
C3PO: Optimized Large Language Model Cascades with Probabilistic Cost Constraints for Reasoning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Large language models (LLMs) have achieved impressive results on complex
2 reasoning tasks, but their high inference cost remains a major barrier to real-world
3 deployment. A promising solution is to use cascaded inference, where small,
4 cheap models handle easy queries, and only the hardest examples are escalated
5 to more powerful models. However, existing cascade methods typically rely on
6 supervised training with labeled data, offer no theoretical generalization guarantees,
7 and provide limited control over test-time computational cost. We introduce **C3PO**
8 (*Cost Controlled Cascaded Prediction Optimization*), a self-supervised framework
9 for optimizing LLM cascades under probabilistic cost constraints. By focusing on
10 minimizing regret with respect to the most powerful model (MPM), C3PO avoids
11 the need for labeled data by constructing a cascade using only unlabeled model
12 outputs. It leverages conformal prediction to bound the probability that inference
13 cost exceeds a user-specified budget. We provide theoretical guarantees on both
14 cost control and generalization error, and show that our optimization procedure is
15 effective even with small calibration sets. Empirically, C3PO achieves state-of-the-
16 art performance across a diverse set of reasoning benchmarks—including GSM8K,
17 MATH-500, and BigBench-Hard—outperforming strong LLM cascading baselines
18 in both accuracy and cost-efficiency. Our results demonstrate that principled,
19 label-free cascade optimization can enable scalable and reliable LLM deployment.

20 1 Introduction

21 Large language models (LLMs) have demonstrated remarkable capabilities in a wide range of
22 reasoning tasks, from arithmetic problem solving to common sense and formal logical reasoning.
23 Techniques such as chain-of-thought prompting have further unlocked multi-step reasoning by
24 encouraging LLMs to provide explicit intermediate rationale steps before arriving at an answer [Wei
25 et al., 2022a]. However, these gains come at a steep cost as each additional token or sample incurs
26 inference time and monetary expenditure, imposing a barrier for real-world deployment of LLMs.

27 A natural remedy is to employ a *cascade* of models of increasing size and accuracy. This way, cheap
28 and small models handle easy queries, deferring only the hard cases to a powerful LLM. From early
29 works in classifier cascades [Viola and Jones, 2001] to more recent efforts in LLM cascading, it has
30 been conclusively shown that cascades can yield dramatic cost savings with minimal accuracy loss.
31 Yet existing LLM cascade approaches share two major limitations. First, they rely on large labeled
32 datasets for meta-model training under budget constraints, which is expensive to collect and slow to
33 adapt to new domains [Zhang et al., 2024]. Second, they offer few theoretical guarantees on either
34 generalization error or worst-case cost overruns, leaving practitioners to tune budgets and thresholds
35 by hand [Yue et al., 2024]. While model weight tuning is not a major hindrance in general cascade
36 settings, the high training costs of LLMs make finetuning costly. Furthermore, the need for test time

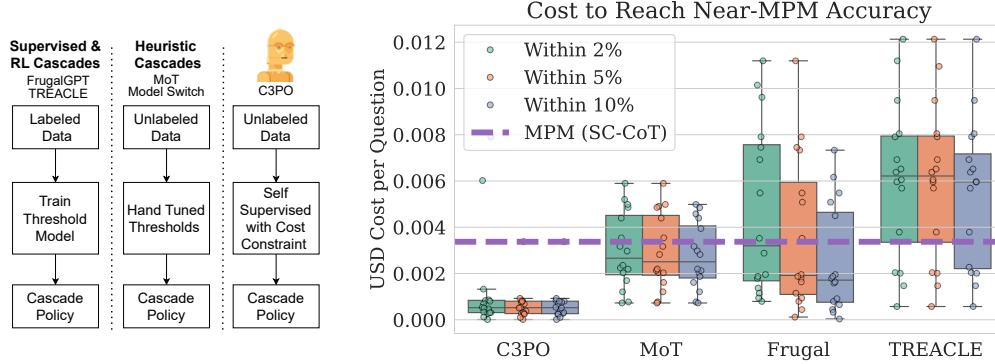


Figure 1: Left: Existing LLM cascading strategies include supervised learning, reinforcement learning, and learning-free heuristics. In contrast, C3PO proposes a novel, cost-constrained, self supervised (label-free) cascade paradigm. Right: Surpassing existing cascade approaches tremendously, C3PO offers markedly superior cost-effectiveness across 16 benchmarks, requiring **less than 20% of the cost of the most powerful model (MPM)** (cost shown in purple) for an accuracy gap of at most 2, 5, and 10% using a LLAMA cascade. In this boxplot, each dot represents a dataset and the whiskers extend to 90% coverage.

adaptation using few examples, motivates development of specialized LLM cascade methods that do not require ground truth labels.

In this work, we introduce *Cost Controlled Cascaded Prediction Optimization (C3PO)*, a self-supervised cascaded inference framework that overcomes these challenges. Rather than depending on labels or costly meta-training, C3PO learns to exit early by exploiting *consistency signals* between a cheap model and the *most powerful model (MPM)* in the cascade. Intuitively, when a small model’s confidence is high, we can safely halt inference—saving MPM’s inference cost—while maintaining its accuracy. In such cases, it is highly likely that both the cheap model and the MPM would be correct. Additionally, when both models are wrong and we still exit early, we accept a controlled loss in accuracy in exchange for significant cost reduction. Only ambiguous cases where the cheap model is expected to disagree with MPM should be forwarded to the next model in the cascade. C3PO makes three key technical contributions:

- **Data-efficient cascade learning.** We show that cascade decision rules can be learned using only a small “self-supervision” pool of unlabeled prompts—fewer than 1% of the examples used by *state-of-the-art* methods such as TREACLE [Zhang et al., 2024]. This self-supervised approach requires no ground-truth labels, enabling easy adaptation to new domains and task distributions.
- **Rigorous theoretical guarantees.** We derive conformal cost bounds that with arbitrary probability control the cascade’s inference cost under a user-specified budget. In parallel, we establish PAC-Bayesian generalization bounds certifying that the learned cascade decision rules generalize.
- **State-of-the-art empirical performance.** C3PO consistently outperforms across a diverse suite of reasoning benchmarks including arithmetic (GSM8K, SVAMP), formal mathematical (MATH-500), and logical reasoning tasks (CommonSenseQA, BIG-Bench-Hard).

2 Related Work

LLM Cascading is a structured inference paradigm in which language models are queried in a sequential manner based on their complexity and cost. The goal is to minimize inference costs while maintaining high answer quality. FrugalGPT [Chen et al., 2024] first demonstrated the efficacy of cascading in reducing LLM inference costs by predicting the likelihood that an LLM output is correct using a BERT-like meta-model. If this prediction exceeds a user-defined threshold, the model’s output is returned as the final answer; otherwise, the next, more capable model is invoked. Zhang et al. [2023] extend this idea by incorporating LLM-based self-verification before escalating to a more powerful model, though this introduces additional delays due to sequential dependencies. AutoMix [Aggarwal et al., 2024] leverages self-consistency sampling to decide whether to defer

to stronger LLMs, learning thresholds over sampled predictions to optimize for quality and cost. Similarly, Mixture-of-Thoughts (MoT) [Yue et al., 2024] integrates external programmatic solvers and majority voting to assess the adequacy of weak model responses before escalation. More recent approaches such as ModelSwitch [Chen et al., 2025a] employ self-consistency as a decision criterion to escalate queries. If samples from the current LLM are inconsistent, a larger model is queried, and a final decision is made via voting over all collected samples. Our method differs crucially from other self consistency based approaches by learning to map self consistency scores to agreement probabilities with the MPM rather than with the ground truth. This eliminates the need for dataset labels and simplifies the cascade learning greatly as we do not need many samples to understand LLM output correlations. TREACLE [Zhang et al., 2024] and Online Cascade [Nie et al., 2024] approach the cascade learning problem using reinforcement and imitation learning, respectively. TREACLE trains a deep Q-network to learn cascade decisions, while Online Cascade distills knowledge from stronger models into weaker ones. However, both methods rely on large supervised training sets and are less practical in label-scarce environments. Additionally, Online Cascade requires fine tuning of model parameters, which is very costly.

LLM routing approaches are another prominent class of multi-LLM inference, but they fall outside the scope of this work. Routers select the best LLM *prior to inference*, often relying on learned or heuristic models to choose among available options based solely on input characteristics. In contrast, cascades defer this decision until *after* an LLM has generated a response, leveraging internal confidence signals to determine whether to exit or escalate. Routing methods like RouteLLM [Ong et al., 2025], SelectLLM [Maurya et al., 2024], and Symbolic-MoE [Chen et al., 2025b] require task-specific training data and often struggle to generalize beyond their training distribution. Cascades, by contrast, offer greater flexibility by basing decisions on the output behavior of models rather than input features alone, at the price of multiple LLM calls per query.

3 Problem Statement

We consider a system comprised of m large language models (LLMs), denoted $\mathcal{M} = \{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_m\}$, each parameterized by a neural conditional generative model $f_j : \mathcal{X} \mapsto \mathcal{A}$ which maps from the prompt space \mathcal{X} to the answer space \mathcal{A} . Additionally, we assume that the j -th model also returns a confidence score $s_j \in [0, 1]$ along with its answer \hat{y}_j . We impose a mild correlation assumption on s_j that it is stochastically increasing as the probability of correctness of \hat{y}_j increases. The models are arranged in the order of increasing inference cost. That is, \mathcal{M}_1 is the cheapest, while \mathcal{M}_m (MPM) is the most expensive, and $\text{cost}(\mathcal{M}_i) \leq \text{cost}(\mathcal{M}_j)$ for $i < j$. For an input x drawn i.i.d. from a distribution \mathcal{X} , model \mathcal{M}_j produces a prediction \hat{y}_j , which is evaluated using the 0-1 regret with respect to the most powerful model’s prediction \hat{y}_m :

$$\ell(\hat{y}_j, \hat{y}_m) = \begin{cases} 0, & \text{if } \hat{y}_j = \hat{y}_m, \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

Here, we are operating in an unsupervised setting without access to the true labels, so the regret is formulated with respect to the most powerful LLM’s answer. That is, we never have access to the ground truth y , but only to the predictions of the m models $\{\hat{y}_j\}_{j=1}^m$.

To reduce inference cost while maintaining high accuracy, we employ a cascade decision rule. Given an input x , the cascade begins by querying the first model, \mathcal{M}_1 . After obtaining the output from \mathcal{M}_j , the system must decide whether to exit early and return \hat{y}_j if it is highly likely to match \hat{y}_m , or to escalate the query to \mathcal{M}_{j+1} if the current output is deemed unreliable. This process continues until an early exit occurs or the query x reaches \mathcal{M}_m , whose answer is always used if no earlier model qualifies. Following empirical results from Yue et al. [2024], we do not include outputs from model \mathcal{M}_j when prompting model \mathcal{M}_{j+1} , as it can potentially increase the error rate and prompting costs. Each model \mathcal{M}_j incurs an inference cost c_j , satisfying $c_1 \leq c_2 \leq \dots \leq c_m$. The cost associated with querying a model reflects monetary API fees, inference latency, or a combination thereof. We use a standard cost model from the literature that mimics real-world API cost structures [Chen et al., 2024]. The costs $\{c_j\}_{j=1}^m$ can be assumed to be fixed, known a priori, and query-independent for our analysis, since we know the per-token costs for different LLMs, the typical prompts and response lengths, and the latency. If the cascade stops at model \mathcal{M}_z , the total inference cost incurred is $\text{Cost} = \sum_{k=1}^z c_k$.

The objective is to design a cascade selection strategy that minimizes the total expected loss, measured as the regret relative to the best model \mathcal{M}_m (Eq. (1)), while ensuring that the expected inference cost

exceeds some prespecified budget C^* with a probability less than or equal to a fixed, small number α . Here $\alpha \in (0, 1)$ denotes the maximally allowed cost-constraint violation rate.

We model the binary decision to exit the cascade after querying \mathcal{M}_j by a learnable thresholding rule applied on s_j . Let $\boldsymbol{\tau} \stackrel{\text{def}}{=} [\tau_1, \dots, \tau_{m-1}, 0]$ be a vector of learnable thresholds of confidence scores $\mathbf{S} \stackrel{\text{def}}{=} [s_1, \dots, s_{m-1}, 1]$ such that each model \mathcal{M}_j has its own learnable confidence score threshold $\tau_j \in [0, 1]$. When s_j takes a value above the critical threshold τ_j this leads to an exit decision at \mathcal{M}_j . Note that for the last model in the cascade, $\tau_m = 0$ and $s_m = 1$. Thus the model always exits, as it is impossible to defer. We define the exit point of the cascade as the first model index for which the model confidence is greater than the threshold: $z(\mathbf{S}, \boldsymbol{\tau}) = \min_j \{j : s_j \geq \tau_j\}$.

To learn the optimal thresholds, we are given N questions sampled from \mathcal{X} for which we do not require ground truth labels. In addition, we have access to sampled responses and their confidence scores from each model for each of these N questions. We call these sample questions and LLM outputs the dataset $\mathcal{D} \stackrel{\text{def}}{=} \{x^i, (\hat{y}_1^i, s_1^i), \dots, (\hat{y}_m^i, s_m^i)\}_{i=1}^N$. At test time, the cascade output is evaluated by the percent agreement with the ground truth labels of a held-out dataset (accuracy). Thus, the primary learning objective is to solve the following optimization problem:

$$\min_{\boldsymbol{\tau}} L(\boldsymbol{\tau}), \text{ where } L(\boldsymbol{\tau}) \stackrel{\text{def}}{=} \mathbb{E}_X \left[\mathbb{E}_{(\hat{Y}_1, S_1), \dots, (\hat{Y}_m, S_m) | X} \ell(\hat{Y}_{z(\mathbf{S}, \boldsymbol{\tau})}, \hat{Y}_m) \right], \quad (2)$$

subject to a cost constraint. This constraint is soft, in the sense that the realized cost of answering a random query X violating the allowable budget C^* can be tolerated provided that such occurrences can only happen with a probability not exceeding α . This constraint is formally specified as:

$$\mathbb{E}_X \left[\mathbb{E}_{\mathbf{S} | X} \left[\mathbb{1} \left\{ \sum_{k=1}^{z(\mathbf{S}, \boldsymbol{\tau})} c_k > C^* \right\} \right] \right] \leq \alpha. \quad (3)$$

4 Methodology

We propose a principled methodology for optimizing exit thresholds in a LLM cascade, subject to a soft constraint on inference costs. The approach consists of three key components: (i) optimizing over thresholds to minimize empirical regret; (ii) conformal prediction-based adjustment for probabilistic cost control; and (iii) establishing generalization guarantees using PAC-Bayesian analysis.

To ensure our optimization and conformal guarantees are valid, we partition the dataset \mathcal{D} into the self-supervision data, \mathcal{D}_{SS} , which are used to learn the thresholds, and calibration data, \mathcal{D}_{Cal} , which are used for evaluating whether a particular threshold vector satisfies the conformal cost guarantee. Thus, we have two non-overlapping splits of the dataset $\mathcal{D} = \mathcal{D}_{\text{SS}} \cup \mathcal{D}_{\text{Cal}} \stackrel{\text{def}}{=} \{x^i, (\hat{y}_1^i, s_1^i), \dots, (\hat{y}_m^i, s_m^i)\}_{i=1}^{N_{\text{SS}} + N_{\text{Cal}}}$ with $N = N_{\text{SS}} + N_{\text{Cal}}$ and $\mathcal{D}_{\text{SS}} \cap \mathcal{D}_{\text{Cal}} = \emptyset$. Furthermore, we define an empirical approximation of Equation (2) using \mathcal{D}_{SS} as:

$$\hat{L}(\boldsymbol{\tau}) = \frac{1}{N_{\text{SS}}} \sum_{i=1}^{N_{\text{SS}}} \ell(\hat{y}_{z(\mathbf{S}^i, \boldsymbol{\tau})}, \hat{y}_m). \quad (4)$$

4.1 Discretization

The empirical 0–1 regret $\hat{L}(\boldsymbol{\tau})$ and cumulative cost are piecewise-constant functions of $\boldsymbol{\tau}$, with discontinuities occurring only when a threshold τ_j crosses a confidence score s_j in the finite dataset \mathcal{D}_{SS} . This structure arises because, while s_j may theoretically take continuous values, the dataset \mathcal{D}_{SS} contains at most N_{SS} distinct confidence scores per model. Sorting these scores for \mathcal{M}_j yields $N_{\text{SS}} + 1$ non-overlapping intervals, and varying τ_j within each such interval, while keeping all other thresholds fixed, does not alter the set of examples that exit at \mathcal{M}_j . As a result, $\hat{L}(\boldsymbol{\tau})$ and the cost remain constant within each interval; changes occur only at the empirical confidence values $\{s_j^{(i)}\}_{i=1}^{N_{\text{SS}}}$. Therefore, the loss surface is flat almost everywhere (with the exception of sampled s_j values) rendering gradients zero and gradient-based optimization ineffective.

Grid Setup For each model \mathcal{M}_j (except for the MPM), we construct a candidate grid of thresholds defined as: $\mathcal{T}_j = \{\frac{k}{K-2} \mid k \in \{0, 1, \dots, K-1\}\}$ with integer K defining the grid resolution. Note

Algorithm 1 Grid Search with Quantile Check

Require: Grids $\prod_{j=1}^{m-1} \mathcal{T}_j$, self-supervision data \mathcal{D}_{SS} , calibration set \mathcal{D}_{Cal} , budget C^* , risk level α

```
1: bestRegret  $\leftarrow \infty$ , bestTau  $\leftarrow (0, 0, \dots, 0)$ 
2: for all  $(\tau_1, \tau_2, \dots, \tau_{m-1}) \in \mathcal{T}_1 \times \dots \times \mathcal{T}_{m-1}$  do
3:   Assemble threshold vector  $\tau$ 
4:    $\hat{L} \leftarrow$  Evaluate eq. (4) on  $\mathcal{D}_{\text{SS}}$ 
5:   Evaluate all instances in  $\mathcal{D}_{\text{Cal}}$  and store the cascade costs  $\{C_j\}$ .
6:   Sort  $\{C_j\}$  to  $C_{(1)} \leq \dots \leq C_{(N)}$ 
7:    $k \leftarrow \lceil (N+1)(1-\alpha) \rceil$ 
8:    $q_{1-\alpha} \leftarrow C_{(k)}$ 
9:   if  $q_{1-\alpha} \leq C^*$  and  $\hat{L} < \text{bestRegret}$  then
10:     bestRegret  $\leftarrow \hat{L}$ 
11:     bestTau  $\leftarrow \tau$ 
12:   end if
13: end for
14: return bestTau
```

162 that this grid contains the edge points 0 and $(K-1)/(K-2)$, potentially allow the cascade system
163 to learn to always exit at a particular model, or to always skip it. This is useful for handling edge
164 cases. For example, when the reasoning problems are trivial, we would wish to always exit at the first
165 model. Alternatively, when a collection of weaker models produce effectively useless predictions, the
166 cascade would perform better by discarding those models entirely.

167 **Optimizing the cascade thresholds** Due to the non-differentiable 0–1 loss, the optimization is
168 performed by an exhaustive or heuristic search over all possible combinations of thresholds τ selected
169 from discrete grids. Note that this is usually not computationally unreasonable, since we are not
170 considering cascades of dozens of LLMs. For each candidate threshold configuration, the cascade
171 is evaluated on \mathcal{D}_{SS} to compute both the average 0–1 loss and the average cumulative cost. The
172 configuration that empirically minimizes $\mathcal{L}(\tau)$, while satisfying the cost constraint, is selected as
173 optimal. Pseudocode is provided in Algorithm 1.

174 4.2 Efficient Conformal Search Strategy

175 From a practical point of view, a brute-force search can be a viable approach in a setting with fewer
176 than 5 models in the cascade. For example, in our experiments, searching for the optimal threshold
177 using a brute-force search for a cascade of 4 LLMs with 10 discrete threshold levels for each of the
178 first 3 of them, takes approximately 0.01 seconds on a M3 CPU for GSM8K using 50 questions in
179 \mathcal{D}_{SS} . Compared to the inference time of LLMs response generation, this is a small amount of time.
180 Cascades are unlikely to be considerably larger than this in practice.

181 **Selecting Grid Resolution under Statistical Indistinguishability** The finite calibration size N_{SS}
182 induces a statistical resolution limit. Given N_{SS} samples, differences in empirical regret smaller than
183 $O(1/\sqrt{N_{\text{SS}}})$ cannot be distinguished using an appropriate statistical test (proof in Supplementary).
184 Thus, over-refinement of the grid does not offer any empirical benefit, yet inflates the size of
185 hypothesis class and degrades the theoretical generalization guarantee. We empirically observed that
186 using <10 grid points for each τ_j works in our experiments, and using more points is not beneficial.

187 **Grid-Search with Quantile Check** To enforce the probabilistic cost constraint $\Pr(C_{\text{test}} \leq C^*) \geq$
188 $(1-\alpha)$ during threshold optimization, we embed a quantile filter within the grid search. For each
189 candidate configuration τ , we evaluate the cascade on \mathcal{D}_{Cal} to compute the cumulative cost C^i
190 for each $i = 1, \dots, N_{\text{Cal}}$. We then sort the costs in the ascending order to obtain order statistics
191 $C_{(1)} \leq \dots \leq C_{(N)}$. We accept τ as a valid candidate to be considered as a cost-constrained
192 minimizer of the regret over \mathcal{D}_{SS} , if and only if the empirical $(1-\alpha)$ cost-quantile $q_{1-\alpha} \leq C^*$.
193 Among all accepted configurations, we select the τ with the minimal empirical regret \hat{L} on \mathcal{D}_{SS} . As
194 shown in Theorem 1, this ensures that the probabilistic cost-constraint is satisfied for a test query.

4.3 Theoretical Results

In this section, we provide important theoretical guarantees for our algorithm. First, in Theorem 1, we show that our search algorithm only returns solutions that are provably under the budget — if any such solutions exist. Second, in Theorem 2, we examine the generalization error of our algorithm using PAC-Bayes theory.

Theorem 1 (Conformal Cost Guarantee). *Let τ be the thresholds and \mathcal{D}_{Cal} denote a calibration set containing N_{Cal} questions and the cascade answers, obtained using thresholds τ . Sort the costs of the cascade on the calibration dataset and define the rank of the budget C^* as $k \stackrel{\text{def}}{=} \min\{p : C_{(p-1)} \leq C^* \leq C_{(p)}\}$. If $k \geq \lceil (N_{\text{Cal}} + 1)(1 - \alpha) \rceil$ is satisfied, then the inference cost C_{test} for a new query can be bounded under exchangeability of calibration and test data with $\Pr(C_{\text{test}} > C^*) \leq \alpha$.*

Proof. See Supplementary Material. □

Theorem 1 demonstrates that, assuming exchangeability of samples from \mathcal{D}_{Cal} and the test set, any new test query can only exceed the budget constraint with probability at most α .

In Theorem 2, we provide PAC-Bayes analysis that bounds the excess test regret due to the evaluation on the learned thresholds from \mathcal{D}_{SS} in terms of the number of models, grid-size, and the number of samples used to minimize the empirical regret.

Theorem 2 (Generalization Bound). *Let $\mathcal{H} = \prod_{j=1}^{m-1} \mathcal{T}_j$ denote the hypothesis class of all threshold combinations with $|\mathcal{H}| = K^{(m-1)}$, where K denotes the grid-size and m is the number of LLMs in the cascade. Let $\mathcal{H}_c \subset \mathcal{H}$ be the subset of thresholds for which the cost-constraint in Eq. (3) is satisfied and τ^* is the learned threshold vector from C3PO using \mathcal{D}_{SS} and \mathcal{D}_{Cal} . Then, for any $\delta \in (0, 1)$, with probability at least $(1 - \delta)$, we have*

$$L(\tau^*) \leq \min_{\tau \in \mathcal{H}_c} L(\tau) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}. \quad (5)$$

Proof. See Supplementary Material. □

Discussion As a concrete example, take $m = 3$, $K = 10$, $N_{\text{SS}} = 150$, and $\delta = 0.05$. Then $\log |\mathcal{H}| = 2 \log 10 \approx 4.61$, and $\log(1/\delta) \approx 2.99$, so the upper bound on excess regret evaluates to

$$\epsilon = \sqrt{\frac{4.61 + 2.99}{300}} \approx 0.159. \quad (6)$$

Thus, the test regret w.r.t. to the answer of the MPM in the cascade can only exceed the minimum regret one could possibly attain while satisfying the cost criterion on the test set, by at most 0.159 with more than 95% certainty. In other words, the percentage of test questions for which C3PO’s answer does not agree with the MPM, but an agreement with MPM can be obtained by solving the optimization problem we propose directly on the test-set under the same budget constraint, cannot exceed 15.9% with probability more than 0.95. Note that, solving the optimization problem directly on the test-set is impractical, since we would need to query all LLMs for each test question to perform the grid search, defeating the purpose of cascade learning using cost-constrained optimization.

The generalization bound degrades logarithmically with K and linearly in $m-1$, and improves at the rate $O(1/\sqrt{N_{\text{SS}}})$. This shows a clear tradeoff. Increasing the grid resolution K or the number of cascade stages m improves fitting of the training data during optimization but may adversely affect generalization guarantees if N_{SS} remains constant. Finally, note that the guarantee requires the calibration set to be i.i.d. from the data distribution and assumes that the hypothesis class \mathcal{H} is fixed prior to observing the calibration data. Under these conditions, we obtain a remarkable result: the empirical risk minimizer over a combinatorially large configuration space enjoys high-probability generalization guarantees with exponential concentration.

235 5 Experiments

236 5.1 Datasets

237 We evaluate our method across a diverse suite of 16 datasets spanning three distinct reasoning cate-
238 gories: logical reasoning, arithmetic reasoning, and mathematical reasoning. The **logical reasoning**
239 group includes tasks designed to assess abstract and commonsense inference capabilities, such as
240 *CommonSenseQA* [Talmor et al., 2019], and *BIG-Bench-Hard* [Suzgun et al., 2023] tasks, including
241 *causal judgement*, *date understanding*, *disambiguationQA*, *formal fallacies*, *geometric shapes*, *movie*
242 *recommendation*, *penguins*, *ruin names*, *snarks*, *sports*, and *temporal sequences*. These benchmarks
243 emphasize multi-step deduction, ambiguity resolution, and contextual judgment. In contrast, the
244 **arithmetic reasoning** group contains datasets like *GSM8K* [Cobbe et al., 2021], *SVAMP* [Patel
245 et al., 2021], and *AQuA* [Ling et al., 2017], which require basic numerical and word-problem-solving
246 skills, often through step-by-step calculation. Finally, the **mathematical reasoning** category targets
247 advanced problem-solving ability, represented by the *MATH-500* [Hendrycks et al., 2021] dataset,
248 which consists of high-school level and competition-style math questions that demand rigorous
249 algebraic and geometric manipulation. These datasets enable a nuanced assessment of reasoning
250 capabilities across domains. All of these benchmarks are available under open-source licenses ¹.

251 5.2 Models

252 We conduct experiments on both open and closed source LLMs. We use the following model families:
253 **LLAMA**: The open weights Meta models comprise Llama 3.2-3B-Instruct, Llama 3.3-70B-Instruct,
254 and Llama-3.1-405B-Instruct. These models have been fine-tuned for instruction following and
255 exhibit competitive performance on reasoning tasks.
256 **QWEN**: The QWEN family includes Qwen2.5-1.5B-Instruct, Qwen2.5-32B-Instruct, and Qwen2.5-
257 72B-Instruct. These open source models balance efficiency with advanced reasoning capabilities.
258 **GPT**: The OpenAI model family includes GPT 3.5 Turbo, GPT-4o-mini, and o3-mini. These models
259 are renowned for their strong few-shot learning capabilities and robust chain-of-thought reasoning.
260 In our experiments, they serve as a benchmark for high-performance proprietary LLMs.

261 5.3 Baselines

262 We compare C3PO with several existing cascade LLM approaches from the literature. **Frugal-**
263 **GPT** [Chen et al., 2024] learns the probability of correctness of a solution using a lightweight,
264 supervised training procedure from DistilBert embedding of the question and output CoT. Subse-
265 quently, it employs a threshold decision rule for cascading several LLMs. **Mixture of Thoughts**
266 **(MoT)** [Yue et al., 2024] is an unsupervised cascading framework involving a weak and a strong
267 LLM, which rely on an LLM’s confidence in the most frequent answer when a single prompt is
268 provided. MoT proposes a family of similar cascade architectures. We select the specific method
269 that matches our setting: CoT-1D-Vote. **TREACLE** [Zhang et al., 2024] learns a reinforcement
270 learning (RL)-based decision rule tailored for LLM reasoning, which requires true labels for a subset
271 of questions for Deep Q-Network (DQN) training. The original TREACLE architecture includes a
272 prompt adaptation component; we remove this in our implementation to adapt to our setting where all
273 models are evaluated on fixed prompts. **Model Switch** [Chen et al., 2025a] is an unsupervised early
274 exit module with the option to defer to a more elaborate model. If no model is sufficiently confident
275 in its prediction, the model outputs a weighted ensemble prediction.

276 Note that all our baselines are applicable to the cascade setting, where we assume that routing
277 decisions are made *after* generating a candidate response. This is a different setting from the routing
278 methods that must make routing decisions before seeing candidate outputs.

279 5.4 Results

280 **Can C3PO outperform existing cascades at different budgets?** We conduct experiments by
281 organizing the models in the three cascades listed in Section 5.2: the LLAMA cascade which consists
282 of four LLAMA-3 models of 1B, 3B, 70B and 405B variants, the QWEN-2.5 cascade 1.5B, 32B,
283 72B models and the GPT cascade consisting of GPT 3.5 Turbo, GPT-4o-mini and GPT-4. All models

¹Apache 2.0 [AQuA] and MIT [CommonSenseQA; SVAMP; GSM8K; MATH-500; BIG-Bench Hard].

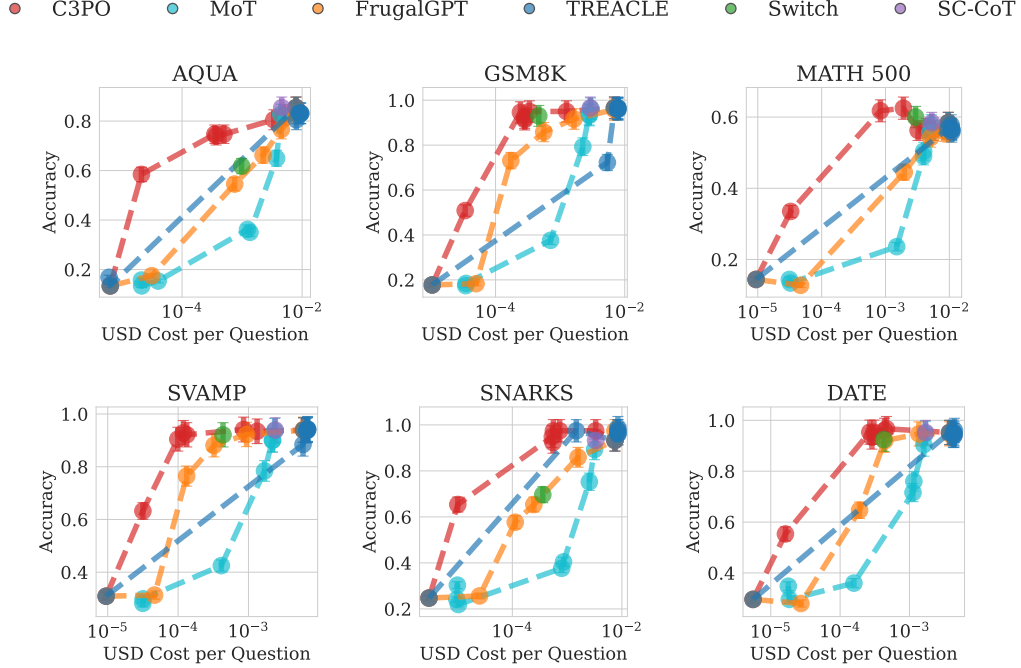


Figure 2: C3PO achieves the best performance for a wide range of cost budgets on the LLAMA cascade. Error bars denote 90% confidence interval. GPT and QWEN cascade results in Supplementary.

are prompted with 8-shot examples and during evaluation output 5 CoT samples. For all cascades we conduct the same experiment independently. Each time, for a given dataset, we deploy C3PO and the baselines with a training set of 100 reasoning problems. Unsupervised methods such as C3PO, Mixture-of-Thoughts, and ModelSwitch receive only the questions and 5 CoT outputs from each LLM per question. Supervised methods such as FrugalGPT and TREACLE also receive the ground truth labels to the reasoning problems. The seeds are fixed such that all methods receive the same questions and sample identical CoTs during training and testing. Additionally, the few shot examples in the prompts are identical for all methods. To evaluate the methods we compare them across a range of given budgets. The range extends from trivially using the cheapest model of the cascade with 5 CoT samples, all the way up to the cost equivalent to going through the entire cascade for all questions with 5 CoT samples per model. We note that since TREACLE can choose to early exit with an arbitrary number of CoT samples it sometimes chooses to exit with fewer than 5 CoTs. It can thus choose to use even lower than the minimal budget in some easy cases.

The results for the LLAMA cascade for AQUA, GSM8K, MATH-500, SVAMP and a subset of BigBenchHard (SNARKS, DATE) are depicted in Fig. 2. We observe that in all cases our model outperforms for most cost configurations. **Notably, for large datasets such as SVAMP and MATH-500 we obtain the top accuracy performance at 10 times lower the cost compared to other cascade baselines.** The performance for other datasets and the QWEN and GPT cascades is qualitatively similar and reported in the Supplementary Material.

How does C3PO compare to Self-Consistency for a comparable maximum allowable budget? Form Fig. 2, we observe that in all cases, C3PO achieves an accuracy, which is either comparable to or higher than that of self-consistency using the MPM at a fraction of the incurred cost. For example, from Figure 2, we observe that **on MATH-500, C-3PO achieves 62.5% accuracy with cost of 0.0019 USD/question, whereas SC using MPM manages to obtain 57% accuracy at a cost of 0.0053 USD/question.** This comparison validates the core philosophy of employing a LLM-cascade as an effective cost-saving alternative to performing self-consistency using a strong LLM.

Does C3PO actually satisfy the theoretical conformal guarantee for cost and the generalization bound? To evaluate whether C3PO satisfies its theoretical guarantee on inference cost, we consider

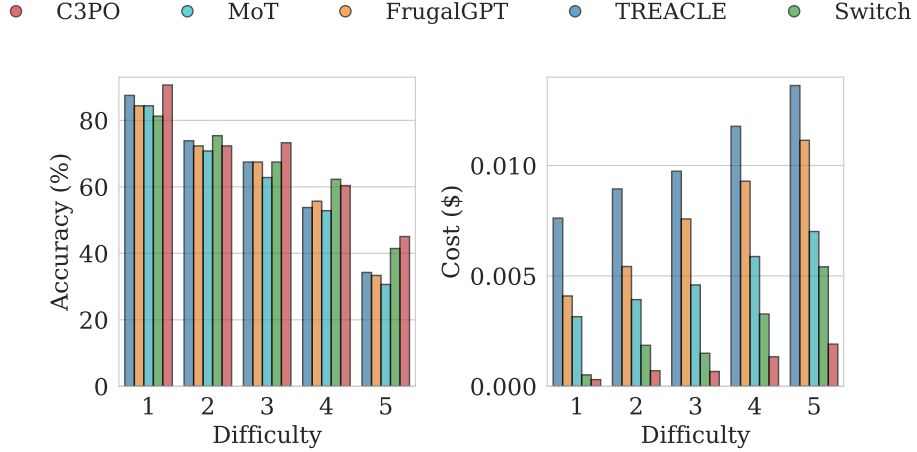


Figure 3: Average accuracies and dollar costs of different algorithms for five different difficulty levels of the MATH-500 dataset using the LLAMA cascade.

a collection of 15 datasets, 2 different cascade configurations (using LLAMA and Qwen models), 5 different target budget (C^*) levels for each cascade, and two different levels of α (0.05 and 0.1). For each combination of dataset, cascade, budget level, and α , we calibrate C3PO using a held-out calibration set and then evaluate it on a disjoint test set. The conformal guarantee implies that the proportion of test examples for which the inference cost exceeds C^* (referred to empirical violation rate) should be at most α . A violation occurs when this is not satisfied. We observe only a single violation in 300 ($15 \times 2 \times 5 \times 2$) runs of C3PO, which provides strong empirical evidence in favor of the validity of the cost guarantee. More detailed results are provided in Supplementary Material. In addition, in each of these runs of C3PO, the generalization bound specified in Theorem 2 is empirically verified to be true, and in fact the train-test gap in regrets is significantly lower.

In comparison to baselines, how does C3PO perform and allocate budget for questions with varying difficulties? Although the main results demonstrate the beneficial cost-accuracy trade-off offered by C3PO on multiple benchmarks, they do not provide any details of the allocation of the computational budget across different questions with different difficulties. For the MATH-500 dataset, we have access to the ‘true’ difficulty of each question as part of the metadata, based on human annotation. Therefore, we conduct a detailed analysis to compare the average accuracy and incurred cost of all baselines for each difficulty level of MATH-500. For this analysis, we evaluate each method using the hyperparameters which resulted in their maximum accuracy in Figure 2. From Figure 3, we observe that for all algorithms, the accuracy decreases and inference cost increases, as the questions become more difficult. C3PO incurs significantly lower cost for each difficulty level, while having the highest overall accuracy.

6 Conclusion, Limitations & Broader Impact

In this work, we introduced C3PO, a novel cascade framework that dynamically decides whether to accept an intermediate model’s prediction or defer to the more expensive next model in the cascade based on learned thresholds and conformal cost calibration. Our methodology combines discrete grid search over thresholds with rigorous conformal guarantees on expected cost and PAC-Bayes generalization bounds, ensuring both theoretical soundness and practical reliability. Beyond these empirical gains, C3PO’s design is highly extensible: one can incorporate richer confidence surrogates, employ more advanced search heuristics, or integrate powerful verification models without altering the core cascade logic. **Limitations:** A limitation of our approach is that the conformal cost guarantee can lead to overly conservative cascade decisions leading the model to under-utilize the cost budget sometimes. **Broader impact:** By demonstrating that principled early-exit strategies can deliver both high accuracy and cost efficiency, our work paves the way for more sustainable and scalable deployment of large language models in real-world reasoning applications.

References

- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Proc. Adv. Neural Info. Proc. Sys. (NeurIPS)*, pages 24824–24837, Dec. 2022a.
- Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *Proc. IEEE Conf. Comp. Vision and Pattern Recognition (CVPR)*, pages 510–518, Kauai, HI, USA, Dec. 2001.
- Xuechen Zhang, Zijian Huang, Ege Onur Taga, Carlee Joe-Wong, Samet Oymak, and Jiasi Chen. Efficient contextual LLM cascades through budget-constrained policy learning. In *Proc. Conf. Neural Info. Proces. Syst. (NeurIPS)*, pages 91691–91722, Vancouver, Canada, Dec. 2024.
- Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. Large language model cascades with mixture of thought representations for cost-efficient reasoning. In *Proc. Int. Conf. Learning Representations (ICLR)*, Vienna, Austria, May 2024.
- Lingjiao Chen, Matei Zaharia, and James Zou. Frugalgpt: How to use large language models while reducing cost and improving performance. *Trans. Machine Learning Research (TMLR)*, 2024.
- Jieyu Zhang, Ranjay Krishna, Ahmed H Awadallah, and Chi Wang. Ecoassistant: Using LLM assistant more affordably and accurately, 2023.
- Pranjal Aggarwal, Aman Madaan, Ankit Anand, Srividya Pranavi Potharaju, Swaroop Mishra, Pei Zhou, Aditya Gupta, Dheeraj Rajagopal, Karthik Kappaganthu, Yiming Yang, Shyam Upadhyay, Manaal Faruqui, and Mausam . Automix: Automatically mixing language models. In *Proc. Conf. Neural Info. Proces. Syst. (NeurIPS)*, pages 131000–131034, Vancouver, Canada, Dec. 2024.
- Jianhao Chen, Zishuo Xun, Bocheng Zhou, Han Qi, Qiaosheng Zhang, Yang Chen, Wei Hu, Yuzhong Qu, Wanli Ouyang, and Shuyue Hu. Do we truly need so many samples? Multi-LLM repeated sampling efficiently scale test-time compute, 2025a.
- Lunyu Nie, Zhimin Ding, Erdong Hu, Christopher Jermaine, and Swarat Chaudhuri. Online cascade learning for efficient inference over streams. In *Proc. Int. Conf. Machine Learning (ICML)*, page 38071–38090, Vienna, Austria, Jul. 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 *Int. Conf. Learning Representations (ICLR)*, 2025.
- Kaushal Kumar Maurya, KV Srivatsa, and Ekaterina Kochmar. SelectLLM: Query-aware efficient selection algorithm for large language models, 2024.
- Justin Chih-Yao Chen, Sukwon Yun, Elias Stengel-Eskin, Tianlong Chen, and Mohit Bansal. Symbolic Mixture-of-Experts: Adaptive skill-based routing for heterogeneous reasoning, 2025b.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *Proc. Conf. N. American Chapter Assoc. Comput. Linguistics*, pages 4149–4158, Minneapolis, MN, USA, Jun. 2019.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In *Proc. Assoc. Comput. Linguistics (ACL)*, pages 13003–13051, Toronto, Canada, Jul. 2023.
- 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, 2021.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. Are NLP models really able to solve simple math word problems? In *Proc. Conf. North Amer. Chapter of the Associ. Comput. Linguistics (ACL)*, pages 2080–2094, Jun. 2021.

394 Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. Program induction by rationale genera-
395 tion: Learning to solve and explain algebraic word problems. In *Proc. Conf. Empirical Methods in*
396 *Natural Language Process. (EMNLP)*, Copenhagen, Denmark, Sep. 2017.

397 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
398 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *Proc.*
399 *Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 1–8, Dec. 2021.

400 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
401 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*,
402 pages 24824–24837, 2022b.

403 Pranjal Aggarwal, Aman Madaan, Yiming Yang, and Mausam. Let’s sample step by step: Adaptive-
404 consistency for efficient reasoning and coding with LLMs. In Houda Bouamor, Juan Pino, and
405 Kalika Bali, editors, *Proc. Conf. Empirical Methods in Natural Language Proces. (ACL)*, pages
406 12375–12396, Singapore, Dec. 2023.

407 Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Bin Sun, Xinglin Wang, Heda Wang, and
408 Kan Li. Turning dust into gold: distilling complex reasoning capabilities from llms by leveraging
409 negative data. In *Proc. AAAI Conf. Artificial Intell.*, 2024a.

410 Jiaze Zhu, Yingtao Shen, Jie Zhao, and An Zou. Path-Consistency: Prefix Enhancement for Efficient
411 Inference in LLM. *arXiv e-prints*, art. arXiv:2409.01281, August 2024. doi: 10.48550/arXiv.2409.
412 01281.

413 DiJia Su, Sainbayar Sukhbaatar, Michael Rabbat, Yuandong Tian, and Qinqing Zheng. Dualformer:
414 Controllable fast and slow thinking by learning with randomized reasoning traces. In *The Thirteenth*
415 *International Conference on Learning Representations*, 2025.

416 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
417 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
418 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
419 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
420 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
421 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
422 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
423 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
424 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
425 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,
426 Meng Li, Miaojuan Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang,
427 Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L.
428 Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang,
429 Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng
430 Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng
431 Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan
432 Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang,
433 Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen,
434 Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li,
435 Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang,
436 Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan,
437 Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia
438 He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong
439 Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha,
440 Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang,
441 Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,
442 Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen
443 Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.

444 Zishun Yu, Tengyu Xu, Di Jin, Karthik Abinav Sankararaman, Yun He, Wenxuan Zhou, Zhouhao
445 Zeng, Eryk Helenowski, Chen Zhu, Sinong Wang, et al. Think smarter not harder: Adaptive
446 reasoning with inference aware optimization. *arXiv preprint arXiv:2501.17974*, 2025.

447 Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. Mixture-of-Agents enhances
448 large language model capabilities. *arXiv preprint arXiv:2406.04692*, 2024a.

449 Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models
450 with pairwise comparison and generative fusion. In *Proc. Assoc. for Computat. Linguistics (ACL)*,
451 2023.

452 Costas Mavromatis, Petros Karypis, and George Karypis. Pack of LLMs: Model fusion at test-time
453 via perplexity optimization. In *First Conference on Language Modeling (COLM)*, 2024.

454 Quang H Nguyen, Duy C Hoang, Juliette Decugis, Saurav Manchanda, Nitesh V Chawla, and Khoa D
455 Doan. Metallm: A high-performant and cost-efficient dynamic framework for wrapping llms.
456 *arXiv preprint arXiv:2407.10834*, 2024.

457 Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and
458 Mengnan Du. The impact of reasoning step length on large language models. In *Findings of the
459 Association for Computational Linguistics: ACL 2024*. Association for Computational Linguistics,
460 2024.

461 Peiji Li, Kai Lv, Yunfan Shao, Yichuan Ma, Linyang Li, Xiaoqing Zheng, Xipeng Qiu, and Qipeng
462 Guo. Fastmcts: A simple sampling strategy for data synthesis, 2025.

463 Dujian Ding, Ankur Mallick, Chi Wang, Robert Sim, Subhabrata Mukherjee, Victor Rühle, Laks V. S.
464 Lakshmanan, and Ahmed Hassan Awadallah. Hybrid LLM: Cost-efficient and quality-aware query
465 routing. In *Int. Conf. Learning Representations (ICLR)*, 2024.

466 Alireza Mohammadshahi, Arshad Rafiq Shaikh, and Majid Yazdani. Routoo: Learning to route to
467 large language models effectively, 2024. URL <https://arxiv.org/abs/2401.13979>.

468 Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou, and Jingren Zhou.
469 Routing to the expert: Efficient reward-guided ensemble of large language models. In Kevin Duh,
470 Helena Gomez, and Steven Bethard, editors, *Proc. Conf. N.A. Chapter of the Assoc. Computat.
471 Linguistics: Human Language Tech.*, pages 1964–1974, Mexico City, Mexico, Jun. 2024.

472 Zesen Zhao, Shuwei Jin, and Z. Morley Mao. Eagle: Efficient training-free router for multi-
473 LLM inference, 2024. URL <https://arxiv.org/abs/2409.15518>.

474 Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, and Kan
475 Li. Escape sky-high cost: Early-stopping self-consistency for multi-step reasoning. In *Int. Conf.
476 Learning Representations (ICRL)*, 2024b.

477 Xinglin Wang, Shaoxiong Feng, Yiwei Li, Peiwen Yuan, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Yao
478 Hu, and Kan Li. Make every penny count: Difficulty-adaptive self-consistency for cost-efficient
479 reasoning. *arXiv preprint arXiv:2408.13457*, 2024b.

480 Jiabao Pan, Yan Zhang, Chen Zhang, Zuozhu Liu, Hongwei Wang, and Haizhou Li. DynaThink: Fast
481 or Slow? A Dynamic Decision-Making Framework for Large Language Models. *arXiv e-prints*,
482 art. arXiv:2407.01009, July 2024. doi: 10.48550/arXiv.2407.01009.

483 Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang,
484 Zengmao Wang, Ziwei Liu, Bo Du, Xianglong Liu, and Dacheng Tao. Dynamic parallel tree
485 search for efficient llm reasoning, 2025.

486 Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. C3ot: Generating shorter chain-of-thought
487 without compromising effectiveness, 2024.

488 Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young Yun.
489 Self-training elicits concise reasoning in large language models, 2025.

490 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong
491 Tian. Training large language models to reason in a continuous latent space, 2024. URL <https://arxiv.org/abs/2412.06769>.

493 Zhenyi Shen, Hanqi Yan, Linhai Zhang, Zhanghao Hu, Yali Du, and Yulan He. Codi: Compressing
494 chain-of-thought into continuous space via self-distillation, 2025a.

495 Xuan Shen, Yizhou Wang, Xiangxi Shi, Yanzhi Wang, Pu Zhao, and Jiuxiang Gu. Efficient reasoning
496 with hidden thinking, 2025b.

497 Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao,
498 and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning.
499 *arXiv preprint arXiv:2501.12570*, 2025.

500 Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai
501 Wang, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reasoning models. *arXiv*
502 *preprint arXiv:2503.04472*, 2025c.

503 Yuxiao Qu, Matthew YR Yang, Amrith Setlur, Lewis Tunstall, Edward Emanuel Beeching, Ruslan
504 Salakhutdinov, and Aviral Kumar. Optimizing test-time compute via meta reinforcement fine-
505 tuning. *arXiv preprint arXiv:2503.07572*, 2025.

506 Neeraj Varshney and Chitta Baral. Model cascading: Towards jointly improving efficiency and
507 accuracy of NLP systems. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Proc.*
508 *Conf. Empirical Methods in Natural Language Proces. (EMNLP)*, pages 11007–11021, Abu Dhabi,
509 United Arab Emirates, Dec. 2022.

510 Guillem Ramírez, Matthias Lindemann, Alexandra Birch, and Ivan Titov. Cache & distil: Optimising
511 API calls to large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar, editors,
512 *Findings of Assoc. Comput. Linguistics (ACL)*, pages 11838–11853, Bangkok, Thailand, Aug.
513 2024.

514 Jasper Dekoninck, Maximilian Baader, and Martin Vechev. A unified approach to routing and
515 cascading for llms. *arXiv preprint arXiv:2410.10347*, 2024.

516 Jinwu Hu, Yufeng Wang, Shuhai Zhang, Kai Zhou, Guohao Chen, Yu Hu, Bin Xiao, and Mingkui
517 Tan. Dynamic ensemble reasoning for LLM experts. *arXiv preprint arXiv:2412.07448*, 2024.

518 Neha Gupta, Harikrishna Narasimhan, Wittawat Jitkittum, Ankit Singh Rawat, Aditya Krishna
519 Menon, and Sanjiv Kumar. Language model cascades: Token-level uncertainty and beyond. In *Int.*
520 *Conf. Learning Representations (ICLR)*, Vienna, Austria, May 2024.

521 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha
522 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
523 models. In *Int. Conf. Learning Representations (ICRL)*, 2023.

C3PO: Optimized LLM Cascades with Probabilistic Cost Constraints for Reasoning

Supplementary material

525

526

Table of Contents

527	A Broader Impact, Limitations and Code	15
528	B Extended Literature Review	16
529	B.1 Overview of Efficient LLM Test Time Compute Allocation Literature	16
530	B.2 Efficient Multi-Model Inference for LLMs	17
531	C Upper Bound on the Minimal Detectable Change in Empirical Regret	19
532	D Conformal Cost Guarantee	19
533	E Generalization Guarantees	20
534	F Inference Cost Model	22
535	G Complete Experimental Results	23
536	G.1 LLAMA Results	23
537	G.2 QWEN Results	25
538	G.3 GPT Results	27
539	G.4 Comparison of C3PO and baselines in terms of performance and budget allocation for questions with varying difficulties in MATH-500	29
541	H Few-shot Prompting Strategy and Prompt Examples	30
542	I Case Study and Analysis	32
543	I.1 Question Level Case Study	32
544	I.2 Detailed study of C3PO’s cost efficiency	33

545

546

547

A Broader Impact, Limitations and Code

The development of **C3PO** (Cost Controlled Cascaded Prediction Optimization) has the potential to significantly improve the accessibility, sustainability, and scalability of large language model (LLM) deployments, especially in real-world reasoning tasks. As LLMs continue to expand in size and capability, their computational and monetary costs pose a critical bottleneck that limits their application in low-resource settings, as well as their environmental sustainability. C3PO addresses this issue by offering a label-free, conformally-controlled framework for optimizing cascaded inference, enabling accurate and efficient LLM use while providing formal guarantees on inference cost.

Social and Environmental Impact. By substantially reducing reliance on powerful LLMs for every query, C3PO can lower the environmental footprint associated with LLM inference. For institutions operating under strict compute or budget constraints such as nonprofits, schools, or healthcare systems, C3PO offers a practical path to deploying high-quality language systems without incurring prohibitive costs. Its ability to generalize from unlabeled data further reduces barriers to adoption in under-resourced domains where labeled data is scarce or unavailable.

Equity and Access. C3PO’s cost-efficiency and domain adaptability can democratize access to high-quality language technologies. For example, resource-constrained platforms (e.g., small online businesses) could adopt C3PO to offer reliable results at minimal cost, expanding reach without compromising performance.

Limitations, Risks and Misuse. As with any method that enables wider LLM deployment, C3PO may inadvertently facilitate the use of language models in settings lacking oversight or appropriate safeguards. Cost-reduction techniques might be applied indiscriminately to critical decision-making domains (e.g., legal or healthcare) without verifying the alignment of model behavior with task requirements. To mitigate these risks, it is important to emphasize that C3PO offers probabilistic control over cost but does not itself guarantee correctness or fairness of model outputs. Proper validation, human-in-the-loop design, and task-specific evaluation must accompany its use in high-stakes applications.

Explainability and Transparency. C3PO’s reliance on interpretable thresholds and its theoretical guarantees on both generalization and cost provide a degree of transparency that many black-box routing or reinforcement-based systems lack. These properties support responsible use and make C3PO amenable to audit and monitoring, especially in safety-critical deployments. Moreover, its minimal data requirements lower the need for large-scale data collection, which may help reduce privacy concerns associated with supervised training pipelines.

Conclusion. C3PO provides a scalable and principled framework for efficient LLM inference that is compatible with privacy-sensitive, low-resource, and budget-constrained environments. Its broader societal benefit lies in making LLM-powered reasoning more affordable and sustainable, while its limitations call for responsible deployment practices that pair efficiency with rigorous domain-specific evaluation. **We make our code available at** <https://anonymous.4open.science/status/C3PO-047D>.

B Extended Literature Review

B.1 Overview of Efficient LLM Test Time Compute Allocation Literature

Fig. 4 presents a comprehensive taxonomy of current approaches to reducing test-time computational cost for large language models (LLMs). At the highest level, the diagram bifurcates into two overarching strategies: those that employ multiple models in a coordinated fashion, and those that operate on a single model but optimize its internal inference process.

On the top branch of Fig. 4, single-model approaches achieve efficiency by modifying the inference process of one LLM. Non-fine-tuning methods adapt the prompt or sampling strategy such as Chain-of-Thought [Wei et al., 2022b] at test time. Some examples include adaptive self-consistency [Aggarwal et al., 2023], early-stopping tests for chain-of-thought samples [Li et al., 2024a], and dynamic path-consistency pruning [Zhu et al., 2024] to reduce redundant reasoning steps without additional training. Fine-tuning based approaches, by contrast, incur an offline training cost to produce a model that internally decides when to exit or which reasoning branches to pursue. Supervised fine-tuning variants embed early-exit logic into the model weights [Su et al., 2025], while reinforcement learning variants use reward signals to guide policy learning for token or chain pruning [DeepSeek-AI et al., 2025, Yu et al., 2025]. DeepSeek-AI et al. [2025] has shown that one can prune CoT generation *inside* a single LLM via an internal reward model. We note that our framework naturally subsumes such approaches as models inside the cascade. As such, any model equipped with an internal CoT stopping rule can participate in the cascade, using our self-supervised thresholding and conformal-cost machinery to decide whether to exit.

On the lower branch of Fig. 4, multiple-model techniques are organized into cascaded pipelines and ensemble methods. Cascaded pipelines invoke a sequence of LLMs ordered by increasing computational expense: an entry-level model attempts the task first, and only if its output fails a decision criterion does the system escalate to a more powerful model. Within this category, unsupervised cascades require no labeled data, relying instead on internal agreement or consistency signals among model outputs. Notable instances include methods that sample multiple reasoning chains and halt when they converge (e.g., Mixture-of-Thoughts [Wang et al., 2024a]), or that compare a weaker model’s output against a stronger model for self-supervised thresholding (as in C3PO). In contrast, supervised and reinforcement-learning cascades use labeled examples or reward signals to train a meta-model or policy that decides when to stop or escalate—for example, FrugalGPT [Chen et al., 2024] uses a fine-tuned classifier to predict correctness, while TREACLE [Zhang et al., 2024] applies deep Q-learning to balance cost against accuracy.

Also under the multiple-model umbrella, ensemble methods combine responses from several LLMs either during inference (by aggregating or voting on outputs from parallel queries) or before inference (by routing each prompt to a single selected model based on input features). During-inference ensembles such as LLM-Blender [Jiang et al., 2023] and PackLLM [Mavromatis et al., 2024] query all candidate models and merge their outputs in real time, often improving robustness at the expense of higher instantaneous compute. Pre-inference routing methods like RouteLLM [Ong et al., 2025] and MetaLLM [Nguyen et al., 2024] train lightweight selectors to choose the most suitable model for each query, thereby avoiding sequential calls but requiring supervised training on a massive training set or online adaptation to maintain accuracy across domains. We note the special case of Automix [Aggarwal et al., 2024]. While Automix is a sequential method in the sense that it calls each candidate LLM one at a time, it also uses an external verifier LLM. Thus, due to this presence of at least 2 LLM models per output we classify it as an ensemble method.

By tracing paths from the root through these subcategories, we observe that C3PO occupies the multiple-model, cascade, unsupervised quadrant. Unlike supervised cascades, it requires no labeled data; unlike heuristic confidence-threshold cascades, it leverages self-supervised agreement metrics and conformal prediction to enforce probabilistic cost bounds; and unlike pure routing or ensemble methods, it dynamically defers per input based on realized outputs. This unified view highlights both the diversity of existing techniques and the unique position of C3PO in providing label-free, theoretically grounded cost control for cascaded LLM inference.

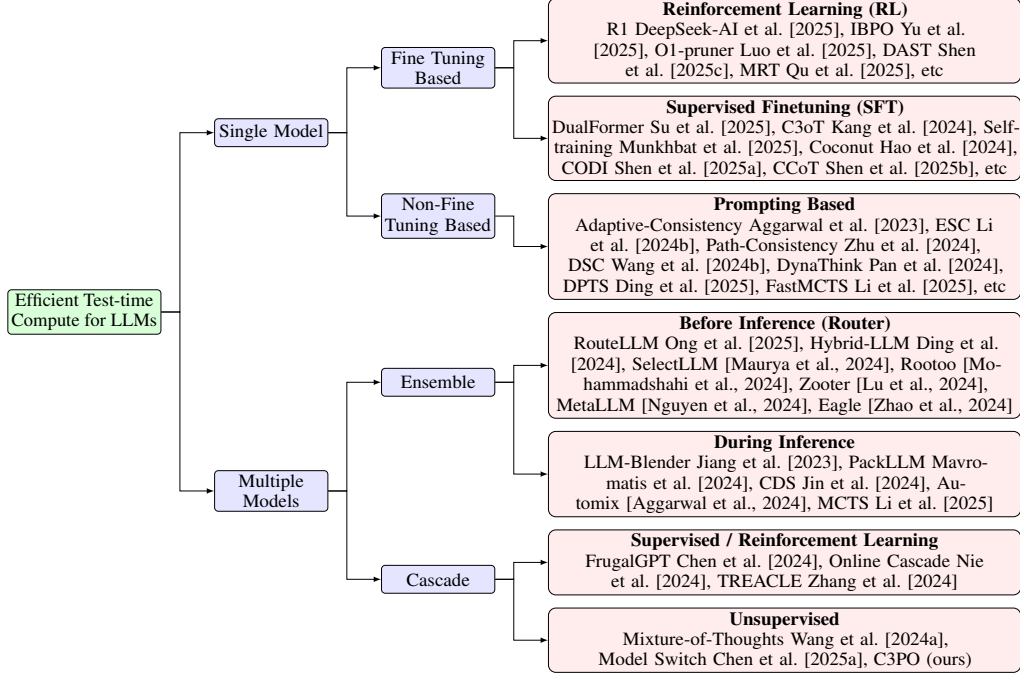


Figure 4: Taxonomy of recent research on efficient test-time compute for LLMs.

B.2 Efficient Multi-Model Inference for LLMs

Inference cost is a critical barrier to deploying large language models (LLMs) at scale. A prominent line of research reduces average compute by *cascading* or *routing* queries through multiple LLMs of increasing size and cost, so that cheaper models handle *easy* inputs and only *hard* cases escalate to more expensive ones. Broadly, existing methods differ in (1) their supervision requirements; (2) decision criteria for stopping or escalating; and (3) whether they provide formal guarantees on cost or accuracy. Below, we review representative approaches, grouping them by supervision regime and decision strategy, and highlight how our method, C3PO, addresses outstanding gaps.

Unsupervised and Heuristic Cascades Early cascade strategies adapted classical classifier-cascade ideas to LLMs without requiring any labeled data. Varshney and Baral [2022] leverage a model’s softmax-based confidence: if a small LLM’s maximum token probability exceeds a heuristic threshold, decoding halts; otherwise, the next model is queried. While simple and label-free, this approach relies on manually tuned thresholds and offers no probabilistic cost control. Yue et al. [2024] propose *Mixture-of-Thoughts (MoT)*, which samples multiple reasoning chains from a weak model and halts when self-consistency among samples is high. MoT improves decision robustness via ensemble-style voting, but incurs extra sampling cost and still lacks formal guarantees on budget overruns.

Neural caching techniques [Ramírez et al., 2024] introduce online distillation of a strong model into a weak one: when the weak model’s confidence is low, the strong model answers and its output augments a cache that refines the weak model over time. This iterative distillation reduces repeated expensive calls, but the cache requires streaming data and provides no worst-case cost bounds. Another line, exemplified by Dekoninck et al. [2024], learns a lightweight *cascade and router hybrid model* using unlabeled internal features (e.g., token entropy) to predict whether to exit or escalate. Under mild distributional assumptions, this router can enforce a user-specified cost bound with high probability, yet it often needs some held-out calibration and does not directly optimize for regret relative to the oracle.

Supervised Meta-Model and Reinforcement Learning Approaches To improve predictiveness of the exit decision, several works resort to supervised learning on labeled meta-data. Chen et al. [2024] train a BERT-style meta-model to estimate each LLM’s correctness probability and stop when confidence exceeds a learned threshold. FrugalGPT delivers strong empirical speedups, but depends

on large annotated datasets and lacks formal cost-violation guarantees. Reinforcement-based methods such as *AutoMix* [Aggarwal et al., 2024] and *DER* [Hu et al., 2024] frame cascade control as a Markov decision process, learning a policy to decide both whether to stop and which successor model to call. These methods achieve fine-grained cost-accuracy trade-offs but incur high meta-training overhead and offer only empirical, rather than provable, guarantees. Additionally, these approaches can be called hybrid router-cascade approaches as they either use more than one LLM during inference (*AutoMix*) or they are free to traverse the cascade in any order they choose (*DER*) as opposed to fixed cheap-to-expensive route allowed in classic cascade methods. Thus by being allowed to jump to an arbitrary model *DER* resembles more of a sequential router rather than a cascade.

Quantile-feature cascades [Gupta et al., 2024] represent a middle ground. Gupta et al. [2024] extract quantile statistics from token-level confidences to select whether to exit early or to escalate to a larger model. An issue with this formulation is that it requires access to model internals (embeddings), and it is hence not applicable to closed source models.

Discussion and Positioning of C3PO Table 1 summarizes key attributes of these methods. Most unsupervised cascades use confidence or consistency heuristics but lack formal cost guarantees; supervised and RL methods offer strong empirical performance at the expense of labeled data and black-box policies; routing methods trade multiple calls for a single, pre-selection step but cannot adapt dynamically post-response. None combine (a) *self-supervised* training from unlabeled agreement signals, (b) *conformal* calibration to bound the probability of exceeding a user budget, and (c) *PAC-Bayesian* generalization guarantees on accuracy regret.

In contrast, C3PO is entirely label-free, using only *self-supervision* from off-the-shelf model outputs; it employs *conformal prediction* to guarantee that inference cost exceeds a user-specified threshold with bounded probability; and it derives *PAC-Bayesian* bounds on accuracy regret. This unique combination makes C3PO the first method to blend unsupervised cascade learning with rigorous cost and generalization guarantees, addressing key limitations of prior work.

Method	Supervision	Signal	Cost Bounds
Self-Consistency [Wang et al., 2023]	None	Sample consensus	None
Mixture-of-Thoughts (MoT) [Yue et al., 2024]	None	Self-consistency voting	None
ESC / Adaptive SC [Li et al., 2024b,a]	None	Stopping tests	None
Neural Caching [Ramírez et al., 2024]	None	Conf. + distill	None
Cascade Routing [Dekoninck et al., 2024]	Weak	Entropy/router	Prob. (w/o labels)
FrugalGPT [Chen et al., 2024]	Full	Meta-model score	Empirical
DER [Hu et al., 2024]	Full	RL policy	Empirical
AutoMix [Aggarwal et al., 2024]	Full	RL policy	Empirical
TREACLE [Zhang et al., 2024]	Full	RL policy	Empirical
Online Cascade [Nie et al., 2024]	Full	RL policy	Empirical
Quantile Cascades [Gupta et al., 2024]	Full	Quantile features	Heuristic
ModelSwitch [Chen et al., 2025a]	None	Vote disagreement	None
RouteLLM [Ong et al., 2025]	Full	Input features	None
C3PO	None	Agreement signal	Conformal

Table 1: Comparison of multi-LLM inference methods: supervision level, decision signal, and cost guarantees. Highlighted rows denote our primary baselines and proposed method. **Our baselines and proposed method are the only methods that (i) are pure cascade methods; (ii) do not fine tune the LLMs; and (iii) can be applied to both open and closed source LLMs.**

C Upper Bound on the Minimal Detectable Change in Empirical Regret

Theorem 3 (Upper Bound on the MDC in Empirical Regret). *Let $\hat{L}(\tau)$ be the empirical 0–1 regret over N_{SS} samples at threshold τ . Denote its standard error by $\hat{\sigma}_{\hat{L}(\tau)} = \sqrt{\hat{L}(\tau)(1 - \hat{L}(\tau))/N_{\text{SS}}}$. Let $z_{(1-\alpha/2)}$ denote the $(1 - \alpha/2)$ quantile of the standard normal distribution, $\mathcal{N}(0, 1)$. Absolute differences between independently estimated empirical regrets at any two thresholds τ and τ' (i.e., $|\hat{L}(\tau) - \hat{L}(\tau')|$), which is smaller than $\Delta_{\min} = z_{(1-\alpha/2)} \sqrt{\hat{\sigma}_{\hat{L}(\tau)}^2 + \hat{\sigma}_{\hat{L}(\tau')}^2}$, are statistically indistinguishable at the confidence level $(1 - \alpha)$ for any $\alpha \in (0, 1)$. In addition, $\Delta_{\min} \leq z_{(1-\alpha/2)} \sqrt{\frac{1}{2N_{\text{SS}}}}$.*

Proof. The empirical regret $\hat{L}(\tau) \in [0, 1]$ is a binomial proportion with variance:

$$\text{Var}(\hat{L}(\tau)) = \frac{L(\tau)(1 - L(\tau))}{N_{\text{SS}}}. \quad (7)$$

where $L(\tau) = \mathbb{E}[\hat{L}(\tau)]$ is the true regret. Since we do not know $L(\tau)$, we cannot compute $\text{Var}(\hat{L}(\tau))$. Instead, we define the standard error of $\hat{L}(\tau)$ by $\hat{\sigma}_{\hat{L}(\tau)}$, which is the square root of the estimated variance, therefore, $\sigma_{\hat{L}(\tau)} = \sqrt{\frac{\hat{L}(\tau)(1 - \hat{L}(\tau))}{N_{\text{SS}}}}$. Note that, $\sigma_{\hat{L}(\tau)} \leq \sqrt{\frac{1}{4N_{\text{SS}}}}$, where the equality is attained at $\hat{L}(\tau) = \frac{1}{2}$.

We define a **hypothesis test for distinguishability**. For thresholds τ and τ' , define the Z -statistic

$$Z = \frac{\hat{L}(\tau) - \hat{L}(\tau')}{\sqrt{\hat{\sigma}_{\hat{L}(\tau)}^2 + \hat{\sigma}_{\hat{L}(\tau')}^2}}. \quad (8)$$

Under $H_0 : L(\tau) = L(\tau')$, and $Z \sim \mathcal{N}(0, 1)$. H_0 can be rejected at level α if $|Z| > z_{(1-\alpha/2)}$.

The critical value of $|\hat{L}(\tau) - \hat{L}(\tau')|$, induced by this hypothesis test is termed Minimal Detectable Change (MDC) in empirical regret and is denoted by Δ_{\min} .

$$\Delta_{\min} = z_{(1-\alpha/2)} \sqrt{\hat{\sigma}_{\hat{L}(\tau)}^2 + \hat{\sigma}_{\hat{L}(\tau')}^2} \quad (9)$$

$$\leq z_{(1-\alpha/2)} \sqrt{\frac{1}{2N_{\text{SS}}}}. \quad (10)$$

As an example, for $\alpha = 0.05$ and worst-case $\hat{L} = 0.5$, $\Delta_{\min} \leq 1.38/\sqrt{N_{\text{SS}}}$.

□

D Conformal Cost Guarantee

Theorem 4 (Conformal Cost Guarantee). *Let τ be the thresholds and \mathcal{D}_{Cal} denote a calibration set containing N_{Cal} questions and the cascade answers, obtained using thresholds τ . Sort the costs of the cascade on the calibration dataset and define the rank of the budget C^* as $k \stackrel{\text{def}}{=} \min\{p : C_{(p-1)} \leq C^* \leq C_{(p)}\}$. If $k \geq \lceil (N_{\text{Cal}} + 1)(1 - \alpha) \rceil$ is satisfied, then the inference cost C_{test} for a new query can be bounded under exchangeability of calibration and test data with $\Pr(C_{\text{test}} > C^*) \leq \alpha$.*

Proof. Let R denote the rank of C_{test} among $\{C_i\}_{i=1}^{N_{\text{Cal}}} \cup \{C_{\text{test}}\}$. By exchangeability, we have,

$$\Pr(R = r) = \frac{1}{N_{\text{Cal}} + 1} \quad \text{for all } r \in \{1, \dots, N_{\text{Cal}} + 1\}. \quad (11)$$

Therefore,

$$\Pr(C_{\text{test}} > C^*) = \Pr(C_{\text{test}} > C_{(k)}) = \Pr(R > k) = \frac{N_{\text{Cal}} + 1 - k}{N_{\text{Cal}} + 1} \leq \alpha, \quad (12)$$

since $k \geq (N_{\text{Cal}} + 1)(1 - \alpha)$ by construction.

□

718 E Generalization Guarantees

719 **Theorem 5** (Generalization Bounds). *Let $\mathcal{H} = \prod_{j=1}^{m-1} \mathcal{T}_j$ denote the hypothesis class of all threshold*
 720 *combinations with $|\mathcal{H}| = K^{(m-1)}$, where K denotes the grid-size and m is the number of LLMs*
 721 *in the cascade. Let $\mathcal{H}_c \subset \mathcal{H}$ be the subset of thresholds for which the cost-constraint in Eq. (3)*
 722 *is satisfied and τ^* is the learned threshold vector from C3PO using \mathcal{D}_{SS} and \mathcal{D}_{Cal} . Then, for any*
 723 *$\delta \in (0, 1)$, with probability at least $(1 - \delta)$, we have*

$$L(\tau^*) \leq \hat{L}(\tau^*) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \quad (13)$$

$$\text{and, } L(\tau^*) \leq \min_{\tau \in \mathcal{H}_c} L(\tau) + 2\sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}. \quad (14)$$

724 *Proof.* For any threshold configuration $\tau \in \mathcal{H}$, we know that $0 \leq L(\tau), \hat{L}(\tau) \leq 1$. The empirical
 725 regret, $\hat{L}(\tau)$ is an unbiased estimator of the ‘true’ regret $L(\tau) = \mathbb{E}[\hat{L}(\tau)]$ on i.i.d. samples from
 726 \mathcal{D}_{SS} . Applying Hoeffding’s inequality, we have, for any τ and $\epsilon > 0$:

$$\Pr(L(\tau) > \hat{L}(\tau) + \epsilon) \leq e^{-2N_{\text{SS}}\epsilon^2}, \quad (15)$$

$$\text{and, } \Pr(\hat{L}(\tau) > L(\tau) + \epsilon) \leq e^{-2N_{\text{SS}}\epsilon^2}. \quad (16)$$

727 Applying an union bound over all $\tau \in \mathcal{H}$, we obtain:

$$\Pr(\exists \tau \in \mathcal{H} : L(\tau) > \hat{L}(\tau) + \epsilon) \leq \sum_{\tau \in \mathcal{H}} \Pr(L(\tau) > \hat{L}(\tau) + \epsilon) \leq |\mathcal{H}|e^{-2N_{\text{SS}}\epsilon^2}, \quad (17)$$

$$\text{and, } \Pr(\exists \tau \in \mathcal{H} : \hat{L}(\tau) > L(\tau) + \epsilon) \leq \sum_{\tau \in \mathcal{H}} \Pr(\hat{L}(\tau) > L(\tau) + \epsilon) \leq |\mathcal{H}|e^{-2N_{\text{SS}}\epsilon^2}. \quad (18)$$

728 Letting the right-hand side be equal to δ , we solve for ϵ :

$$\delta = |\mathcal{H}|e^{-2N_{\text{SS}}\epsilon^2} \Rightarrow \epsilon = \sqrt{\frac{\log(|\mathcal{H}|/\delta)}{2N_{\text{SS}}}} = \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}. \quad (19)$$

729 Therefore, with probability at least $(1 - \delta)$, we have the uniform bounds

$$\forall \tau \in \mathcal{H}, L(\tau) \leq \hat{L}(\tau) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \quad (20)$$

$$\text{and, } \hat{L}(\tau) \leq L(\tau) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \quad (21)$$

730 Since the bound in eq. (20) holds for all $\tau \in \mathcal{H}$, it holds for the learned threshold vector from C3PO,
 731 τ^* . This proves the claim stated in eq. (13). Essentially, it shows that the ‘true’ test regret at C3PO’s
 732 learned threshold cannot exceed the empirical regret on \mathcal{D}_{SS} by more than ϵ with a high probability,
 733 greater than $(1 - \delta)$.

734 Now, let us denote by $\tilde{\tau}$ the minimizer of true risk $L(\tau)$ over $\mathcal{H}_c \subset \mathcal{H}$, where \mathcal{H}_c is the subset of
 735 thresholds for which the cost-constraint in Eq. (3) is satisfied.

736 Now, w.p. at least $(1 - \delta)$, we have,

$$L(\boldsymbol{\tau}^*) \leq \widehat{L}(\boldsymbol{\tau}^*) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \text{ from eq. (13), since, } \boldsymbol{\tau}^* \in \mathcal{H}_c \quad (22)$$

$$\leq \widehat{L}(\tilde{\boldsymbol{\tau}}) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \text{ since, } \boldsymbol{\tau}^* \text{ is the minimizer of } \widehat{L}(\boldsymbol{\tau}) \text{ over } \mathcal{H}_c \quad (23)$$

$$\leq L(\tilde{\boldsymbol{\tau}}) + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}} + \sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \text{ using the bound in eq. (21)} \quad (24)$$

$$= \min_{\boldsymbol{\tau} \in \mathcal{H}_c} L(\boldsymbol{\tau}) + 2\sqrt{\frac{(m-1) \log K - \log \delta}{2N_{\text{SS}}}}, \text{ since, } \tilde{\boldsymbol{\tau}} \text{ is the minimizer of } L(\boldsymbol{\tau}) \text{ over } \mathcal{H}_c, \quad (25)$$

737 which proves the claim in eq (14). This shows that the ‘true’ regret at C3PO’s learned threshold
 738 cannot exceed the minimum attainable test regret by more than 2ϵ with a high probability, greater
 739 than $(1 - \delta)$.

740

□

F Inference Cost Model

In our experiments, we used various LLAMA, QWEN, and GPT models. All three GPT models are closed-source, but can be publicly accessed using the OpenAI API². The LLAMA and QWEN models are open-source, although in practice, we used a commercial API service³ to query these models due to hardware constraints. The pricing information for LLAMA and QWEN can be found at <https://nebius.com/prices-ai-studio> and the GPT token cost details are available at <https://openai.com/api/pricing/>. The detailed per token costs for different LLM families are available in Tables 2, 3, 4.

For any single query, the total cost of sampling multiple CoTs is computed by adding the dollar costs for prompting, i.e., the cost of processing the system prompt and input prompt, and the total cost of generating multiple CoTs. We conducted all experiments during April and May 2025 and used the pricing information available at that time for dollar cost accounting throughout the paper.

Although token costs change over time (typically at least once a year), we argue that this does not invalidate any conclusions from this work for the following reason. When pricing changes (usually becoming cheaper) the API providers ensure that the ratio of token costs between different LLMs within the same family remains consistent before and after the pricing change as the main driver of costs is the model memory which is constant.

This guarantees that the dollar cost of a single query using a fixed prompt scales predictably, and that cost calculations for all baseline algorithms and C3PO are uniformly affected. For example, a 20% reduction in token pricing corresponds to a 20% reduction in cost across all methods considered in this paper. Therefore, any comparison between two LLM cascades yields the same conclusion regardless of pricing changes.

Table 2: Cost per million tokens for LLaMA models. Price source: <https://nebius.com/prices-ai-studio>

LLaMA Model	Input Cost (\$/M tokens)	Output Cost (\$/M tokens)
LLaMA 3.2 1B-Instruct	0.005	0.01
LLaMA 3.2 3B-Instruct	0.01	0.02
LLaMA 3.3 70B-Instruct	0.13	0.40
LLaMA 3.1 405B-Instruct	1.00	3.00

Table 3: Cost per million tokens for QWEN models. Price source: <https://nebius.com/prices-ai-studio>

Qwen Model	Input Cost (\$/M tokens)	Output Cost (\$/M tokens)
Qwen 2.5 1B-Instruct	0.02	0.06
Qwen 2.5 32B-Instruct	0.06	0.20
Qwen 2.5 72B-Instruct	0.13	0.40

Table 4: Cost per million tokens for GPT models. Price source: <https://platform.openai.com/docs/pricing>

GPT Model	Input Cost (\$/M tokens)	Output Cost (\$/M tokens)
GPT 3.5-Turbo	0.50	1.50
GPT 4o-mini	0.15	0.60
OpenAI o3-mini	1.10	4.40

²<https://openai.com/api>

³<https://docs.nebius.com/studio/inference/api>

763 G Complete Experimental Results

764 In Figure 2 of the main paper, we compared the performance of all baselines and the proposed C3PO
 765 on 6 datasets using the LLAMA cascade. Here, we present the complete experimental results on 16
 766 datasets and 3 LLM cascades. Accuracy vs. cost plots using LLAMA, QWEN, and GPT cascades
 767 on 16 datasets are shown in Figures 5, 7, and 9 respectively. In addition, for each of the 3 LLM
 768 cascades, we include summary boxplots to compare the distributions of required inference costs to
 769 reach near-MPM accuracy across 16 datasets, to compare C3PO with the baseline algorithms. These
 770 boxplots for the LLAMA, QWEN and GPT cascades are shown in Figures 6, 8, and 10 respectively.

771 G.1 LLAMA Results

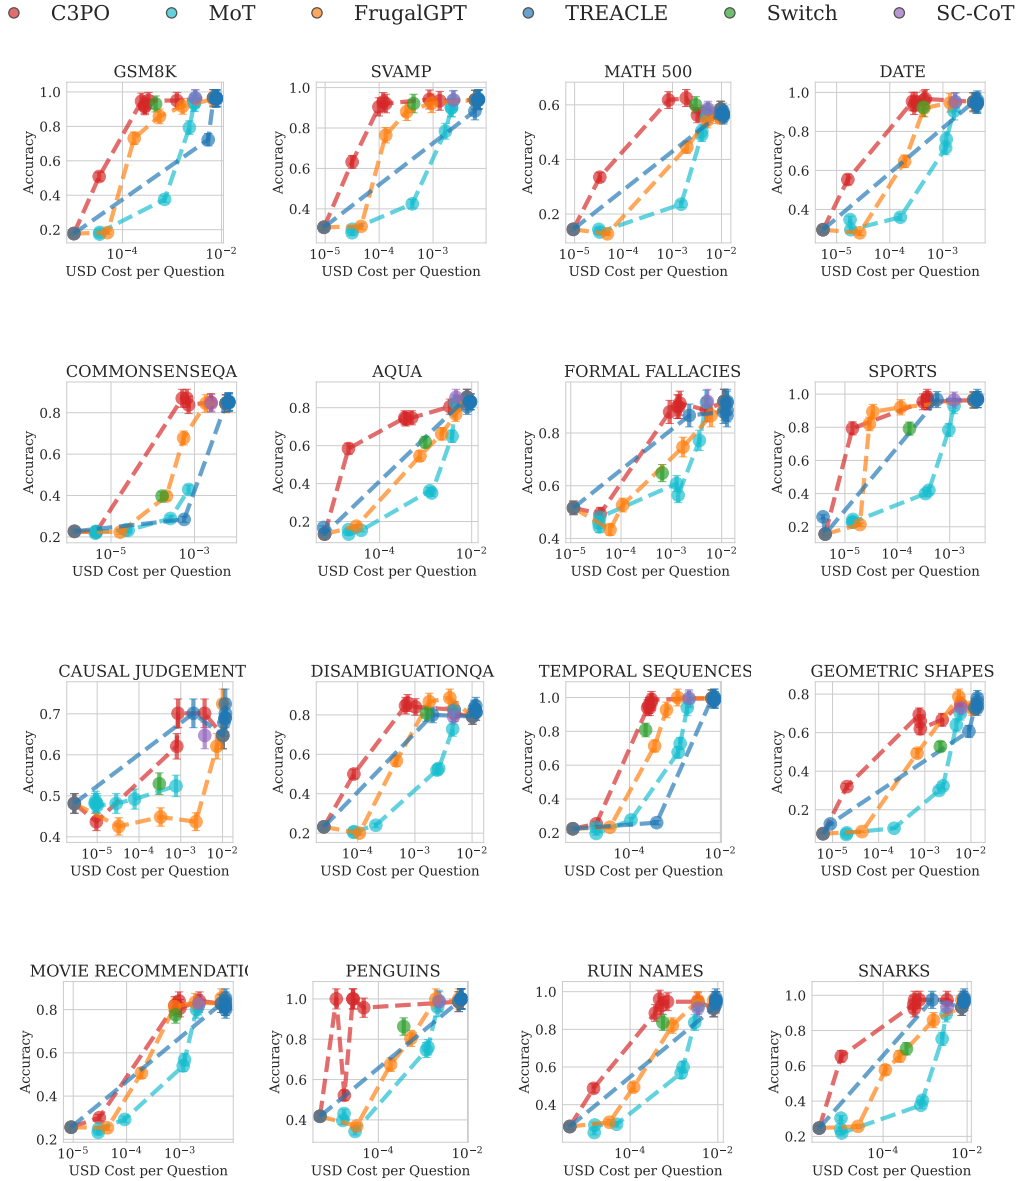


Figure 5: Accuracy vs. dollar cost of different algorithms for 16 datasets using the LLAMA cascade

772 From Figure 5, we observe that the proposed C3PO outperforms its competitors for most datasets and
 773 cost configurations. The accuracy improvement offered by C3PO in comparison to baselines is more

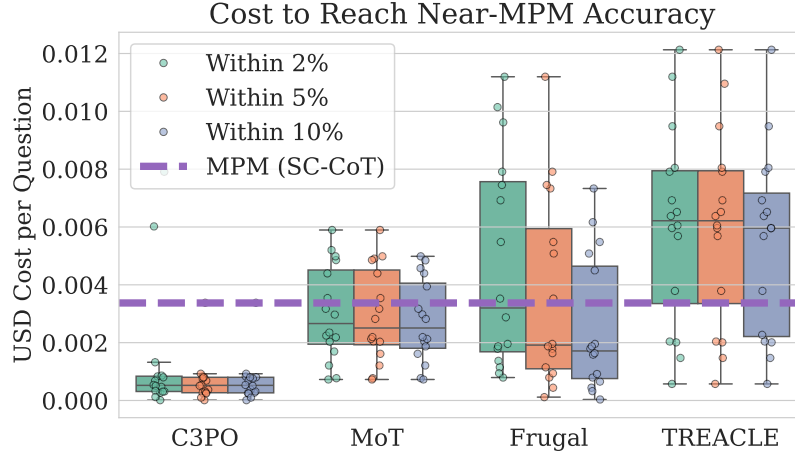


Figure 6: Surpassing existing cascade approaches tremendously, C3PO offers markedly superior cost-effectiveness across 16 benchmarks, requiring **less than 20% of the cost of the most powerful model (MPM)** (cost shown in purple) for an accuracy gap of at most 2, 5, and 10% using a LLAMA cascade. In this boxplot, each dot represents a dataset and the whiskers extend to 90% coverage.

774 pronounced for relatively lower dollar costs, demonstrating C3PO’s efficacy in resource-constrained
775 settings. This is further supported by Figure 6, which shows that across different datasets, C3PO
776 requires significantly lower costs compared to the baselines to achieve near-MPM accuracy. Similar
777 conclusions can be drawn for QWEN and GPT cascades from Figures 7, 8, 9, and 10.

778 Note that for the GPT cascade, we use o3-mini as the MPM as it is the most expensive of the three
779 GPT models we use to form the cascade. However, for some datasets, e.g. *sports*, *disambiguationQA*,
780 *movie recommendation*, o3-mini’s accuracy is significantly lower than that of a much cheaper GPT-
781 4o-mini. As a result, accuracies of all algorithms deteriorate with increasing budgets for those
782 datasets.

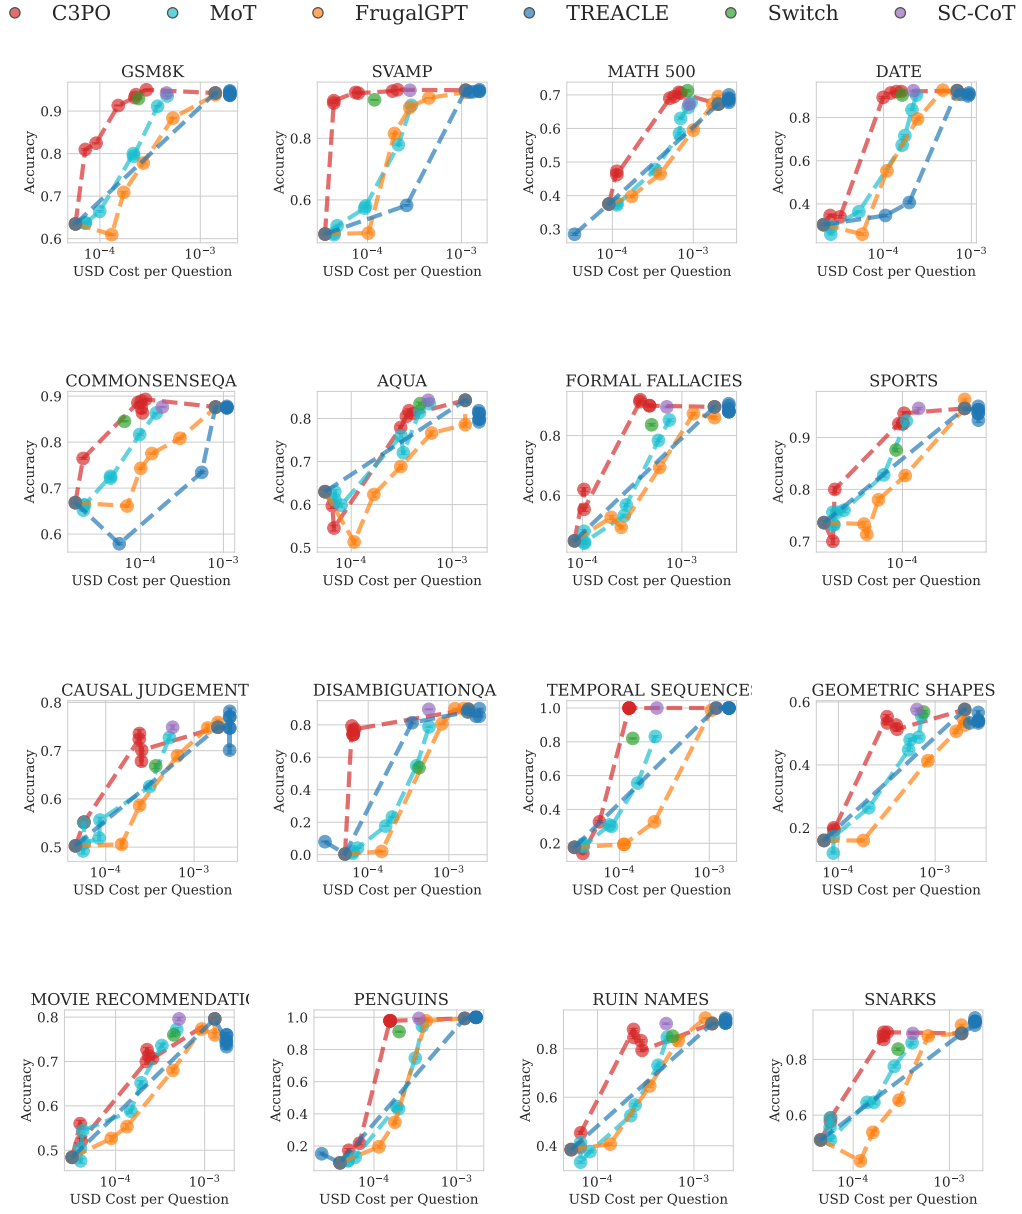


Figure 7: Accuracy vs. dollar cost of different algorithms for 16 datasets using the QWEN cascade

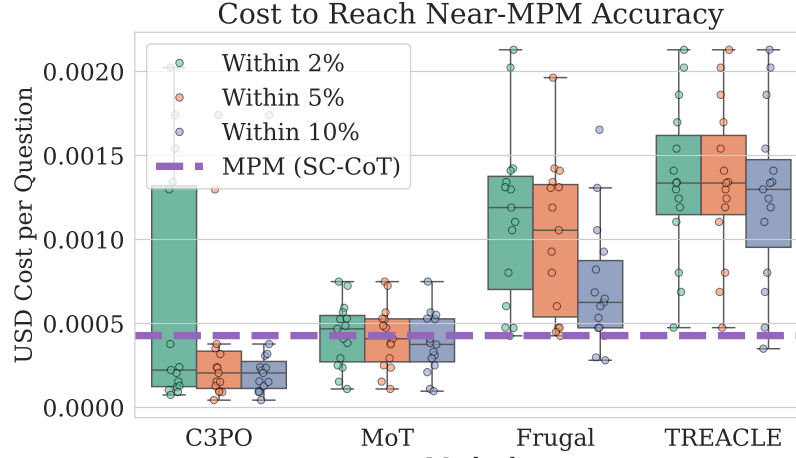


Figure 8: Surpassing existing cascade approaches tremendously, C3PO offers markedly superior cost-effectiveness across 16 benchmarks, requiring **less than 60% of the cost of the most powerful model (MPM)** (cost shown in purple) for an accuracy gap of at most 2, 5, and 10% using a QWEN cascade. In this boxplot, each dot represents a dataset and the whiskers extend to 90% coverage.

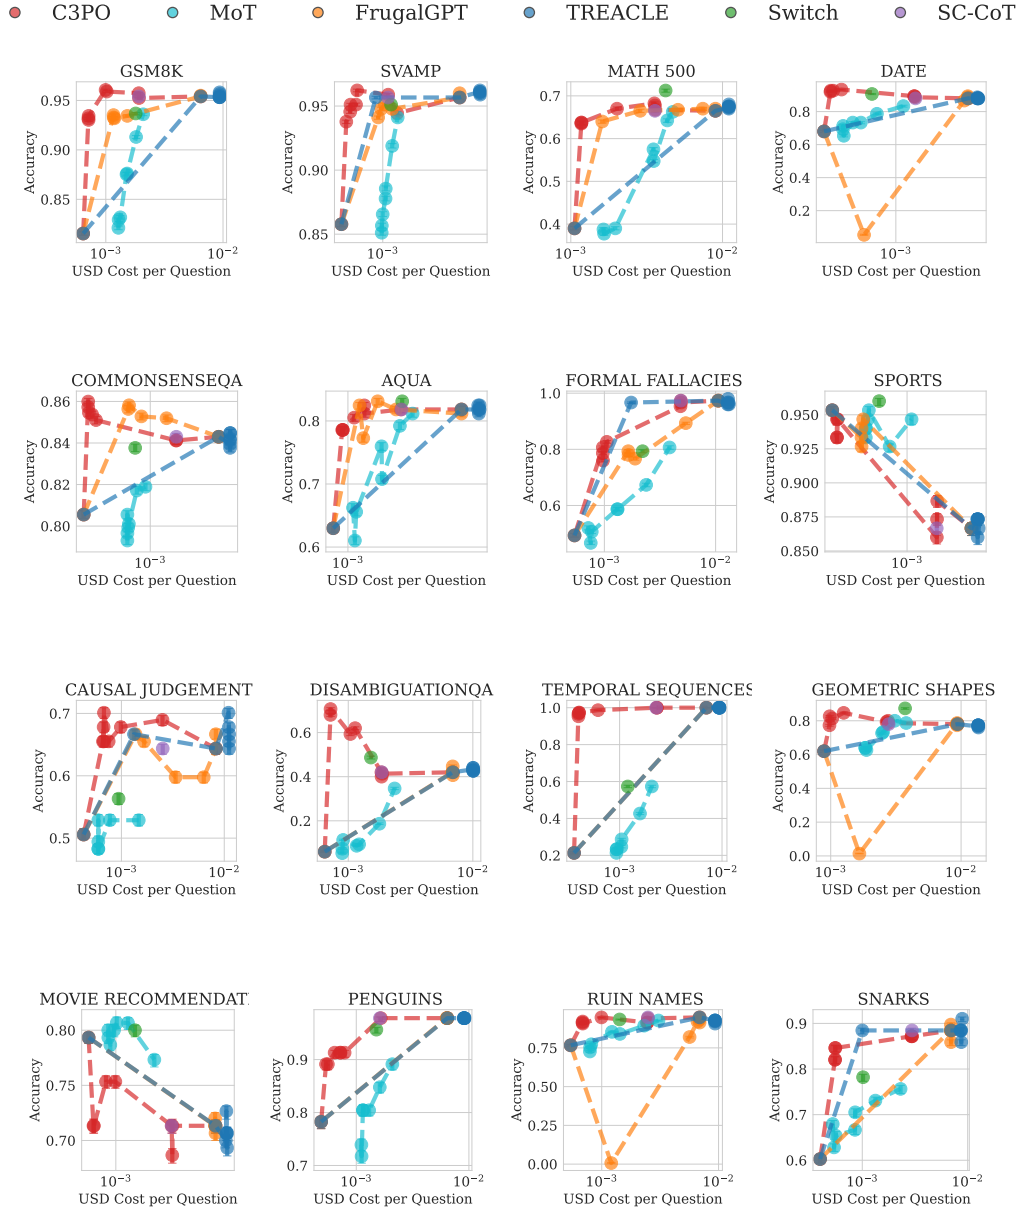


Figure 9: Accuracy vs. dollar cost of different algorithms for 16 datasets using the GPT cascade

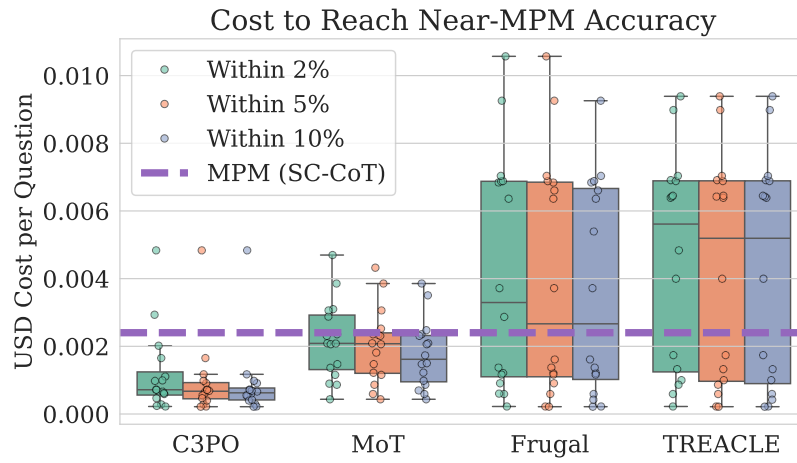


Figure 10: Surpassing existing cascade approaches tremendously, C3PO offers markedly superior cost-effectiveness across 16 benchmarks, requiring **less than 50% of the cost of the most powerful model (MPM)** (cost shown in purple) for an accuracy gap of at most 2, 5, and 10% using a GPT cascade. In this boxplot, each dot represents a dataset and the whiskers extend to 90% coverage.

785 **G.4 Comparison of C3PO and baselines in terms of performance and budget allocation for**
786 **questions with varying difficulties in MATH-500**

787 In Figure 3 in the main paper, we showed that using the LLAMA cascade, C3PO outperforms most
788 of the baselines at almost all difficulty levels with significantly lower inference costs. Similar results
789 are obtained for the QWEN and GPT cascades and are shown in Figure 11.

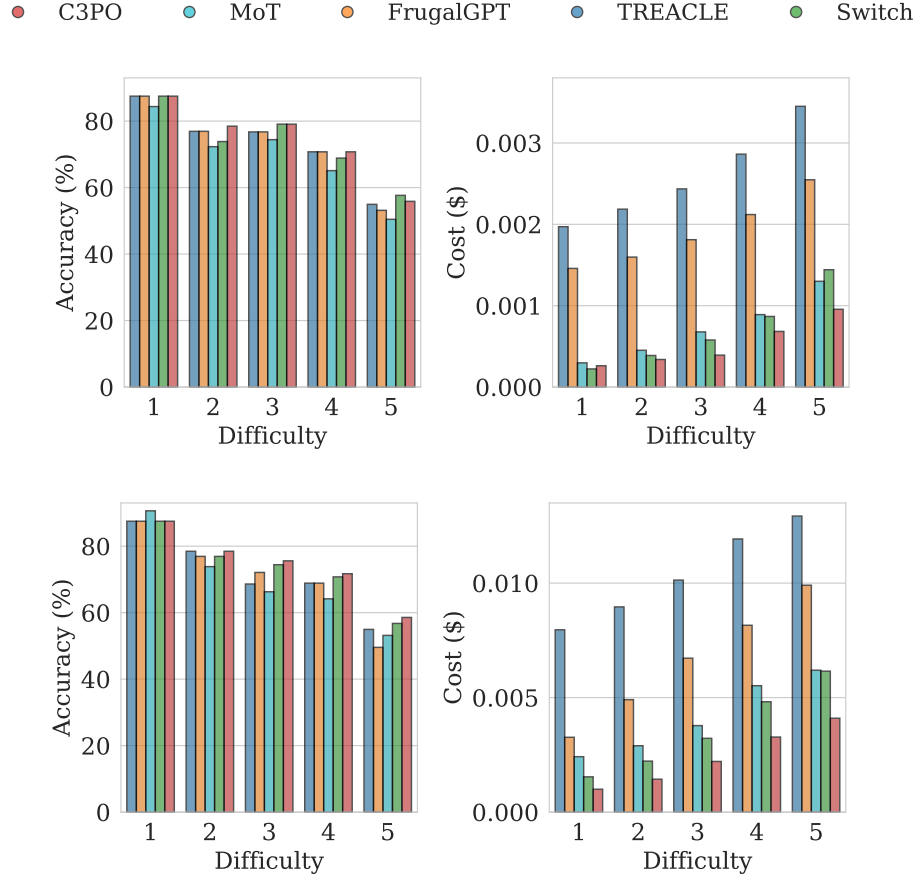


Figure 11: Average accuracies and dollar costs of different algorithms for five different difficulty levels of the MATH-500 dataset using the QWEN (top) and GPT (bottom) cascade.

H Few-shot Prompting Strategy and Prompt Examples

To ensure a fair comparison across all evaluated methods, we use a fixed few-shot prompting strategy that provides consistent formatting and reasoning guidance. Each model receives the same demonstration examples, with uniform instructions, answer formatting, and step-by-step rationales.

The prompt begins with a general instruction to use the boxed answer format (in LaTeX) and then presents a sequence of fully worked-out examples covering diverse mathematical topics, such as geometric series, completing the square, and factoring. This context is followed by the final question to be solved.

Below, we show an excerpt of the few-shot prompts used in our evaluations of free form answer datasets:

Free Form Answer Mathematical Question Prompt Template (MATH-500)

You are given a mathematical question. Solve the following problem and output correct result in a given format. Output answer at the end, use format $\boxed{\text{answer in LaTeX}}$. Think step-by-step.

Question: Kevin Kangaroo begins hopping on a number line at 0. He wants to get to 1, but he can hop only $\frac{1}{3}$ of the distance. Each hop tires him out so that he continues to hop $\frac{1}{3}$ of the remaining distance. How far has he hopped after five hops? Express your answer as a common fraction.

Let's think step by step. Kevin hops $\frac{1}{3}$ of the remaining distance with every hop. His first hop takes $\frac{1}{3}$ closer. For his second hop, he has $\frac{2}{3}$ left to travel, so he hops forward $(\frac{2}{3})(\frac{1}{3})$. For his third hop, he has $(\frac{2}{3})^2$ left to travel, so he hops forward $(\frac{2}{3})^2(\frac{1}{3})$. In general, Kevin hops forward $(\frac{2}{3})^{k-1}(\frac{1}{3})$ on his k th hop. We want to find how far he has hopped after five hops. This is a finite geometric series with first term $\frac{1}{3}$, common ratio $\frac{2}{3}$, and five terms. Thus, Kevin has hopped $\frac{\frac{1}{3}(1-(\frac{2}{3})^5)}{1-\frac{2}{3}} = \frac{211}{243}$. Thus, the answer is $\boxed{\frac{211}{243}}$.

Question: What is the area of the region defined by the equation $x^2 + y^2 - 7 = 4y - 14x + 3$?

Let's think step by step. We rewrite the equation as $x^2 + 14x + y^2 - 4y = 10$ and then complete the square, resulting in $(x+7)^2 - 49 + (y-2)^2 - 4 = 10$, or $(x+7)^2 + (y-2)^2 = 63$. This is the equation of a circle with center $(-7, 2)$ and radius $\sqrt{63}$, so the area of this region is $\pi r^2 = \boxed{63\pi}$.

Question: If $x^2 + y^2 = 1$, what is the largest possible value of $|x| + |y|$?

Let's think step by step. If (x, y) lies on the circle, so does $(x, -y)$, $(-x, -y)$, and $(-x, y)$, (which all give the same value of $|x| + |y|$), so we can assume that $x \geq 0$ and $y \geq 0$. Then $|x| + |y| = x + y$. Squaring, we get $(x+y)^2 = x^2 + 2xy + y^2 = 1 + 2xy$. Note that $(x-y)^2 \geq 0$. Expanding, we get $x^2 - 2xy + y^2 \geq 0$, so $2xy \leq x^2 + y^2 = 1$. Hence, $1 + 2xy \leq 2$, which means $x + y \leq \sqrt{2}$. Equality occurs when $x = y = \frac{1}{\sqrt{2}}$, so the maximum value of $|x| + |y|$ is $\boxed{\sqrt{2}}$.

Question: {{ question }}
Let's think step by step.

We also used datasets with multiple choice questions. For these datasets such as CommonSenseQA we used prompts of the style shown on the next page.

Multiple Choice Answer Prompt Template (CommonSenseQA)

You are given the following question. Solve it correctly and output a single letter (A, B, C, D, E or F) corresponding to the correct answer to the question. Output answer at the end, use format `\boxed{single letter (A, B, C, D, E)}`

`< singleletter(A,B,C,D,E) >`

Think step-by-step.

Question: What do people use to absorb extra ink from a fountain pen?

Answer Choices:

- A) shirt pocket
- B) calligrapher's hand
- C) inkwell
- D) desk drawer
- E) blotter

The answer must be an item that can absorb ink. Of the above choices, only blotters are used to absorb ink.

Thus, the answer is `\boxed{E}`.

Question: What home entertainment equipment requires cable?

Answer Choices:

- A) radio shack
- B) substation
- C) television
- D) cabinet

The answer must require cable. Of the above choices, only television requires cable.

It means the correct answer is `\boxed{C}`.

:

Question: Google Maps and other highway and street GPS services have replaced what?

Answer Choices:

- A) united states
- B) mexico
- C) countryside
- D) atlas

The answer must be something that used to do what Google Maps and GPS services do, which is to give directions. Of the above choices, only atlases are used to give directions.

The correct answer is `\boxed{D}`.

Question: Before getting a divorce, what did the wife feel who was doing all the work?

Answer Choices:

- A) harder
- B) anguish
- C) bitterness
- D) tears
- E) sadness

The answer should be the feeling of someone getting divorced who was doing all the work. Of the above choices, the closest feeling is bitterness.

Thus, the answer is `\boxed{C}`.

Question: `{{ question }}`

The answer

803

804 Shown above are the prompts for MATH-500 and CommonSenseQA datasets as examples of the
805 prompt templates for free form and multiple choice reasoning, respectively. All LLMs employed in
806 our study are prompted with the exact same few-shot demonstrations for evaluating all baselines and
807 our proposed method. This strategy ensures a fair evaluation setting, where the differences in empirical
808 performances between any pair of LLM-cascades can be solely attributed to the effectiveness of their
809 decision rules and crucially not to the utilization of unfair prompting advances.

810 I Case Study and Analysis

811 I.1 Question Level Case Study

812 This case study highlights a geometric reasoning problem where the goal is to compute the sum of
813 distances from the intersection points of a circle and hyperbola to a fixed external point. While smaller
814 models like LLaMA 1B and 3B produce inconsistent or incorrect answers, and TREACLE escalates
815 unnecessarily to the largest (and most expensive) model, our method (C3PO) successfully identifies
816 the smallest model (LLaMA 70B) that produces the correct answer. This demonstrates C3PO’s ability
817 to balance cost and accuracy by leveraging confidence-based early exits with probabilistic guarantees.

Case Study: Geometric Intersection Problem (MATH-500 Question #106)

Question: The graphs of $x^2 + y^2 + 6x - 24y + 72 = 0$ and $x^2 - y^2 + 6x + 16y - 46 = 0$ intersect at four points. Compute the sum of the distances from these four points to the point $(-3, 2)$.

Model Responses and Sample Rationales:

- **LLaMA 1B:** “8”, “6”, “18.67”, “18.24”, “ $\sqrt{7}$ ”
Sample rationale: To solve this problem, we need to find the intersection points and then calculate the distances from these points to the point $(-3, 2)$. [...] The intersection points are estimated and plugged into the distance formula. The model computes the distance and returns $\boxed{18.67}$, the answer is wrong due to numerical errors.
- **LLaMA 3B:** “75°”, “ $2\sqrt{41} + 18$ ”, “17”, “ $10\sqrt{3} + 2\sqrt{19}$ ”, “40”
Sample rationale: The model attempts trigonometric simplifications unrelated to the geometric problem, mistakenly treating it as an angle computation task. The conclusion $\boxed{75^\circ}$ is clearly off-topic, illustrating reasoning drift.
- **LLaMA 70B:** “40”, “ $16\sqrt{5}$ ”, “40”, “40”, “24”
Sample rationale: The model rewrites both equations into standard circle and hyperbola form, identifies the center $(-3, 12)$ and point $(-3, 2)$ as vertically aligned, and correctly computes total distance as $4 \times 10 = \boxed{40}$.
- **LLaMA 405B:** “40”, “36”, “40”, “40”, “40”
Sample rationale: Completing the square, the model identifies the circle centered at $(-3, 12)$ and hyperbola centered at $(-3, 8)$, recognizes symmetric geometry and uses translation and distance argument to correctly deduce total distance is $\boxed{40}$.

Cascade Behavior Analysis:

- **FrugalGPT** exits too early at **LLaMA 3B** and returns an incorrect answer.
- **TREACLE** escalates to the **LLaMA 405B** model and answers correctly, but at significantly higher cost.
- **C3PO (Ours)** exits at **LLaMA 70B**, which is the *cheapest* model to consistently return the correct answer (“40”).

Why C3PO wins: It optimally balances cost and accuracy. While weaker models return a mix of incorrect and inconsistent answers, and TREACLE incurs maximum cost, C3PO identifies the minimal model in the cascade capable of producing the correct answer.

818

I.2 Detailed study of C3PO’s cost efficiency

Although Figures 5, 7, and 9 demonstrate C3PO’s superior cost efficiency compared to baselines at the dataset level, they do not provide explicit and fine-grained information of C3PO’s cost allocation across different questions within each dataset.

Here, we conduct a detailed study to examine how C3PO spends its inference cost budget for questions where it succeeds in providing the correct answer and where it fails to do so. Specifically, for each baseline and LLM-cascade, we divide the test set into the following four non-overlapping categories; ‘very bad’, ‘bad’, ‘good’, and ‘very good’. A question falls into the ‘very bad’ category if C3PO fails to answer it correctly in spite of spending a higher inference cost compared to a baseline. If C3PO provides an incorrect answer to a question but incurs a lower cost compared to a baseline as well, then that question is ‘bad’. Clearly, having a ‘bad’ question is better than having a ‘very bad’ question in terms of efficient usage of inference cost. Similarly, ‘good’ questions refer to those test set examples, where C3PO provides correct responses at the expense of increased cost compared to a baseline. Finally, ‘very good’ questions are the ones, for which a correct answer from C3PO coincides with a lower cost compared to a baseline.

We report the distribution of these four categories of questions in the AQuA dataset for each LLM cascade and baselines in Figure 12. We observe that in comparison to most baselines, when C3PO provides an incorrect answer, it is more likely to save inference cost. Additionally, for the majority of the questions, where C3PO’s answer matches the true answer, it incurs reduced cost. This demonstrates the universal nature of C3PO’s cost effectiveness across different types of questions.

Similar results are obtained for other datasets and are included in Figures 13, 14, 15, 16.

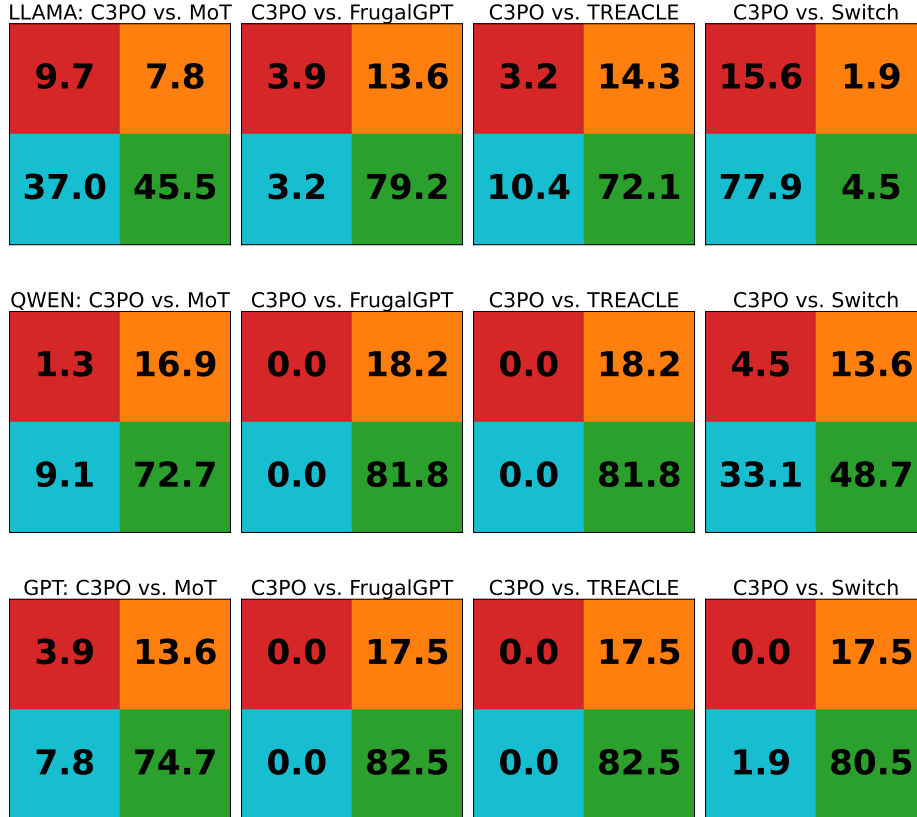


Figure 12: Percentage of ‘very bad’, ‘bad’, ‘good’, and ‘very good’ questions for each baseline and LLM cascade from the AQuA dataset

LLAMA: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
1.0 4.6	0.2 5.3	0.2 5.3	2.7 2.9
2.8 91.7	0.0 94.4	0.0 94.4	43.4 51.0

QWEN: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
2.4 2.1	0.0 4.6	0.3 4.2	2.8 1.8
47.0 48.4	0.0 95.4	4.2 91.2	86.8 8.7

GPT: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
0.2 4.4	0.0 4.7	0.0 4.7	0.1 4.6
0.3 95.0	0.0 95.3	0.0 95.3	0.3 95.0

Figure 13: Percentage of ‘very bad’, ‘bad’, ‘good’, and ‘very good’ questions for each baseline and LLM cascade from the SVAMP dataset. For each baseline and LLM-cascade, we divide the test set into the following four non-overlapping categories; ‘very bad’, ‘bad’, ‘good’, and ‘very good’. A question falls into the ‘very bad’ category if C3PO fails to answer it correctly in spite of spending a higher inference cost compared to a baseline. If C3PO provides an incorrect answer to a question but incurs a lower cost compared to a baseline as well, then that question is ‘bad’. Clearly, having a ‘bad’ question is better than having a ‘very bad’ question in terms of efficient usage of inference cost. Similarly, ‘good’ questions refer to those test set examples, where C3PO provides correct responses at the expense of increased cost compared to a baseline. Finally, ‘very good’ questions are the ones, for which a correct answer from C3PO coincides with a lower cost compared to a baseline. We observe that in comparison to most baselines, when C3PO provides an incorrect answer, it is more likely to save inference cost. Additionally, for the majority of the questions, where C3PO’s answer matches the true answer, it incurs reduced cost. This demonstrates the universal nature of C3PO’s cost effectiveness across different types of questions.

LLAMA: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
0.2 40.0	0.0 40.2	0.0 40.2	8.5 31.8
0.2 59.5	0.0 59.8	0.0 59.8	22.5 37.2

QWEN: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
10.0 19.5	0.0 29.5	0.0 29.5	11.0 18.5
9.5 61.0	0.0 70.5	0.0 70.5	33.2 37.2

GPT: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
6.2 26.8	0.0 33.0	0.0 33.0	13.0 20.0
17.5 49.5	0.0 67.0	0.2 66.8	39.8 27.2

Figure 14: Percentage of ‘very bad’, ‘bad’, ‘good’, and ‘very good’ questions for each baseline and LLM cascade from the **MATH-500** dataset. For each baseline and LLM-cascade, we divide the test set into the following four non-overlapping categories; ‘very bad’, ‘bad’, ‘good’, and ‘very good’. A question falls into the ‘very bad’ category if C3PO fails to answer it correctly in spite of spending a higher inference cost compared to a baseline. If C3PO provides an incorrect answer to a question but incurs a lower cost compared to a baseline as well, then that question is ‘bad’. Clearly, having a ‘bad’ question is better than having a ‘very bad’ question in terms of efficient usage of inference cost. Similarly, ‘good’ questions refer to those test set examples, where C3PO provides correct responses at the expense of increased cost compared to a baseline. Finally, ‘very good’ questions are the ones, for which a correct answer from C3PO coincides with a lower cost compared to a baseline. We observe that in comparison to most baselines, when C3PO provides an incorrect answer, it is more likely to save inference cost. Additionally, for the majority of the questions, where C3PO’s answer matches the true answer, it incurs reduced cost. This demonstrates the universal nature of C3PO’s cost effectiveness across different types of questions.

LLAMA: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
0.0 4.6	0.0 4.6	0.0 4.6	1.8 2.8
0.9 94.5	0.0 95.4	0.0 95.4	57.0 38.4

QWEN: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
0.6 5.8	0.0 6.4	0.0 6.4	2.5 3.9
10.7 82.9	0.0 93.6	0.0 93.6	47.6 46.0

GPT: C3PO vs. MoT	C3PO vs. FrugalGPT	C3PO vs. TREACLE	C3PO vs. Switch
1.7 3.0	0.0 4.7	0.0 4.7	2.5 2.2
57.9 37.4	0.0 95.3	0.0 95.3	70.5 24.9

Figure 15: Percentage of ‘very bad’, ‘bad’, ‘good’, and ‘very good’ questions for each baseline and LLM cascade from the **GSM8K** dataset. For each baseline and LLM-cascade, we divide the test set into the following four non-overlapping categories; ‘very bad’, ‘bad’, ‘good’, and ‘very good’. A question falls into the ‘very bad’ category if C3PO fails to answer it correctly in spite of spending a higher inference cost compared to a baseline. If C3PO provides an incorrect answer to a question but incurs a lower cost compared to a baseline as well, then that question is ‘bad’. Clearly, having a ‘bad’ question is better than having a ‘very bad’ question in terms of efficient usage of inference cost. Similarly, ‘good’ questions refer to those test set examples, where C3PO provides correct responses at the expense of increased cost compared to a baseline. Finally, ‘very good’ questions are the ones, for which a correct answer from C3PO coincides with a lower cost compared to a baseline. We observe that in comparison to most baselines, when C3PO provides an incorrect answer, it is more likely to save inference cost. Additionally, for the majority of the questions, where C3PO’s answer matches the true answer, it incurs reduced cost. This demonstrates the universal nature of C3PO’s cost effectiveness across different types of questions.

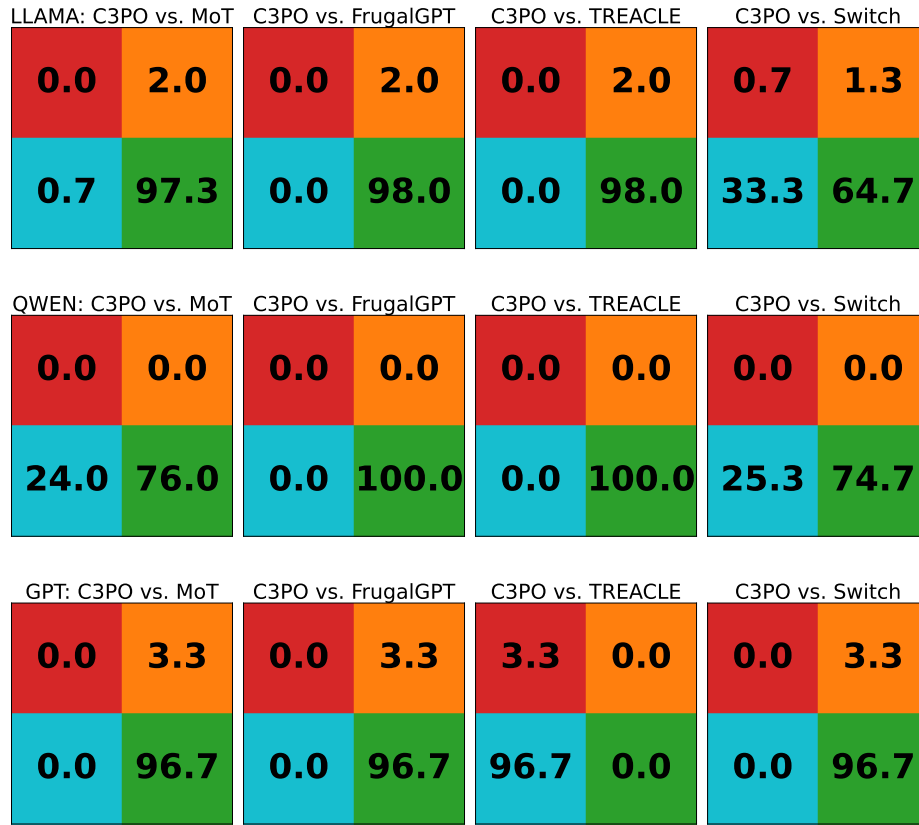


Figure 16: Percentage of ‘very bad’, ‘bad’, ‘good’, and ‘very good’ questions for each baseline and LLM cascade from the **BigBench Temporal Sequences** dataset.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: Every claim in the paper is supported with experiments in Section 5 and proofs in Supplementary Material.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: There is a discussion of limitations of our work in the Conclusion section.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We state the assumptions clearly in the proofs.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide all details required to reproduce our results: we describe our method in detail and provide a full pseudocode description of our approach. . In addition, the LLMs used in our work are all publicly accessible. We include our code for running our experiments in a zip file in the supplementary materials.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We make all experimental data and code available and commit to making it publicly available upon acceptance.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We follow standard splits and implementations from previous works. Where we make any changes we clearly state them.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We conduct repeated trials and provide 90% confidence interval error bars in our figures.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

889 Answer: [Yes]
890 Justification: We detail our computer hardware in the paper. Our method requires very little
891 compute. All LLM queries are performed by API calls.

892 **9. Code Of Ethics**
893 Question: Does the research conducted in the paper conform, in every respect, with the
894 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>
895 Answer: [Yes]
896 Justification: We have read and confirm that our work fully complies with the Code of
897 Ethics.

898 **10. Broader Impacts**
899 Question: Does the paper discuss both potential positive societal impacts and negative
900 societal impacts of the work performed?
901 Answer: [Yes]
902 Justification: We have an Appendix section dedicated to broader impacts.

903 **11. Safeguards**
904 Question: Does the paper describe safeguards that have been put in place for responsible
905 release of data or models that have a high risk for misuse (e.g., pretrained language models,
906 image generators, or scraped datasets)?
907 Answer: [NA]
908 Justification: We do not release new datasets or models.

909 **12. Licenses for existing assets**
910 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
911 the paper, properly credited and are the license and terms of use explicitly mentioned and
912 properly respected?
913 Answer: [Yes]
914 Justification: We include references to each of the benchmark datasets used in the paper, as
915 well as the LLMs used for the experiments. The licenses are listed for each dataset.

916 **13. New Assets**
917 Question: Are new assets introduced in the paper well documented and is the documentation
918 provided alongside the assets?
919 Answer: [Yes]
920 Justification: We do not introduce a new dataset, but our source code includes comments
921 and will be released under a MIT license upon acceptance.

922 **14. Crowdsourcing and Research with Human Subjects**
923 Question: For crowdsourcing experiments and research with human subjects, does the paper
924 include the full text of instructions given to participants and screenshots, if applicable, as
925 well as details about compensation (if any)?
926 Answer: [NA]
927 Justification: This work does not involve crowdsourcing or research with human subjects.

928 **15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human
929 Subjects**
930 Question: Does the paper describe potential risks incurred by study participants, whether
931 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
932 approvals (or an equivalent approval/review based on the requirements of your country or
933 institution) were obtained?
934 Answer: [NA]
935 Justification: This work does not require research with human subjects.

936 **16. Declaration of LLM usage**

937 Question: Does the paper describe the usage of LLMs if it is an important, original, or
938 non-standard component of the core methods in this research? Note that if the LLM is used
939 only for writing, editing, or formatting purposes and does not impact the core methodology,
940 scientific rigorousness, or originality of the research, declaration is not required.

941 Answer: [NA]

942 Justification: The core method development in this research does not involve LLMs as any
943 important, original, or non-standard components.

944