NOT ONLY A HELPER, BUT ALSO A TEACHER: INTER-ACTIVE LLM CASCADE

Anonymous authorsPaper under double-blind review

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

030

032033034

035

037

040

041

042

043

044

045

046

047

048

051

052

ABSTRACT

Large Language Models (LLMs) vary widely in their capabilities, with larger models often having better performance but higher cost: choosing an LLM model often involves trading off performance and cost. The *LLM Cascade* is a paradigm that defers difficult queries from weak/cheap to strong/expensive models. This approach is nonadaptive: the deferral decision is trained offline. When confronted with similar or repeated queries, the LLM Cascade may then repeatedly consult the expensive model and incur higher cost. To improve the cascading efficiency, we propose *Inter-Cascade*, an online and interactive LLM Cascade that extends the role of strong model from a backup helper to a long-term teacher. In our system, when a strong model resolves a difficult query, it also distills its solution into a generalized, reusable problem-solving strategy that boosts the weak model on subsequent queries. Adding strategies to queries enables the weak model to dynamically improve its performance over time, avoiding computationally and time-intensive fine-tuning. Empirically, compared with standard LLM Cascade baselines across multiple benchmarks, the Inter-Cascade significantly improves the accuracy of the weak model (by up to 33.06 absolute percentage points) and the overall system (by up to 5.53 absolute percentage points), while reducing the calls to strong models (by up to 48.05% relative reduction) and saving the corresponding fees (by up to 49.63% relative reduction). 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. LLMs with stronger performance are generally larger in size, and the converse holds as well (Kaplan et al., 2020). Larger models often achieve better performance on more challenging tasks but are correspondingly more expensive. Depending on their expected workload, cost-sensitive users may wish to use *weaker* (and cheaper) models that suffice for simple queries and reserve the use of *stronger* (and expensive) models for more complex queries. In a prototypical example, a weaker model may run on a mobile device such as a phone while a stronger model may run in a cloud-based server owned by another service: the cost of the stronger model can include latency, monetary charges, or both. Since accuracy alone is not the only performance metric of interest, practical deployment scenarios require balancing multiple objectives including efficiency, latency, reliability, and network resource usage (Zhou et al., 2024; Khatun & Brown, 2024; Gundla & Atthuluri, 2025; Yan & Ding, 2025; Zhou et al., 2024).

The *LLM Cascade* has emerged as a widespread LLM paradigm in which weaker models handle routine queries and *defer* uncertain cases to stronger models in a sequential order (Chen et al., 2024). Deferral depends on a *deferral function*, typically estimates a confidence score that decides when to send queries to the strong model. The current approach focuses on improving this deferral decision by training or adjusting the threshold for the confidence score to decide when to defer (Shen et al., 2024; Rayan & Tewari, 2025; Ong et al., 2025; Zellinger et al., 2025; Zellinger & Thomson, 2025; Xia et al., 2024; Nie et al., 2024; Jung et al., 2025). The strong model, weak model, and deferral function are optimized prior to deployment. After training and deployment, the system follows the same pipeline for any incoming query: processing is not adaptive to the workflow during deployment.

In this paper we improve on this "one size fits all" approach by developing a new online adaptation strategy to assist the weaker model in dynamically balancing cost and accuracy. Our key insight is to exploit input similarity so the strong model can help the weaker one adapt and handle similar queries locally. The essence of the approach is to let the weaker model learn from the stronger model online: the weaker model can build a "crib sheet" using prior queries to do prompt engineering at the input that will guide the weaker model to the correct solution locally.

As an example of the concept, consider scenarios that have clusters of similar or even identical problems. Datasets like GSM-Plus (Li et al., 2024), an extension of a math problem dataset GSM8K (Cobbe et al., 2021), contains eight variants for each problem. Identical questions are repeatedly asked: e.g., "Which is larger, 9.9 or 9.11?" Users were once eager to keep asking this question to any newly released LLM (Korzhov, 2023; Schnabel, 2024; Junco, 2025). If a weak models consistently fails on similar or recurring problems, the system must repeatedly consult the strong model each time, which is wasteful/costly. The unidirectional processing in LLM Cascade (where weak models can only offload) prevents weak models from taking advantage of *feedback* from the strong models' capabilities during the generation/inference phase. Furthermore, fine-tuning weaker models to overcome such failures is often impractical: fine-tuning requires substantial memory (e.g., finetune Qwen3-235B requires 130 GB VRAM for LoRA or 2560 GB for full-parameter fine-tuning (Yang et al., 2025)) and must re-train again when distributions shift. For API-based models, fine-tuning may not even be available.

Our approach is influenced by in-context learning (ICL) (Dong et al., 2024) or few-shot prompting (Parnami & Lee, 2022), which can partially alleviate this limitation. In those approaches, by carefully selecting demonstrations or instructions one can enhance the reasoning capacity of weaker models without retraining. However, existing ICL and few-shot methods often rely on manual prompt design or retrieval from either fixed or manually updated database, making them inflexible as query distributions evolve. The goal in this work is to develop a framework the enables cascaded LLMs to interact adaptively: the weak model can benefit from *in situ* reasoning generated by the strong model to improve its own performance during real-world query streams. As a metaphor, the weak model uses the strong model for "on the job" training to improve longer-term performance. We discuss this paradigm's relationship to other related work in the Appendix B.

Primary contributions. We improve on existing LLM Cascade approaches. (1) We propose a new framework, Inter-Cascade, for online and interactive LLM Cascades in which the strong model serves as both a backup helper and longer-term teacher. The strong model can both answer difficult queries and provide feedback that can be reused by the weak model to generalize problem-solving for future queries. These strategies are stored in a local database which the weak model uses as an auxiliary input: its performance is improved on similar future queries by leveraging the strategies generated from the strong model. In this way the strong LLM "teaches" the the weak LLM how to resolve these queries on its own. We think of this approach as a kind of in-context knowledge distillation approach to LLM Cascade using similarity-based memory. (2) We propose a theoretical model and show that without changing the deferral rule, adding strategies helps the weak LLM's confidence score better approximate its probability of correctness. This shows that using strategies can provably guarantee the higher accuracy of the overall system. (3) We compare to the LLM Cascade (Jung et al., 2025) and show that Inter-Cascade improves the overall system accuracy by up to 33.06 absolute percentage points and the overall system by up to 5.53 absolute percentage points, while reducing the usage of strong models by up to 48.05% relative reduction, This can reduce the corresponding fees by up to 49.63% relative reduction with the same guarantee of risk tolerance and error level over all benchmarks.

Crucially, our framework is general and modular: it applies to both API-only models and open-source models, and can be combined with any deferral function or any number of LLMs in cascade. We make the full implementation of Inter-Cascade under an open source license.

2 IMPROVING THE LLM CASCADE

We first describe the standard LLM Cascade (Chen et al., 2024) and revisit the accuracy bound and calibration method for the deferral threshold proposed by 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.

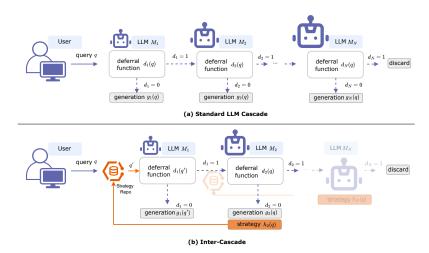


Figure 1: (a) Pipeline of standard LLM Cascade systems. (b) Pipeline of Inter-Cascade. The unique components in Inter-Cascade are painted in orange. For the sake of clarity and readability, we only present the case of two LLMs Inter-Cascade system and the scalable parts beyond two LLMs are rendered in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

2.1 STANDARD LLM CASCADE

Figure 1(a) shows the general N-LLM Cascade system (Chen et al., 2024). Each LLM $M_i: i \in [N]$ contains two key components. One is the *generation function* $g_i: \mathcal{Q} \to \mathcal{A}$, where \mathcal{Q} is the space of queries and \mathcal{A} is the space of answers. The other is *deferral function* $d_i: \mathcal{Q} \to \{0,1\}$, 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 $\preccurlyeq_{\text{wbc}}$ ("weaker but cheaper") to compare models (see Appendix C) and assume that in the cascade, $M_1 \preccurlyeq_{\text{wbc}} M_2 \preccurlyeq_{\text{wbc}} \dots \preccurlyeq_{\text{wbc}} M_N$. For each query $q \in \mathcal{Q}$, the first LLM M_1 takes the query q and gives a final answer q if deferral function q (q) of otherwise q defers this query to the next LLM q if q if q if q 1. If q 1 takes the query from q 1, it repeats the same process and so do the other LLMs except the last model q 2. Recent studies propose different deferral functions q3 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 q4 the Weak LLM and q5 the Strong LLM. One common choice of deferral function is:

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

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

with probability at least $1-\delta$, as proved in Theorem 1 of their work, where $a_{\rm true}$ 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, a fixed-sequence testing (Bauer, 1991) procedure is first conducted, which selects the largest threshold λ from a calibration set, such that $P\left(g_i(q)=a_{\rm true}\mid c(q)\geq\lambda\right)$ 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.

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

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 $n(\lambda) \leftarrow \sum_{(q,a) \in D_{\text{cal}}} \mathbf{1}\{c(q) \ge \lambda\}$

162

163

164

165

166

167

168

169

170

171

172 173 174

175 176 177

178

179

180

181

182

183

184

185

187

188

189

190 191

192

193

194

195

196

197

199 200

201

202

203

204

205

206

207

208

209

210 211

212

213

214

215

 $\hat{R}(\lambda) \leftarrow \frac{1}{n(\lambda)} \sum_{(q,a) \in D_{\text{cal}}} \mathbf{1} \{ g_i(q) \neq a_{\text{true}} \land c(q) \geq \lambda \}$ 4:

 $\hat{R}^+(\lambda) \leftarrow \sup\{R : \Pr[\operatorname{Bin}(n(\lambda), R) \le n(\lambda)\hat{R}(\lambda)] \ge \delta\}$

if $\hat{R}^+(\lambda) \leq \alpha$ then return λ 6:

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 $1-\delta$.

2.2 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 $h \colon \mathcal{Q} \to \mathcal{S}$, 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. The detailed description of strategy repository is depicted below:

Strategy Repository. The Strategy Repository Repo is formally defined as a collection of querystrategy pairs: 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 $sim: \mathcal{Q} \times \mathcal{Q} \to [0,1]$ be a ranking function. Let the Top-k indices (sorted by decreasing similarity) be

$$TopIndex(q) \triangleq (t_1, t_2, \dots, t_k),$$

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

Remark 2.1. 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.

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, 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 (1). Our algorithm can be extended to multi-LLM cases, the corresponding Algorithm 3 is shown in Appendix D.

Algorithm 2 Inter-Cascade Inference Pipeline

Discard current query q_i

216

217

218

219

220

221 222 223

224

225

226

227

228

229

230

231

232

233

234

235

236

237

238

239 240 241

242

243

244

245

246

247

248

249 250

251

253

254

255

256

257

258

259

260

261 262

263 264

265

266

267

268 269 12:

```
Input: Test set \mathcal{T} = \{q_1, \dots, 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: 0 = \text{handle locally}, 1 = \text{defer/forward}.
 1: for i \leftarrow 1 to I do
           \begin{aligned} [s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}] &\leftarrow f(q_i, \text{Repo}) \\ q_i' &\leftarrow [q_i, s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}] \end{aligned} 
                                                                \triangleright Top- k strategies matching from Repo
 3:
                                                                               > concatenate query and strategies
         if d_1(q_i') = 0 then
 4:
                                                                                             ▶ Weak LLM decision
 5:
              generate answer a_i \leftarrow g_1(q_i')
                                                                                  6:
         else
 7:
              if d_2(q_i) = 0 then
                                                                                           ▶ Defer to Strong LLM
 8:
                   s_{\text{new}} \leftarrow h(q_i)

    Strong LLM synthesizes a new strategy

 9:
                   Repo \leftarrow Repo \cup \{(q_i, s_{\text{new}})\}
                                                                  ▷ Send back strategy to Weak LLM and store
                   generate answer a_i \leftarrow g_2(x_i)
10:
                                                                                          else
11:
                                                           ▶ None of LLMs are confident to answer the query
```

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 $bn := n'(\lambda)$, $b \in [1, \infty)$, 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 2.2 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 2.2. Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ . Fix $\delta \in (0,1)$ and an integer $n \geq 1$. For $x \in \{0,1,\ldots,n\}$, $\epsilon \in (0,1]$, and $b \in [1,\infty)$. Suppose that $\min\{\epsilon x + 1, n - \epsilon x\}$ is moderately large and $1 - \delta$ is not an extreme tail, 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 Normal cumulative distribution function, when n is large enough, the decrease of the risk under same level of tolerance is given by,

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

The proof of this theorem is in Appendix F. Theorem 2.2 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 G.1 and the proof are shown in Appendix G.

3 EXPERIMENTS

3.1 BENCHMARKS

In our experiments, we use two categories of datasets. The first category consists of reasoning-focused scientific benchmarks, including *GSM-Symbolic* (Mirzadeh et al., 2025), *GSM-Plus* (Li et al., 2024), and *MetaMath* (Yu et al., 2024), selected to evaluate performance on tasks requiring logical reasoning. The second category includes factual benchmark, represented by *NASA-History-MCQ* (Fleith, 2025), chosen to assess performance on tasks with lower reasoning demands. Using both categories allows for a more comprehensive assessment across tasks of different difficulty levels and types. The detailed descriptions of selected benchmarks are in Appendix I. The prompt template and an example problem for each benchmark are provided in Appendix J.

3.2 EXPERIMENTAL SETTINGS

Inter-Cascade. On all benchmarks, Gemini-2.0-flash consistently outperforms GPT-3.5-turbo (see ICTable 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

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

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.

3.3 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. $C_w^{\rm total}$ denotes the total number of queries answered correctly by the Weak LLM. Let ${\rm Tok}_J$ and ${\rm Tok}_O$ be the tokens consumed by Jung's method and our proposed Inter-Cascade pipeline, and ${\rm Cost}_J$ and ${\rm Cost}_O$ denote their corresponding costs. The evaluation metrics are summarized in Table 2.

Table 2: Evaluation Metrics

Metric	Formula	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	$(\operatorname{Tok}_J - \operatorname{Tok}_O)/\operatorname{Tok}_J$
Cost Reduction	$(\operatorname{Cost}_J - \operatorname{Cost}_O)/\operatorname{Cost}_J$		

3.4 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%-5.53% increase in Pipeline Accuracy on reasoning benchmarks and a 0.76% increase on the non-reasoning factual NASA-History benchmark. The Strong LLM Call Rate is reduced on all benchmarks, with reductions ranging from 4.41% to 28.53%. These results indicate that *Inter-Cascade* pipeline is beneficial across different categories of tasks and particularly effective for reasoning-intensive tasks.

Effectiveness of Similarity-Based Retrieval. To isolate the effect of strategy selection, we include a control variant in which *Inter-Cascade* selects strategies at random. Across datasets, its performance generally falls between the *Inter-Cascade* and Jung's pipeline (see Table 3), demonstrating the benefit of similarity-based retrieval. Although outside the scope of this work, one possible future direction is to further refine the selection of strategies, which would involve verifying whether all of the top-k retrieved strategies are relevant to the given queries. The accuracy of the *Inter-Cascade* (random strategies) differs from Jung's by -2.43% to +1.93%, and its Strong LLM Call Rate shows only a modest reduction, ranging from 1.59% to 5.17%.

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 two pipelines across all benchmarks. The improvements are particularly pronounced on reasoning benchmarks, with gains of 14.88%, 16.2%, and 33.06% on MetaMath, GSM-Plus, and GSM-Symbolic, respectively, while still achieving an improvement of 0.48% on the non-reasoning NASA-History benchmark. 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 that Inter-Cascade enhances not only the Weak LLM's accuracy but also its confidence on correct predictions, a crucial factor in converting local improvements into overall pipeline gains.

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 H. 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

Table 3: Results across datasets using different pipelines. "Jung" denotes Jung's LLM-Cascade, "Our (Random)" denotes Inter-Cascade with random strategies, 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, $\alpha_w=0.6, \delta_w=0.65$. (b) GSM-Plus: For the Strong LLM, $\alpha_s=0.2, \delta_s=0.8, \delta_s=0.8, \delta_s=0.8$. For the Weak LLM, $\alpha_w=0.6, \delta_w=0.6, \delta_w=0.48$. (c) MetaMath: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w=0.4, \delta_w=0.6, \delta_w=0.61$. (d) NASA-History: No threshold is applied for the Strong LLM. For the Weak LLM, $\alpha_w=0.2, \delta_w=0.7, \delta_w=0.87$.

Data Pipeline		Pipeline Acc. (%)↑	Strong Call (%)↓	Cov. (%)	
	Jung	66.04	59.37	86.31	
GSM-Symb.	Our (Random)	63.61	54.20	87.90	
	Our (Retrieval)	70.37	30.84	90.35	
GSM-Plus	Jung	52.78	46.29	93.57	
	Our (Random)	53.63	43.64	94.10	
	Our (Retrieval)	58.31	32.44	94.79	
MetaMath	Jung	65.92	49.22	100.00	
	Our (Random)	67.85	45.99	100.00	
	Our (Retrieval)	70.76	34.98	100.00	
NASA-Hist.	Jung	71.88	26.68	100.00	
	Our (Random)	71.32	25.09	100.00	
	Our (Retrieval)	72.64	22.54	100.00	

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

Data	Pipeline	Weak Acc. (%)↑	Weak Corr. Accpt. (%)↑
	Jung	15.04	12.34
GSM-Symb.	Our (Random)	17.40	15.27
	Our (Retrieval)	48.10	46.09
GSM-Plus	Jung	22.46	19.13
	Our (Random)	25.51	22.38
	Our (Retrieval)	38.66	35.73
MetaMath	Jung	35.43	28.52
	Our (Random)	38.64	32.66
	Our (Retrieval)	50.31	45.31
NASA-Hist.	Jung	66.22	55.37
	Our (Random)	65.22	55.56
	Our (Retrieval)	66.70	58.40

Table 5: Token and API cost reductions across datasets for Inter-Cascade compared with Jung's pipeline.

Dataset	Token Reduction	Cost Reduction	Dataset	Token Reduction	Cost Reduction
GSM-Symb.	47.80%	49.63%	MetaMath	31.66%	33.44%
GSM-Plus	29.98%	30.41%	NASA-Hist.	15.47%	15.75%

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.

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. Table 5 shows the percentage reductions in token usage and corresponding API costs compared with Jung's pipeline. The main overhead of our approach lies in maintaining the Repo. A simple solution is to periodically prune it, as query distributions may shift. More promisingly, accumulated queries and responses

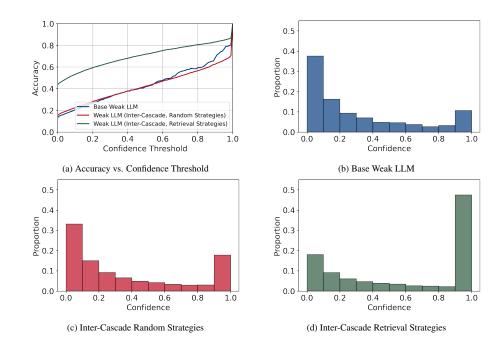


Figure 2: GSM-Symbolic dataset: (a) Accuracy as a function of the confidence threshold for the base Weak LLM, Inter-Cascade with random strategies, and Inter-Cascade with retrieval strategies, and (b) - (d) their corresponding confidence histograms. Our Inter-Cascade (Retrieval) consistently concentrates 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).

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.

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.

4 Conclusion

We propose *Inter-Cascade*, an online and interactive *LLM Cascade* framework that enables Weak LLM to learn online from Strong LLM's prior experience without fine-tuning. Inter-Cascade improves both the accuracy of Weak LLM and the overall system, while reducing the reliance on Strong LLM, saving computation, monetary cost, and latency (when Strong LLM is deployed on remote server), compared with current LLM Cascade.

Inter-Cascade provides a general and scalable framework for multi-LLM systems, which can be implemented with different LLMs and cascade layers. Despite the promising performance of Inter-Cascade, further improvements can still be achieved by proposing better methods of generating strategy, better algorithms to evaluate similarity and mechanisms to prevent mismatch in future work. Inter-Cascade is also naturally well-suited for distributed systems, where local Weak LLM owners can teach and boost their model in a tailored way by sending customized queries to Strong LLM. Another future work direction arises from Inter-Cascade's potential to bridge online and offline learning. While augmenting the system performance by incorporating queries with related strategies during online operation, the generated strategy repository can be exported as local training set for periodic finetuning, permanently improving the capability of Weak LLM. We hope Inter-Cascade inspires future research on building more interactive LLM Cascades or other multi-LLM systems.

REFERENCES

- Peter Bauer. Multiple testing in clinical trials. Statistics in medicine, 10(6):871–890, 1991.
- 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*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=cSimKw5p6R.
- 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. URL https://doi.org/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. Technical Report 2110.14168, arXiv, November 2021. URL https://doi.org/10.48550/arXiv.2110.14168.
- 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. URL https://doi.org/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. URL https://aclanthology.org/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. Technical Report 2401.08281, arXiv, 2025. URL https://arxiv.org/abs/2401.08281.
- 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*, 2023.
- 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. Technical report, arXiv, February 2025. URL https://doi.org/10.48550/arXiv.2404.16130.
- 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. URL https://huggingface.co/datasets/patrickfleith/NASA-History-MCQ.
- Naresh Kumar Gundla and Sri Harsha Atthuluri. Multi-node inference architectures for low-latency LLM serving. In 2025 International Conference on Advanced Computing Technologies (ICoACT), pp. 01–06. IEEE, 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. Technical Report 1503.02531, arXiv, March 2015. URL https://doi.org/10.48550/arXiv.1503.02531.
 - Geoffrey Irving, Paul Christiano, and Dario Amodei. AI safety via debate. Technical Report 1805.00899, arXiv, October 2018. URL https://doi.org/10.48550/arXiv.1805.00899.
 - 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. URL https://aclanthology.org/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. URL https://doi.org/10.1109/TBDATA.2019.2921572.
 - 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. URL https://doi.org/10.1109/TNNLS.2022.3223018.
 - Raul Junco. Every repeated LLM call is money on fire, June 2025.

 URL https://newsletter.systemdesignclassroom.com/p/
 every-repeated-llm-call-is-money-on-fire.
 - 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. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/08dabd5345b37fffcbe335bd578b15a0-Paper-Conference.pdf.
 - 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. Technical Report 2001.08361, arXiv, 2020. URL https://arxiv.org/abs/2001.08361.
 - 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.
 - Aisha Khatun and Daniel G. Brown. Trutheval: A dataset to evaluate llm truthfulness and reliability. Technical Report 2406.01855, arXiv, 2024. URL https://arxiv.org/abs/2406.01855.
 - **Dmitry** Korzhov. Stop calling out your LMM for every sim-2023. URL ilar question!, July https://uxplanet.org/ stop-calling-out-your-lmm-for-every-similar-question-50940c64b3df.
 - 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. URL https://doi.org/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. URL https://aclanthology.org/2024.acl-long.163/.
 - 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. URL https://aclanthology.org/2024.lrec-main.593/.
 - 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. URL https://doi.org/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. URL https://aclanthology.org/2023.findings-emnlp.334/.
- 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. URL https://openreview.net/forum?id=AjXkRZIvjB.
- 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. URL https://proceedings.mlr.press/v119/mozannar20b.html.
- 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. URL https://openreview.net/forum?id=vo9t20wsmd.
- 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*, 2025. URL https://openreview.net/forum?id=zl0HLZOJC9.

- 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. URL https://openreview.net/forum?id=8sSqNntaMr.
- 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. URL https://aclanthology.org/2023.emnlp-main.180/.
- Archit Parnami and Minwoo Lee. Learning from few examples: A summary of approaches to few-shot learning, 2022. URL https://arxiv.org/abs/2203.04291.
- 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. Technical Report 2502.01459, arXiv, February 2025. URL https://doi.org/10.48550/arXiv.2502.01459.
- 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. URL https://aclanthology.org/D19-1410/.
- Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints for Thin Deep Nets. Technical Report 1412.6550, arXiv, March 2015. URL https://doi.org/10.48550/arXiv.1412.6550.
- 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. URL https://aclanthology.org/2022.naacl-main.191/.
- Tobias Schnabel. Stop Wasting LLM Tokens, August 2024. URL https://towardsdatascience.com/stop-wasting-llm-tokens-a5b581fb3e6e/.
- 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. URL https://aclanthology.org/2024.acl-long.701/.
- 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. URL https://doi.org/10.1609/aaai.v39i27.35063.
- Joshua Strong, Pramit Saha, Yasin Ibrahim, Cheng Ouyang, and Alison Noble. Expert-agnostic learning to defer, 2025b. URL https://arxiv.org/abs/2502.10533.
- 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 (GEM*²), pp. 404–430, July 2025. URL https://aclanthology.org/2025.gem-1.33/.
- 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*, 2025. URL https://openreview.net/forum?id=CI4sCBMXjP.
- 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. URL https://openreview.net/forum?id=sJvhwDtFhQ.
- 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. URL https://aclanthology.org/2024.acl-long.511/.
- Yu Wu and Anand Sarwate. Learning to help: Training models to assist legacy devices. Technical Report 2409.16253, arXiv, 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. URL https://openreview.net/forum?id=NCgTbt2j1F.
- Yu Xia, Fang Kong, Tong Yu, Liya Guo, Ryan A. Rossi, Sungchul Kim, and Shuai Li. Which Ilm 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. URL https://doi.org/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. URL https://openreview.net/forum?id=gjeQKFxFpZ.
- Xiao Yan and Yi Ding. Are we there yet? A measurement study of efficiency for LLM applications on mobile devices. In *Proceedings of the 2nd International Workshop on Foundation Models for Cyber-Physical Systems & Internet of Things*, pp. 19–24, 2025.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. Technical Report 2505.09388, arXiv, 2025.

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. URL https://openreview.net/forum?id=N8N0hqNDRt.

- Michael J. Zellinger and Matt Thomson. Rational tuning of LLM cascades via probabilistic modeling. *Transactions on Machine Learning Research*, 2025. URL https://openreview.net/forum?id=YCBVcGSZeR.
- Michael J. Zellinger, Rex Liu, and Matt Thomson. Cost-Saving LLM Cascades with Early Abstention. Technical Report 2502.09054, arXiv, February 2025. URL https://doi.org/10.48550/arXiv.2502.09054.
- Hao Zhou, Chengming Hu, Ye Yuan, Yufei Cui, Yili Jin, Can Chen, Haolun Wu, Dun Yuan, Li Jiang, Di Wu, et al. Large language model (llm) for telecommunications: A comprehensive survey on principles, key techniques, and opportunities. *IEEE Communications Surveys & Tutorials*, 2024.
- 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. URL https://aclanthology.org/2025.findings-acl.475/.

A CLARIFICATION: USE OF LLMS ON AIDING OR POLISHING WRITING

We used ChatGPT and Gemini solely as writing assistants to correct the typos and grammars, help polish the language, improve clarity, and refine the presentation of this manuscript. The LLMs did not contribute to the conception of ideas, design of experiments, execution of analyses, or interpretation of results.

B EXTENDED RELATED WORK

LLM Cascade LLM Cascade is firstly proposed by Chen et al. (2024) 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., 2025; 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.

865

866

867

868

870

871

872

873

874

875

876

877

878

879 880

882

883

884

885

887

888

889

890

891

892

893

894 895

897

898

899

900

901

902

903

904 905 906

907908909

910

911

912

913

914

915

916 917

Other related topics Retrieval-Augmented Generation (RAG) (Lewis et al., 2020; Edge et al., 2025; Wang et al., 2025) is an approach that combines pre-trained parametric and non-parametric memory for language generation. There are also studies work on training a retriever to pick up context for queries (Rubin et al., 2022; Margatina et al., 2023). RAG and related works focus on building accurate retrieval algorithms for existing database while in Inter-Cascade, we dynamically update the strategy repository during the inference phase and the strategies are generated by Strong LLM according to the users history query-style preference. There are also a weak model and strong model in Speculative 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., 2023; 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, \dots, q_I\} \subseteq \mathcal{Q}; LLM M_n with deferral function d_n, generation function
     g_n, strategy repository \text{Repo}_n and strategy generator h_n.
     Deferral convention: 0 = \text{handle locally}, 1 = \text{defer/forward}.
 1: 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)
                   [s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}] \leftarrow f_n(q_i, \text{Repo}_n)
q_i' \leftarrow [q_i, s_i^{t_1}, s_i^{t_2}, ..., s_i^{t_k}]
 6:
                                                                     \triangleright Find most relevant top-k strategies to q_i
 7:
                                                                                8:

    ▶ Last LLM doesn't maintain Repo

 9:
              (Deferral Decision)
10:
              if d_n(q_i') = 0 then
11:
                   generate answer a_i \leftarrow g_1(q_i')
12:
                                                                                 s_{\text{new}} \leftarrow h(q_i)
13:
                   \operatorname{Repo}_{\leq n} \leftarrow \operatorname{Repo}_{\leq n} \cup \{s_{\text{new}}\}\
14:

    ▷ Add strategy to all the weaker LLMs

              else
15:
                   if n < N then
16:
17:
                       Pass
                                                                                                  ▷ Defer to next level
18:
                   else
19:
                       Discard current query q_i
                                                            ▶ None of LLMs are confident to answer the query
```

C 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 $\operatorname{Perf}(M)$, which can be instantiated by measures such as the expected accuracy or negative loss on \mathcal{D} . Similarly, we let $\operatorname{Cost}(M)$ 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 $\operatorname{Cost}(M)$. We say that M_w is weaker than M_s if $\operatorname{Perf}(m_w) \leq \operatorname{Perf}(m_s)$, and that it is cheaper if $\operatorname{Cost}(m_w) \leq \operatorname{Cost}(m_s)$. To simplify notation, we introduce the shorthand relation

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

if and only if

$$\operatorname{Perf}(M_w) \leq \operatorname{Perf}(M_s)$$
 and $\operatorname{Cost}(M_w) \leq \operatorname{Cost}(M_s)$,

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 $M_1 \leq_{\text{wbc}} M_2 \leq_{\text{wbc}} ... \leq_{\text{wbc}} M_N$.

D 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.

E PROOF: CLOPPER-PERSON UPPER BOUND AS A BETA QUANTILE

In the lemma below, we apply the Clopper-Pearson upper bound to rewrite $R^+(\lambda)$, yielding a clearer form that facilitates computation. This helps the proof of Theorem 2.2 and Theorem G.1.

Lemma E.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

$$X \sim \text{Bin}(n(\lambda), R(\lambda)),$$

and $x \in \{0, 1, ..., n(\lambda)\}$ is the number of error observed. Write $\widehat{R}(\lambda) = x/n(\lambda)$. For a fixed $\delta \in (0, 1)$, define the one-sided $(1 - \delta)$ upper confidence limit by

$$\widehat{R}^+(\lambda) := \sup \left\{ p \in [0,1] : \Pr_p(\operatorname{Bin}(n(\lambda), p) \le x) \ge \delta \right\}.$$

Then

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

with the usual edge conventions $\operatorname{Beta}^{-1}(1-\delta;1,n)=1-\delta^{1/n}$ when x=0 and $\widehat{R}^+(\lambda)=1$ when $x=n(\lambda)$.

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

$$F(p^*) = P(\operatorname{Bin}(n(\lambda), p^*) \le x) = \delta. \tag{3}$$

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

$$P(X \le x) = \sum_{k=0}^{x} {n(\lambda) \choose k} p^k (1-p)^{n(\lambda)-k} = 1 - I_p(x+1, n(\lambda) - x),$$

where $I_p(a,b)$ is the CDF of Beta(a,b) at p. Plugging this into equation 3 gives

$$I_{p^*}(x+1, n(\lambda) - x) = 1 - \delta,$$

so p^* is the $(1-\delta)$ quantile of the Beta $(x+1, n(\lambda)-x)$ distribution:

$$p^* = \text{Beta}^{-1}(1 - \delta; x + 1, n(\lambda) - x).$$

This equals $\widehat{R}^+(\lambda)$ by definition. The stated edge cases follow from $F(p)=(1-p)^{n(\lambda)}$ when x=0 and from monotonicity when $x=n(\lambda)$.

F PROOF: UNCHANGED THRESHOLD

Theorem F.1. Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ . Fix $\delta \in (0,1)$ and an integer $n \geq 1$. For $x \in \{0,1,\ldots,n\}$, $\epsilon \in (0,1]$, and $b \in [1,\infty)$. Suppose that $\min\{\epsilon x+1,\,n-\epsilon x\}$ is moderately large and $1-\delta$ is not an extreme tail, 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 Normal cumulative distribution function, when n is large enough, the decrease of the risk under same level of tolerance is given by,

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

Proof. We use a Beta function 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 approximation to Beta function to evaluate the decrease of α by definition. For the convenience of statement of our theories, we define that $\alpha(\epsilon,b)$ as the 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 E.1, we know that

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

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

$$P(\operatorname{Bin}(bn, p_1) \le \epsilon x) \le P(\operatorname{Bin}(n, p_1) \le x) = \delta,$$

because lowering the threshold ($\epsilon x \leq bx$) 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 E.1, it is equivalent to that $P(\mathrm{Bin}(bn,p_2) \leq \epsilon x) = \delta$. It follows that $P(\mathrm{Bin}(bn,p_2) \leq \epsilon x) = \delta \geq P(\mathrm{Bin}(bn,p_1) \leq \epsilon x)$, which implies that $p_2 \leq p_1$. Hence the new upper bound $p_2 = \mathrm{Beta}^{-1}(1-\delta; x'+1, n'-x')$ satisfies $p_2 \leq p_1$. This shows the statement (a).

(b) Write

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

In the large–sample, interior regime, e.g., $\min\{\epsilon x+1, n-\epsilon x\}\gg 1$ and x/n bounded away from 0 and 1,

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

This is by the approximation to Beta distribution by normal distribution. Calculate $\alpha(1,1) - \alpha(\epsilon,b)$ demonstrate the result of theorem.

G 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

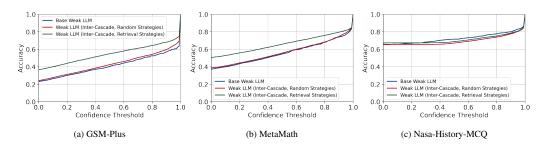


Figure 3: Accuracy as a function of the confidence threshold for the base Weak LLM and for the Weak LLM within the Inter-Cascade using random and retrieval strategies across three benchmarks.

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 G.1.

Theorem G.1. Suppose that $\widehat{R}^+(\lambda)$ is a monotonic decreasing function of λ . Fix $\delta \in (0,1)$ and an integer $n \geq 1$. For $x \in \{0,1,\ldots,n\}$ and $\epsilon \in (0,1]$. Suppose that $\min\{\epsilon x+1, n-\epsilon x\}$ is moderately large and $1-\delta$ is not an extreme tail, then:

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

$$\alpha(\epsilon) < \alpha(1)$$
.

(b) Normal approximation for the amount of decrease. Let $z := \Phi^{-1}(1 - \delta)$, for ϵ near 1,

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

Hence the decrease is approximately linear in $(1 - \epsilon)$ with the coefficient in brackets; in particular, when $x \le n/2$ 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 2.2, 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 2.2, notice that

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

For $i=\epsilon x+1$, $j=n-\epsilon x$, the Beta (i,j) mean and variance are $\mu_{\epsilon}=i/(i+j)$ and $\sigma_{\epsilon}^2=ij/[(i+j)^2(i+j+1)]$. Approximating the $(1-\delta)$ quantile by the Normal formula gives $\alpha(\epsilon)=\mu_{\epsilon}+z\sigma_{\epsilon}+O(1/n)$. Differentiate at $\epsilon=1$ to obtain the first-order change:

$$\frac{d\mu_\epsilon}{d\epsilon}\Big|_{\epsilon=1} = \frac{x}{n+1}, \qquad \frac{d\sigma_\epsilon}{d\epsilon}\Big|_{\epsilon=1} = \frac{1}{2(n+1)\sqrt{n+2}} \cdot \frac{(n-1-2x)x}{\sqrt{(x+1)(n-x)}}.$$

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

H CONFIDENCE DISTRIBUTION

Figures 3 and 4 present results for the GSM-Plus, MetaMath, and Nasa-History-MCQ datasets, complementing the GSM-Symbolic analyses in the main text.

Figure 3 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

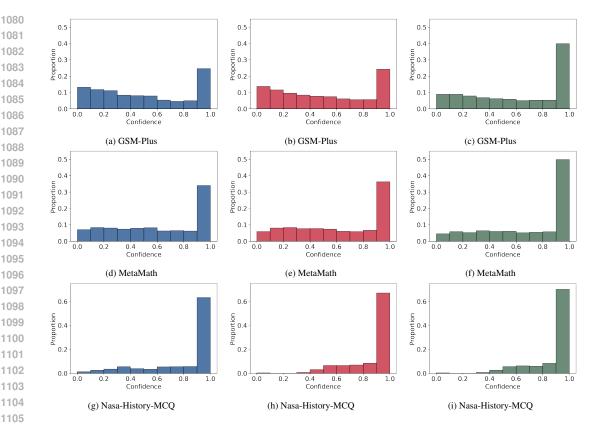


Figure 4: Confidence histograms for three benchmarks. Columns correspond to (a)(d)(g) the base Weak LLM, (b)(e)(h) the Weak LLM within the Inter-Cascade using random strategies, and (c)(f)(i) the Weak LLM within the Inter-Cascade using retrieval strategies. Across all datasets, the Inter-Cascade with 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.

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 4 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.

I 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 J.

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 inser-

tion, 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 J.

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 10,000 additional problems as the test set. The prompt template and an example problem are provided in Appendix J.

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 J.

J PROMPT TEMPLATES AND EXAMPLES

Table 6 and Table 7 present the strategy-free prompt templates for the four datasets, along with one example question per dataset. Table 8 - Table 11 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 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 6: Strategy-free prompt template with example questions from GSM-Symbolic, GSM-Plus, and Meta-**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:** many minutes will it take for the fog bank to cover the whole city? **GSM-Plus Example Question:** there than white shirts with no collars?

[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

[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

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?

	Table 7: Strategy-free prompt template with example question from NASA-History-MCQ
_	
Pı	rompt Template:
Ва	ased on the question below, please strictly follow this format when answering:
	Start with [Strategy] section explaining the general approach for solving similar problems;
	End with [Answer] section containing ONLY the single capital letter of the correct option (exactly
OI	ne of: A, B, C, D). Do NOT include the option text or any explanation.
E	xample:
[Ç	Question]: Which of the following was a direct result of the Treaty of Versailles (1919)?
	. The outbreak of World War II
	The United States declaring war on Germany and entering World War I. The establishment of the Fascist regime in Italy
	The creation of the League of Nations
[S	trategy]: Recall that the treaty ended World War I and included peacekeeping measures. Compare
	otions with the treaty: only the League of Nations was directly established by it.
ĮΑ	answer]: D
N	ow answer this question:
	Question]: {question}
	trategy]:
ĮΑ	answer]:
_	
N	asa-History-MCQ Example Question:
[Ç	Question]: What three distinct stages comprise current psychological support protocols for US astro-
na	auts?
	. Individual, group, family
	Training, mission, debriefing Cognitive, emotional, behavioral
	. Preflight, in-flight, postflight

1296 1297 1298 1299 1300 1301 Table 8: Strategy-based prompt template with example input from GSM-Symbolic 1302 1303 **Prompt Template:** 1304 Using the following problem examples: 1305 {strategy} 1306 Based on the question below, please strictly follow this format when answering: 1307 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, 1309 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1310 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1311 Now answer this question: 1312 [Question]: {question} 1313 [Strategy]: [Answer]: 1315 1316 1317 **GSM-Symbolic Example Input:** 1318 Using the following problem examples: 1319 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 1320 many minutes will it take for the fog bank to cover the whole city? 1321 [Strategy]: First, determine how many 2-mile segments are in the 24-mile city. Divide the total 1322 distance of the city by the length of each segment. Then, multiply the number of segments by the time 1323 it takes to cover each segment to find the total time. 1324 [Answer]: 1116 1325 2. [Question]: A fog bank rolls in from the ocean to cover a city. It takes 114 minutes to 1326 cover every 3 miles of the city. If the city is 99 miles across from the oceanfront to the opposite 1327 inland edge, how many minutes will it take for the fog bank to cover the whole city? 1328 [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 1330 miles = 99/38 minutes. Simplify the fraction 99/38 = 2.60526315789. Multiply 114 by 99/3 to 1331 get the answer 114 * (99/3) = 114 * 33 = 3762. 1332 [Answer]: 3762 1333 1334 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; 1335 2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, 1336 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1337 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1338 1339 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 1340 miles of the city. If the city is 95 miles across from the oceanfront to the opposite inland edge, how 1341 many minutes will it take for the fog bank to cover the whole city? [Strategy]: [Answer]: 1344

1350 1351 1352 1353 1354 Table 9: Strategy-based prompt template with example input from GSM-Plus 1355 1356 **Prompt Template:** 1357 Using the following problem examples: 1358 {strategy} 1359 1360 Based on the question below, please strictly follow this format when answering: 1361 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, 1362 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1363 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1364 1365 Now answer this question: [Question]: {question} [Strategy]: 1367 [Answer]: 1368 1369 1370 **GSM-Plus Example Input:** 1371 Using the following problem examples: 1372 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 1373 there than white shirts with no collars? 1374 [Strategy]: Let $\,W\,$ be the number of white shirts. The number of white shirts with collars is $\,W/2\,$, 1375 so the number of white shirts without collars is W - W/2 = W/2. The number of floral shirts is 1376 50. The number of floral shirts with buttons is 20, so the number of floral shirts without buttons is 1377 50-20=30. We want to find the difference between the number of floral shirts without buttons 1378 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 1379 difference. 1380 [Answer]: None 1381 1382 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? 1384 [Strategy]: 1. Calculate the number of white shirts with collars: 40/2 = 20. 2. Calculate the number 1385 of white shirts without collars: 40 - 20 = 20. 3. Calculate the number of floral shirts without 1386 buttons: 20 + 10 = 30. 4. Calculate the number of floral shirts with buttons: 50 - 30 = 20. 1387 [Answer]: 20 1388 Based on the question below, please strictly follow this format when answering: 1389 1. Start with [Strategy] section explaining the general approach for solving similar problems; 1390 2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, 1391 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1392 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1393 Now answer this question: 1394 [Question]: A clothing store has 60 white shirts and 70 floral shirts. Two-thirds of the white shirts 1395 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]: 1398

1404 1405 1406 1407 1408 Table 10: Strategy-based prompt template with example input from MetaMath 1409 1410 **Prompt Template:** 1411 Using the following problem examples: 1412 {strategy} 1413 1414 Based on the question below, please strictly follow this format when answering: 1415 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, 1416 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1417 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1418 1419 Now answer this question: [Question]: {question} 1420 [Strategy]: 1421 [Answer]: 1422 1423 1424 MetaMath Example Input: 1425 Using the following problem examples: 1426 1. [Question]: Martha has x crayons. She lost half of them, so she bought a new set of 20 crayons. 1427 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 1428 x/2 crayons left. Then she bought 20 new crayons, so she had x/2 + 20 crayons. We are given that 1429 she has 29 crayons in total, so x/2 + 20 = 29. Subtracting 20 from both sides gives x/2 = 9. 1430 Multiplying both sides by 2 gives x = 18. 1431 [Answer]: 18 1432 2. [Question]: Justin and Sabrina at Pine Grove Elementary have 50 pencils combined. Justin 1433 has x more than twice as many pencils as Sabrina. Sabrina has 14 pencils. What is the value of 1434 unknown variable x? 1435 [Strategy]: Let J be the number of pencils Justin has and S be the number of pencils Sabrina has. 1436 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 1437 the first equation to find J: J+14=50, so J=50-14=36. Now we can substitute J=361438 and S = 14 into the second equation: 36 = 2(14) + x, so 36 = 28 + x. Solving for x, we get 1439 x = 36 - 28 = 8. 1440 [Answer]: 8 1441 1442 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; 1443 2. End with [Answer] section containing ONLY the value. (1) Do NOT include units such as minutes, 1444 feet, etc.; (2) If the question asks for a percentage, ONLY provide the number (e.g., answer 20 instead 1445 of 20%); (3) Do NOT include any explanations; (4) If there is no answer, RETURN None as the value. 1446 1447 Now answer this question: [Question]: Lara bought 52 stems of flowers at the flower shop. She gave 15 flowers to her mom 1448 and gave x more flowers than she gave to her mom to her grandma. She put the rest in a vase. How 1449 many stems of flowers did Lara put in the vase? If we know the answer to the above question is 16, 1450 what is the value of unknown variable x? 1451 [Strategy]: 1452 [Answer]:

1458 Table 11: Strategy-based prompt template with example input from NASA-History-MCQ 1459 1460 **Prompt Template:** 1461 Using the following problem examples: 1462 {strategy} 1463 Based on the question below, please strictly follow this format when answering: 1464 1. Start with [Strategy] section explaining the general approach for solving similar problems; 1465 2. End with [Answer] section containing ONLY the single capital letter of the correct option (exactly 1466 one of: A, B, C, D). Do NOT include the option text or any explanation. 1467 Now answer this question: 1468 [Question]: {question} 1469 [Strategy]: 1470 [Answer]: 1471 1472 **Nasa-History-MCQ Example Input:** 1474 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? 1476 A. Specialized dietary supplements to combat isolation 1477 B. Automated exercise routines tailored to reduce stress 1478 C. Psychological support hardware and software 1479 D. Advanced life support systems with mood stabilizers [Strategy]: The question asks about specific hardware or software that aids psychological well-being 1480 during long-duration spaceflights, beyond communication and care packages. We need to evaluate 1481 each option to see if it fits this description. Option A focuses on dietary supplements, which are not 1482 hardware or software. Option B describes automated exercise routines, which could involve software 1483 and hardware. Option C is too general, simply restating the question. Option D focuses on life support 1484 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. 1485 [Answer]: B 1486 1487 2. [Question]: What is the anticipated effect of constraints inherent in lunar and Martian mis-1488 sions on psychological support approaches? 1489 A. Greater emphasis on real-time communication with Earth-based support teams B. Increased reliance on virtual reality and AI companionship to mitigate isolation 1490 C. A shift towards highly individualized psychological profiles and tailored interventions 1491 D. A return to the mindset and strategies of earlier explorers and their families 1492 [Strategy]: The question asks about the impact of constraints in lunar and Martian missions on psy-1493 chological support. These constraints include isolation, limited resources, communication delays, and 1494 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 1495 and a need for personalized approaches due to the unique challenges faced by each astronaut. Options 1496 A and D are less likely because of communication delays and the differences between modern space 1497 missions and earlier explorations. Option B is plausible, but option C is more comprehensive as it 1498 addresses the need for personalized support, which is crucial given the constraints. 1499 [Answer]: C 1500 Based on the question below, please strictly follow this format when answering: 1501 1. Start with [Strategy] section explaining the general approach for solving similar problems; 1502 2. End with [Answer] section containing ONLY the single capital letter of the correct option (exactly 1503 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 1506 astronauts? 1507 A. Individual, group, family B. Training, mission, debriefing 1509 C. Cognitive, emotional, behavioral D. Preflight, in-flight, postflight 1510 [Strategy]: 1511

[Answer]: