM, reaching a better trade-off between the Accuracy and Cost in LLM Cascade systems.

L   Prompt Templates and Examples

Table 18 and Table 19 present the strategy-free prompt templates for the four datasets, along with one example question per dataset. Table 20 - Table 23 show the strategy-based prompt templates and example inputs for each dataset. In our experiments, the number of strategies is set to $k = 2$; these strategies and

Table 17: New LLM Pairs (Weak LLM: Gemini-2.0-flash; Strong LLM: Gemini-2.5-flash) Results on Weak LLM across GSM-Symbolic dataset. Reported metrics are Weak LLM Accuracy (Weak Acc.) and Weak Correct Accepted (Weak Corr. Accpt.). Parameter settings are the same as in Table 16.

Data	Pipeline	Weak Acc. (%)↑	Weak Corr. Accpt. (%)↑
GSM-Symbolic	Jung et al.	64.20	63.40
	Our (Retrieval)	**77.00**	**76.80**

Table 18: Strategy-free prompt template with example questions from GSM-Symbolic, GSM-Plus, and MetaMath

---

Prompt Template:

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Example:
[Question]: $x + y = 10, y = 4$, what is x?
[Strategy]: To solve for x, isolate x by subtracting y from both sides of the equation. $x = 10 - y = 10 - 4 = 6$.
[Answer]: 6

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

GSM-Symbolic Example Question:

[Question]: A fog bank rolls in from the ocean to cover a city. It takes 495 minutes to cover every 95 miles of the city. If the city is 95 miles across from the oceanfront to the opposite inland edge, how many minutes will it take for the fog bank to cover the whole city?

GSM-Plus Example Question:

[Question]: A clothing store has 60 white shirts and 70 floral shirts. Two-thirds of the white shirts have collars, and 30 of the floral shirts have buttons. How many more floral shirts with no buttons are there than white shirts with no collars?

MetaMath Example Question:

[Question]: Lara bought 52 stems of flowers at the flower shop. She gave 15 flowers to her mom and gave x more flowers than she gave to her mom to her grandma. She put the rest in a vase. How many stems of flowers did Lara put in the vase? If we know the answer to the above question is 16, what is the value of unknown variable x?

---

their corresponding answers are generated by the Strong LLM. Since the pipeline operates without human intervention, all strategies that exceed the Strong LLM confidence threshold λ_s are accepted. Consequently, the Repo may contain incorrect strategies or answers. Nonetheless, the results in Table 3 and Table 4 demonstrate the effectiveness of λ_s and the robustness of our proposed Inter-Cascade pipeline.

Table 19: Strategy-free prompt template with example question from NASA-History-MCQ

Prompt Template:

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the single capital letter of the correct option (exactly one of: A, B, C, D). Do NOT include the option text or any explanation.

Example:
[Question]: Which of the following was a direct result of the Treaty of Versailles (1919)?
A. The outbreak of World War II
B. The United States declaring war on Germany and entering World War I
C. The establishment of the Fascist regime in Italy
D. The creation of the League of Nations
[Strategy]: Recall that the treaty ended World War I and included peacekeeping measures. Compare options with the treaty: only the League of Nations was directly established by it.
[Answer]: D

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

Nasa-History-MCQ Example Question:

[Question]: What three distinct stages comprise current psychological support protocols for US astronauts?
A. Individual, group, family
B. Training, mission, debriefing
C. Cognitive, emotional, behavioral
D. Preflight, in-flight, postflight

Table 20: Strategy-based prompt template with example input from GSM-Symbolic

---

Prompt Template:

Using the following problem examples:
{strategy}

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

GSM-Symbolic Example Input:

Using the following problem examples:
1. [Question]: A fog bank rolls in from the ocean to cover a city. It takes 93 minutes to cover every 2 miles of the city. If the city is 24 miles across from the oceanfront to the opposite inland edge, how many minutes will it take for the fog bank to cover the whole city?
[Strategy]: First, determine how many 2-mile segments are in the 24-mile city. Divide the total distance of the city by the length of each segment. Then, multiply the number of segments by the time it takes to cover each segment to find the total time.
[Answer]: 1116

2. [Question]: A fog bank rolls in from the ocean to cover a city. It takes 114 minutes to cover every 3 miles of the city. If the city is 99 miles across from the oceanfront to the opposite inland edge, how many minutes will it take for the fog bank to cover the whole city?
[Strategy]: First, determine the rate at which the fog bank covers the city in miles per minute. Then, multiply this rate by the total distance of the city to find the total time it takes to cover the city. The rate is 3 miles / 114 minutes = $1/38$ miles per minute. The total time is ($1/38$ miles/minute) $*$ 99 miles $= 99/38$ minutes. Simplify the fraction $99/38 = 2.60526315789$. Multiply 114 by $99/3$ to get the answer $114 * (99/3) = 114 * 33 = 3762$.
[Answer]: 3762

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: A fog bank rolls in from the ocean to cover a city. It takes 495 minutes to cover every 95 miles of the city. If the city is 95 miles across from the oceanfront to the opposite inland edge, how many minutes will it take for the fog bank to cover the whole city?
[Strategy]:
[Answer]:

Table 21: Strategy-based prompt template with example input from GSM-Plus

Prompt Template:

Using the following problem examples:
{strategy}

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

GSM-Plus Example Input:

Using the following problem examples:
1. [Question]: A clothing store has some white shirts and 50 floral shirts. Half of the white shirts have collars, and 20 of the floral shirts have buttons. How many more floral shirts with no buttons are there than white shirts with no collars?
[Strategy]: Let W be the number of white shirts. The number of white shirts with collars is $W/2$, so the number of white shirts without collars is $W - W/2 = W/2$. The number of floral shirts is 50. The number of floral shirts with buttons is 20, so the number of floral shirts without buttons is $50 - 20 = 30$. We want to find the difference between the number of floral shirts without buttons and the number of white shirts without collars, which is $30 - W/2$. However, we are not given the number of white shirts, W. Therefore, we cannot determine the exact numerical value of the difference.
[Answer]: None

2. [Question]: A clothing store has 40 white shirts and 50 floral shirts. Half of the white shirts have collars, and some of the floral shirts have buttons. There are 10 more floral shirts with no buttons than white shirts with no collars. How many floral shirts have buttons?
[Strategy]: 1. Calculate the number of white shirts with collars: $40/2 = 20$. 2. Calculate the number of white shirts without collars: $40 - 20 = 20$. 3. Calculate the number of floral shirts without buttons: $20 + 10 = 30$. 4. Calculate the number of floral shirts with buttons: $50 - 30 = 20$.
[Answer]: 20

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: A clothing store has 60 white shirts and 70 floral shirts. Two-thirds of the white shirts have collars, and 30 of the floral shirts have buttons. How many more floral shirts with no buttons are there than white shirts with no collars?
[Strategy]:
[Answer]:

Table 22: Strategy-based prompt template with example input from MetaMath

Prompt Template:

Using the following problem examples:
{strategy}

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

MetaMath Example Input:

Using the following problem examples:
1. [Question]: Martha has x crayons. She lost half of them, so she bought a new set of 20 crayons. Martha has 29 crayons in total after the purchase. What is the value of unknown variable x?
[Strategy]: Let x be the number of crayons Martha initially had. She lost half of them, so she had $x/2$ crayons left. Then she bought 20 new crayons, so she had $x/2 + 20$ crayons. We are given that she has 29 crayons in total, so $x/2 + 20 = 29$. Subtracting 20 from both sides gives $x/2 = 9$. Multiplying both sides by 2 gives $x = 18$.
[Answer]: 18

2. [Question]: Justin and Sabrina at Pine Grove Elementary have 50 pencils combined. Justin has x more than twice as many pencils as Sabrina. Sabrina has 14 pencils. What is the value of unknown variable x?
[Strategy]: Let J be the number of pencils Justin has and S be the number of pencils Sabrina has. We are given that $J + S = 50$ and $S = 14$. We are also given that Justin has x more than twice as many pencils as Sabrina, which can be written as $J = 2S + x$. We can substitute $S = 14$ into the first equation to find J: $J + 14 = 50$, so $J = 50 - 14 = 36$. Now we can substitute $J = 36$ and $S = 14$ into the second equation: $36 = 2(14) + x$, so $36 = 28 + x$. Solving for x, we get $x = 36 - 28 = 8$.
[Answer]: 8

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value.

Now answer this question:
[Question]: Lara bought 52 stems of flowers at the flower shop. She gave 15 flowers to her mom and gave x more flowers than she gave to her mom to her grandma. She put the rest in a vase. How many stems of flowers did Lara put in the vase? If we know the answer to the above question is 16, what is the value of unknown variable x?
[Strategy]:
[Answer]:

Table 23: Strategy-based prompt template with example input from NASA-History-MCQ

Prompt Template:

Using the following problem examples:
{strategy}

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the single capital letter of the correct option (exactly one of: A, B, C, D). Do NOT include the option text or any explanation.

Now answer this question:
[Question]: {question}
[Strategy]:
[Answer]:

Nasa-History-MCQ Example Input:

Using the following problem examples:
1. [Question]: Beyond communication and care packages, what specific types of hardware or software aid psychological well-being during long-duration spaceflights?
A. Specialized dietary supplements to combat isolation
B. Automated exercise routines tailored to reduce stress
C. Psychological support hardware and software
D. Advanced life support systems with mood stabilizers
[Strategy]: The question asks about specific hardware or software that aids psychological well-being during long-duration spaceflights, beyond communication and care packages. We need to evaluate each option to see if it fits this description. Option A focuses on dietary supplements, which are not hardware or software. Option B describes automated exercise routines, which could involve software and hardware. Option C is too general, simply restating the question. Option D focuses on life support systems with mood stabilizers, which are not necessarily hardware or software designed specifically for psychological well-being. Therefore, option B is the most specific and relevant answer.
[Answer]: B

2. [Question]: What is the anticipated effect of constraints inherent in lunar and Martian missions on psychological support approaches?
A. Greater emphasis on real-time communication with Earth-based support teams
B. Increased reliance on virtual reality and AI companionship to mitigate isolation
C. A shift towards highly individualized psychological profiles and tailored interventions
D. A return to the mindset and strategies of earlier explorers and their families
[Strategy]: The question asks about the impact of constraints in lunar and Martian missions on psychological support. These constraints include isolation, limited resources, communication delays, and the need for self-sufficiency. Considering these limitations, the most likely effect would be a greater reliance on technologies that can provide support in the absence of immediate Earth-based assistance and a need for personalized approaches due to the unique challenges faced by each astronaut. Options A and D are less likely because of communication delays and the differences between modern space missions and earlier explorations. Option B is plausible, but option C is more comprehensive as it addresses the need for personalized support, which is crucial given the constraints.
[Answer]: C

Based on the question below, please strictly follow this format when answering:
1. Start with [Strategy] section explaining the general approach for solving similar problems;
2. End with [Answer] section containing ONLY the single capital letter of the correct option (exactly one of: A, B, C, D). Do NOT include the option text or any explanation.

Now answer this question:
[Question]: What three distinct stages comprise current psychological support protocols for US astronauts?
A. Individual, group, family
B. Training, mission, debriefing
C. Cognitive, emotional, behavioral
D. Preflight, in-flight, postflight
[Strategy]:
[Answer]: