
AdaCoT: Pareto-Optimal Adaptive Chain-of-Thought Triggering via Reinforcement Learning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Large Language Models (LLMs) have demonstrated remarkable capabilities but
2 often face challenges with tasks requiring sophisticated reasoning. While Chain-
3 of-Thought (CoT) prompting significantly enhances reasoning, it indiscriminately
4 generates lengthy reasoning steps for all queries, leading to substantial compu-
5 tational costs and inefficiency, especially for simpler inputs. To address this
6 critical issue, we introduce AdaCoT (Adaptive Chain-of-Thought), a novel frame-
7 work enabling LLMs to adaptively decide when to invoke CoT. AdaCoT framed
8 adaptive reasoning as a Pareto optimization problem that seeks to balance model
9 performance with the costs associated with CoT invocation (both frequency and
10 computational overhead). We propose a reinforcement learning (RL) based method,
11 specifically utilizing Proximal Policy Optimization (PPO), to dynamically control
12 the CoT triggering decision boundary by adjusting penalty coefficients, thereby
13 allowing the model to determine CoT necessity based on implicit query complex-
14 ity. A key technical contribution is Selective Loss Masking (SLM), designed to
15 counteract decision boundary collapse during multi-stage RL training, ensuring
16 robust and stable adaptive triggering. Experimental results demonstrate that Ada-
17 CoT successfully navigates the Pareto frontier, achieving substantial reductions
18 in CoT usage for queries not requiring elaborate reasoning. For instance, on our
19 production traffic testset, AdaCoT reduced CoT triggering rates to as low as 3.18%
20 and decreased average response tokens by 69.06%, while maintaining high per-
21 formance on complex tasks. This substantial token decrease directly translates to a
22 significant reduction in inference computational load. AdaCoT pioneers adaptive
23 CoT triggering, offering a practical and principled solution for developing more
24 efficient, responsive, and cost-effective LLMs, particularly crucial for interactive
25 and resource-sensitive applications.

26 1 Introduction

27 Large Language Models (LLMs) have garnered substantial attention due to their remarkable ability
28 to encode extensive world knowledge from vast corpora [1], enabling impressive performance
29 across diverse tasks such as question answering, creative writing, and summarization. Despite
30 these successes, LLMs often demonstrate limitations in tasks requiring sophisticated reasoning,
31 including solving complex mathematical problems and intricate coding puzzles. To mitigate this,
32 recent methodologies have employed Chain-of-Thought (CoT) prompting [2], which encourages
33 models to explicitly generate step-by-step reasoning prior to producing final answers. This approach
34 significantly enhances the reasoning capability of models, even achieving human-expert levels in
35 certain domains [3, 4, 5, 6].

36 However, employing CoT prompting also poses critical challenges during inference [7]. Specifically,
37 it substantially increases the number of tokens generated, even for simple queries that do not benefit
38 from elaborate reasoning, such as straightforward arithmetic questions. This indiscriminate token
39 expense consequently raises deployment costs and reduces inference efficiency. Ideally, a model
40 should adaptively determine when detailed reasoning is necessary. For instance, simple queries like
41 "What is 1+1?" should be answered immediately without additional reasoning steps, whereas more
42 complex queries require deeper and step-by-step reasoning. An adaptive strategy would thus optimize
43 token usage, balancing cost-efficiency with response quality.

44 Recently, a few efforts has made attempts towards this direction. These approaches can be broadly
45 categorized into three main directions. (1) Incorporate length penalties or brevity rewards during
46 the reinforcement learning (RL) stage to encourage shorter, more concise reasoning paths [8, 9, 10,
47 11, 12, 13]. (2) Restructure CoT outputs through post-processing such as iterative summarization
48 or pruning [14, 15, 16, 17, 18, 19, 20, 21]. (3) Employ explicit user instructions or hand-crafted
49 selection mechanisms to control the use of CoT [22, 23, 24, 25, 26, 27]. Despite their contributions,
50 they mainly focus on monotonic reasoning reduction, failing to account for the nuanced variability
51 in query complexity, i.e., treating simple and difficult prompts adaptively. Moreover, they lack a
52 principled optimization framework to guide balancing response quality against deployment cost.

53 To address these limitations, we introduce AdaCoT (Adaptive Chain-of-Thought), a novel approach
54 grounded in formal mathematical analysis. Our key insight is framing adaptive reasoning as a
55 multi-objective optimization problem with two competing goals: maximizing response accuracy and
56 minimizing deployment costs. Specifically, we formalize this balance through Pareto optimization,
57 seeking optimal trade-offs between reasoning complexity and inference efficiency. Such a mathemati-
58 cal framework provides clear theoretical grounding for dynamically adapting CoT triggering based
59 on query complexity.

60 Leveraging this formalization, we propose an RL-based strategy explicitly designed around the Pareto
61 optimization framework, enabling effective control of the model's decision boundary for initiating
62 CoT prompting. During training, the RL agent dynamically assesses the complexity of incoming
63 user queries to determine the necessity and extent of reasoning steps. By carefully designing the
64 reward function to incorporate penalty coefficients, we encourage the RL agent to seek solutions
65 along the Pareto frontier, explicitly optimizing trade-offs between response accuracy and token
66 expenditure. This structured exploration enables the model to effectively discern when detailed
67 reasoning is beneficial, thereby systematically enhancing inference efficiency and significantly
68 reducing deployment costs.

69 The proposed AdaCoT framework delivers substantial benefits in LLM operational efficiency. By
70 empowering models to selectively engage CoT, AdaCoT can reduce triggering rates to as low as
71 3.18% and cut average response tokens by 69.1% in production settings. This significant reduction
72 in computational load is achieved while maintaining strong performance on 15 widely-adopted
73 benchmarks. For example, AdaCoT can achieve a 62.8% average score using only a 53.3% CoT rate,
74 closely rivaling the 65.0% score of a model that always employs CoT. These improvements directly
75 translate to more cost-effective and responsive LLM systems.

76 2 The AdaCoT Framework

77 Enabling a large language model (LLM) to dynamically decide whether to invoke Chain-of-Thought
78 (CoT) reasoning based on the complexity of user queries is a critical task, which allows LLMs to
79 allocate computational resources more rationally, i.e., spending tokens on complex reasoning tasks
80 while avoiding unnecessary overhead for simple queries.

81 To achieve this, we introduce AdaCoT, a unified framework for adaptive reasoning. The central
82 insight is that the decision to employ CoT prompting can be naturally cast as an optimization problem
83 involving two competing goals: (1) maximizing response accuracy and (2) minimizing deployment
84 costs. To capture this trade-off, we frame the task as a Pareto optimization problem, providing a
85 principled foundation for balancing reasoning quality against computational efficiency. Based on this
86 formulation, we propose an RL-based control strategy that governs the invocation of CoT reasoning.
87 Specifically, we train a policy model learning to assess the complexity of each query and decide
88 whether CoT reasoning should be used. During RL training, the policy model is optimized using a
89 reward signal carefully designed to reflect the Pareto trade-off, incorporating both task performance

90 and token efficiency. Through this RL-based mechanism, AdaCoT learns to allocate reasoning effort
 91 adaptively, yielding high-quality responses with minimal computational cost.

92 2.1 Adaptive Reasoning as a Pareto Optimization Problem

93 We formulate the adaptive reasoning challenge as a Pareto optimization problem, aiming to simultaneously
 94 maximize model performance and minimize CoT usage. Let $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ be a dataset
 95 of query-response pairs, where x_i is the input query and y_i is the ground truth response. Let f_θ be an
 96 LLM parameterized by θ , and let $r_\theta(x_i)$ be the response generated by the model for input x_i .

97 To CoT usage is measured by the CoT triggering rate $T(\theta)$, defined as the proportion of responses
 98 that include reasoning:

$$99 T(\theta) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[\text{HasReasoning}(r_\theta(x_i))] \quad (1)$$

100 where $\mathbf{1}[\cdot]$ is the indicator function and $\text{HasReasoning}(\cdot)$ determines if a response contains explicit
 101 CoT steps (e.g., non-empty content within `<think>...</think>` tags).

102 On the other hand, model performance $P(\theta)$ is defined as the average score on a set of evaluation
 103 metrics:

$$104 P(\theta) = \frac{1}{M} \sum_{j=1}^M \text{Score}_j(\theta) \quad (2)$$

105 where M is the number of evaluation instances/metrics and $\text{Score}_j(\theta)$ is the model's score on the
 106 j -th evaluation.

107 Putting CoT usage $T(\theta)$ and model performance $P(\theta)$ together, the objective is to find model
 108 parameters θ^* that achieve an optimal trade-off:

$$109 \theta^* = \arg \max_{\theta} \{\lambda_P \cdot P(\theta) - \lambda_T \cdot T(\theta)\} \quad (3)$$

110 or more generally, to find solutions on the Pareto frontier of $(P(\theta), 1 - T(\theta))$. The hyperparameters
 111 λ_P and λ_T (or their implicit ratio) control the relative importance of performance versus CoT
 112 reduction. AdaCoT focuses on providing mechanisms to explore this frontier.

113 2.2 Training Pipeline for AdaCoT

114 The AdaCoT training pipeline integrates supervised fine-tuning (SFT) as an initialization phase,
 115 followed by multi-stage reinforcement learning (RL) to refine the adaptive CoT triggering behavior.

116 2.2.1 Data Preparation and Supervised Fine-Tuning (SFT) as Warm-up

117 To provide the model with an initial understanding of when CoT might be beneficial, we perform
 118 a data preparation stage. This is achieved by leveraging an auxiliary model, guided by a set of
 119 predefined principles (e.g., query complexity, expected reasoning depth, domain; see Appendix B). In
 120 our implementation, we use an internal 15B-parameter model to generate these annotations; however,
 121 the framework is model-agnostic and can be instantiated using any sufficiently capable LLM with
 122 basic instruction-following abilities. Queries are labeled as either likely benefiting from CoT or
 123 likely suitable for a direct answer. This principled, automated labeling process is more consistent and
 124 scalable than manual annotation.

125 The SFT training data is then structured as follows: For queries labeled as benefiting from CoT,
 126 responses retain the full reasoning process: `<think>reasoning_steps</think>answer`. For
 127 queries labeled as not requiring CoT, responses omit explicit reasoning but maintain structural
 128 consistency: `<think></think>answer`. This SFT stage serves as a "warm-up", equipping the
 129 model with a foundational capability to distinguish between these two response styles. The consistent
 130 use of `<think></think>` tags is crucial for maintaining response format integrity.

131 2.2.2 Reinforcement Learning for Adaptive CoT Control

132 The RL stage is pivotal for fine-tuning AdaCoT's adaptive reasoning capabilities. We design a reward
 133 function $R(x, r)$ for an input query x and generated response r :

$$134 R(x, r) = R_{\text{base}}(x, r) - \alpha_1 \cdot P_{\text{miss}}(x, r) - \alpha_2 \cdot P_{\text{over}}(x, r) - \gamma \cdot P_{\text{fml}}(r) \quad (4)$$

131 where $R_{\text{base}}(x, r)$ is the base reward reflecting response quality, $P_{\text{miss}}(x, r)$ is a binary penalty for
132 reasoning omission, $P_{\text{over}}(x, r)$ is a binary penalty for reasoning overuse, $P_{\text{fmt}}(r)$ is a binary penalty
133 for format errors, and $\alpha_1, \alpha_2, \gamma$ are non-negative penalty coefficients. By adjusting α_1 and α_2 , we
134 steer AdaCoT towards different CoT triggering decision boundaries, allowing exploration of the
135 Pareto frontier.

136 2.2.3 Addressing Decision Boundary Collapse with Selective Loss Masking

137 A significant challenge in multi-stage RL, particularly when fine-tuning on specialized datasets with
138 skewed CoT distributions (e.g., mathematical datasets where CoT is almost always beneficial), is the risk of the adaptive CoT triggering capability becoming unstable or collapsing. The model might
139 revert to a homogeneous behavior, either always or never triggering CoT, thereby losing the nuanced
140 decision-making learned in earlier, more balanced training stages. We term this phenomenon decision
141 boundary collapse. This is particularly problematic if the final RL stage has significant bias, as it can
142 lead to the model almost completely losing its adaptive triggering capability.

144 To address decision boundary collapse, AdaCoT incorporates Selective Loss Masking (**SLM**). SLM
145 aims to preserve the CoT triggering ratio and distribution established during SFT or prior RL stages. It
146 achieves this by selectively masking the loss contribution from the pivotal "decision token" during RL
147 phases prone to distribution bias. This decision token is defined as the token immediately succeeding
148 the `<think>` tag.

149 The modified policy gradient loss under SLM, \mathcal{L}_{SLM} , is computed by excluding the loss component
150 associated with this decision token:

$$\mathcal{L}_{\text{SLM}} = \sum_{k \neq k_{\text{decision}}} \ell_k \quad (5)$$

151 where ℓ_k is the original loss component for the k -th token, and k_{decision} is the index of the decision
152 token.

153 3 Experiments

154 We conducted extensive experiments to evaluate the AdaCoT framework, focusing on its ability to
155 navigate the performance-cost trade-off, the effectiveness of its adaptive triggering mechanism, and
156 its impact on inference efficiency. This section details our experimental setup, presents the main
157 results, and analyzes the findings.

158 3.1 Experimental Setup

159 For our base model, we utilized our internal 15B/150B parameter Mixture-of-Experts (MoE) [28, 29]
160 model. The AdaCoT post-training process comprised an initial SFT stage, followed by a two-stage
161 RL procedure: first, a Mathematics-Focused RL stage (RL-Math) concentrated on complex, rule-
162 verifiable problems, and second, a General Domain RL stage (RL-General) which incorporated
163 broader data and a trained reward model. We compared our **AdaCoT RL Models** (Exp1-Exp4)
164 against several baselines: a **Full CoT SFT Baseline** (SFT model always generating CoT), a **Full**
165 **CoT RL Baseline** (RL model derived from the Full CoT SFT, always generating CoT), a **No CoT**
166 **SFT Baseline** (SFT model never generating CoT), a **No CoT RL Baseline** (RL model derived from
167 the No CoT SFT, never generating CoT), and our **AdaCoT SFT Model** (our model after only the
168 SFT stage, also referred to as Adaptive SFT Model).

169 The SFT and RL training datasets were constructed to cover a diverse range of domains, including
170 mathematics, reasoning, professional disciplines (e.g., law, medicine), dialogue, creative writing, and
171 general knowledge question answering. Both SFT and RL data were labeled for CoT necessity using
172 the principle-guided assessment detailed in Appendix B. In the SFT dataset, approximately 67% of
173 the samples were labeled as requiring CoT, while in the RL dataset, this proportion was around 40%.
174 During SFT, queries identified as not requiring CoT were formatted with empty `<think></think>`
175 tags. In the RL-Math stage, which is particularly prone to decision boundary collapse, we employed
176 Selective Loss Masking (SLM), as described in Section 2.2.3. For the RL-General stage, we applied
177 penalties according to Equation 4, systematically varying the α_1 and α_2 coefficients to explore
178 different points on the Pareto frontier. Proximal Policy Optimization (PPO) [30] was used for all RL
179 policy updates.

180 For evaluation, we used 15 diverse open-source benchmark datasets to assess overall performance,
181 measured by the average score. To balance internal iteration efficiency with evaluation accuracy, some
182 of these datasets underwent up-sampling or down-sampling, or the number of inference runs per test
183 sample was adjusted (with the final score being an average over multiple inferences). These datasets
184 include LiveBench [31], MMLU Pro [32], SuperGPQA [33], GPQA [34], Chinese SimpleQA [35],
185 SimpleQA [36], AIME24 & AIME25, MATH [37], OlympiadBench [38], SweBench Agentless [39],
186 LiveCodeBench [40], KOR-Bench [41], ProcBench [42], and SysBench [43]. The detailed per-
187 dataset scores, which form the basis for our average score calculations, are presented in Appendix A
188 (Table 4). To specifically assess CoT triggering decisions on typical user queries, we curated a high-
189 quality balanced test set of 1000 prompts. These prompts were labeled for CoT necessity using the
190 same principle-guided assessment as our SFT/RL training data and subsequently underwent manual
191 verification to ensure label accuracy. On this set, we report CoT Triggering Accuracy, F1-score,
192 Precision, and Recall, where the positive class indicates that CoT is required. Other metrics included
193 the CoT triggering rate on the benchmark datasets and the average response token num on production
194 setting.

195 3.2 Results and Analysis

196 Our results demonstrate AdaCoT’s ability to effectively control CoT invocation, leading to improved
197 efficiency while maintaining strong performance.

198 3.2.1 Pareto Frontier Analysis

199 We trained four variants of our AdaCoT RL model (Exp1-Exp4) by varying the penalty coefficients α_1
200 (for missing CoT) and α_2 (for overusing CoT). The specific coefficients were: Exp1 ($\alpha_1 = 0.1, \alpha_2 =$
201 0.3), Exp2 ($\alpha_1 = 0.2, \alpha_2 = 0.3$), Exp3 ($\alpha_1 = 0.3, \alpha_2 = 0.3$), and Exp4 ($\alpha_1 = 0.3, \alpha_2 = 0.1$).
202 The format error penalty γ was consistently set to 1.0. Figure 1 illustrates the average score plotted
203 against the CoT triggering rate for these models and the baselines, based on the average performance
204 across our 15 benchmark datasets (detailed in Appendix Table 4). The No CoT SFT baseline achieved
205 an average score of 43.6% with 0% CoT usage, while the No CoT RL baseline improved this to
206 47.7% at 0% CoT. The AdaCoT SFT Model (Adaptive SFT) registered a 57.1% average score at a
207 61.3% CoT rate.

208 The AdaCoT RL models trace a compelling Pareto frontier. AdaCoT RL Exp1 (43.1% CoT, 59.7%
209 score) and AdaCoT RL Exp2 (53.3% CoT, 62.8% score) demonstrate significant performance gains
210 over the AdaCoT SFT model while operating at lower or comparable CoT rates. Notably, AdaCoT
211 RL Exp2 achieves a 62.8% average score, approaching the Full CoT RL baseline (65.0% score,
212 100% CoT) with nearly half the CoT usage. As we increase the CoT triggering rate, AdaCoT RL
213 Exp3 (65.4% CoT, 64.3% score) and AdaCoT RL Exp4 (67.7% CoT, 64.4% score) further push
214 performance, closely rivaling the Full CoT RL baseline’s score but with approximately 32-35% less
215 CoT invocation. Moreover, these results indicate that despite the fixed CoT triggering labels within
216 the SFT/RL data, adjusting the combination of penalty coefficients during the RL phase enables the
217 final RL model to learn triggering strategies that transcend these initial labeling constraints.

218 This highlights AdaCoT’s effectiveness in navigating the trade-off between performance and CoT cost.
219 However, it is also observable that while the AdaCoT RL models achieve substantial efficiency gains
220 and define a superior Pareto curve compared to simpler baselines, they encounter a slight performance
221 bottleneck when their triggering rates are pushed higher. Specifically, even the highest-performing
222 adaptive models (Exp3 and Exp4, with scores of 64.3% and 64.4% respectively) do not surpass the
223 average score of the Full CoT RL baseline (65.0%). This suggests that while AdaCoT excels at
224 reducing CoT for a vast majority of queries without compromising much on average performance,
225 and indeed offers a better score-to-cost ratio, the absolute peak average performance achieved by
226 a model specialized to always use CoT (Full CoT RL) remains marginally higher. This indicates
227 that the adaptive mechanism, while highly effective, might not perfectly replicate or exceed the
228 performance of an always-on CoT strategy in every single scenario contributing to the average, thus
229 not fully crossing this specific optimal boundary for maximum average score. This could be due to
230 the inherent complexities of learning a universally optimal triggering heuristic or the RL optimization
231 finding a balance that prioritizes the significant cost savings available across the query distribution.

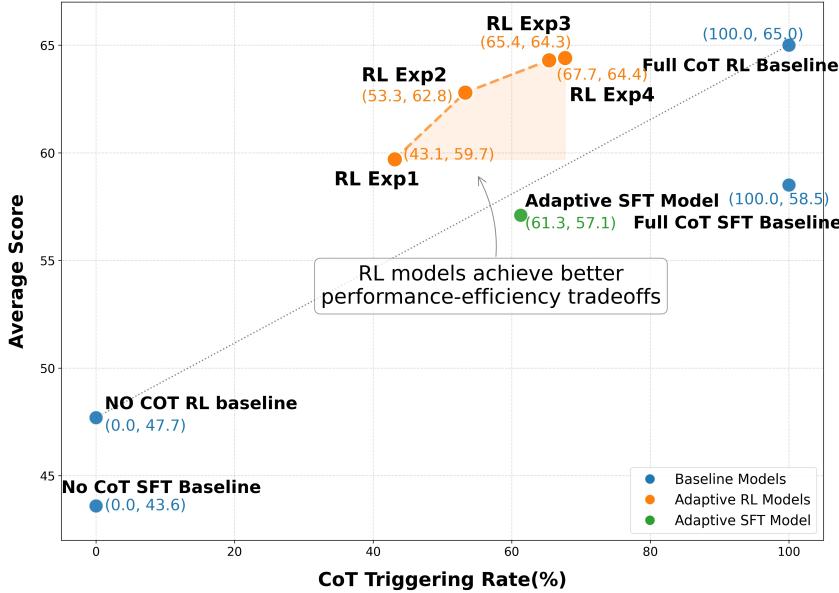


Figure 1: Average Score vs. CoT Triggering Rate across 15 widely-adopted benchmarks. Blue points represent baseline models. The green point is the AdaCoT SFT model. Orange points represent AdaCoT RL models trained with varying penalty coefficients, forming an improved Pareto frontier (indicated by the orange dashed line and shaded region) over the baselines. The dotted line connects the No CoT RL baselines to the Full CoT RL baseline, illustrating a simpler trade-off curve.

232 3.2.2 Adaptive CoT Triggering Performance and Ablation Studies on Daily-Use Queries

233 We evaluated the CoT triggering capabilities of AdaCoT at various training stages using our curated
 234 1000-prompt daily-use test set. Table 1 presents these results, which include an ablation study for
 235 SLM and an assessment of the meta-reasoning strategy (discussed further in Section 4.2) at the SFT
 236 stage.

Table 1: CoT triggering performance on the 1000 daily-use prompt test set across different AdaCoT stages and configurations (positive class: requires CoT). RL-Math is the Mathematics-Focused RL stage; RL-General refers to the final models (Exp1-4).

Model Stage / Variant	Accuracy	F1-Score	Recall	Precision
AdaCoT SFT Model	0.795	0.750	0.616	0.959
AdaCoT RL Model (Exp1 - RL-General)	0.657	0.484	0.322	0.975
AdaCoT RL Model (Exp2 - RL-General)	0.816	0.814	0.804	0.823
AdaCoT RL Model (Exp3 - RL-General)	0.809	0.789	0.716	0.879
AdaCoT RL Model (Exp4 - RL-General)	0.678	0.535	0.370	0.963
RL-Math (without SLM)	0.506	0.669	1.000	0.503
RL-Math (with SLM)	0.813	0.781	0.670	0.938
AdaCoT SFT Model (with Meta-Reasoning)	0.858	0.840	0.762	0.935

237 The AdaCoT SFT model itself provides a strong baseline for adaptive triggering. The results clearly
 238 show that the RL-Math stage without SLM suffers from severe decision boundary collapse: the
 239 model defaults to triggering CoT (Recall=1.0) but with poor precision (0.503) and consequently low
 240 overall accuracy (0.506). Applying SLM during the RL-Math stage effectively preserves the adaptive
 241 capability learned during SFT, maintaining high precision (0.938) and achieving significantly better
 242 accuracy (0.813). The final AdaCoT RL models (Exp1-4, emerging from the RL-General stage)
 243 demonstrate how adjusting the RL penalty coefficients (α_1, α_2) allows for fine-tuning of the decision

244 boundary. AdaCoT RL Model Exp2, for example, achieves a well-balanced F1-score of 0.814. The
245 incorporation of a meta-reasoning strategy at the SFT stage also shows a notable improvement in
246 triggering performance, a point elaborated in Section 4.2.

247 **3.2.3 Response Length Reduction and Efficiency Gains**

248 The adaptive reasoning enabled by AdaCoT translates into significant reductions in computational
249 costs. Table 2 shows the average response length and CoT triggering rates for AdaCoT RL Model
250 Exp2 (selected for its balanced performance on the daily-use set and strong average benchmark
251 performance) when applied to our production traffic testset, which reflects natural, unfiltered user
252 query distributions.

Table 2: Average response token num (with reduction noted) and CoT triggering rate on production traffic testset for AdaCoT RL Model Exp2 vs. Full CoT RL Baseline.

Platform	Model / Mode	Avg. Response Tokens	CoT Triggering Rate
Mobile	Full CoT RL Baseline	377.18	100.00%
	AdaCoT RL Model Exp2 (Adaptive)	116.70 (↓69.1%)	3.18%
PC	Full CoT RL Baseline	1376.31	100.00%
	AdaCoT RL Model Exp2 (Adaptive)	405.25 (↓70.6%)	12.50%

253 As evidenced, AdaCoT RL Model Exp2 achieves very low CoT triggering rates in a production
254 setting (3.18% on mobile devices, 12.50% on PCs). This dramatic reduction from the 100% CoT
255 usage of a non-adaptive model translates directly into substantial computational savings.

256 **4 Discussion and Future Work**

257 **4.1 Design Considerations and Limitations**

258 AdaCoT offers a pragmatic approach to adaptive reasoning by combining principle-guided initial data
259 labeling with RL-based optimization of the CoT decision boundary. This methodology was chosen
260 to circumvent inherent challenges in purely autonomous CoT trigger learning, such as information
261 asymmetry in assessing counterfactual benefits and the difficulty of quantifying quality degradation
262 from CoT omission, particularly for subjective tasks.

263 While AdaCoT is a promising initial step, several limitations exist. The optimal CoT triggering
264 strategy is relative to the base model’s capabilities, necessitating recalibration for different LLMs.
265 Our current binary CoT invocation (on/off) simplifies a continuous spectrum of reasoning depths and
266 styles, potentially limiting nuance. Domain generalization remains a challenge, as CoT necessity can
267 vary significantly across knowledge areas, and the framework currently lacks personalization for user
268 verbosity preferences. Moreover, the initial principle-guided labeling requires continuous refinement.
269 Our Pareto analysis (Section 3.2.1) also indicates that while AdaCoT significantly improves efficiency
270 and nears the peak average performance of specialized always-on CoT models, a small performance
271 gap persists, highlighting the difficulty for adaptive mechanisms to achieve absolute maximum
272 performance across all query types.

273 Acknowledging the limitations of the current framework, we anticipate that future research will offer
274 valuable critiques and further refine these initial explorations. Areas warranting deeper investigation
275 include more granular control over reasoning, such as adaptive reasoning length where models
276 dynamically adjust verbosity, or more nuanced triggering mechanisms beyond a simple binary
277 decision. We believe such continued efforts by the community are crucial for developing more
278 sophisticated and efficient reasoning strategies, potentially addressing the observed performance gap
279 while maximizing efficiency and enhancing nuanced control over LLM reasoning.

280 **4.2 Meta-Reasoning for Reasoning Decisions**

281 We explored an innovative "meta-reasoning" approach during the SFT stage to enhance AdaCoT’s
282 CoT triggering decisions. This involved the model first generating a brief internal assessment of
283 the input query’s complexity before deciding whether to proceed with full CoT, as illustrated by

284 the response structures in Figure 2. Incorporating this strategy into the AdaCoT SFT model led
 285 to a notable improvement in CoT triggering performance on our daily-use test set: the F1-score
 286 increased from 0.750 to 0.840 (Table 1). This result suggests that explicitly prompting the model
 287 to first assess query complexity can significantly enhance its subsequent decision-making regarding
 288 CoT invocation.

Response Format with Meta-Reasoning	
Straightforward Query	<pre><think>This is a straightforward question about X, I think I can answer directly. [empty]</think> {answer}</pre>
Complex Query	<pre><think>This is a relatively complex question about Y, I need to think carefully. [formal reasoning process]</think> {answer}</pre>

Figure 2: Example response structure incorporating explicit meta-reasoning for CoT decisions.

299 An interesting and serendipitous discovery with the meta-reasoning SFT model was an emergent
 300 capability for user-prompt controllability over CoT. Users could, to some extent, influence whether
 301 the model engaged in CoT by including explicit cues in their prompts about the desired level of
 302 reasoning (e.g., "please think step-by-step" to encourage CoT, or "give a direct answer" to discourage
 303 it). While this controllability was not perfectly accurate across all scenarios, it points towards
 304 a promising avenue for developing more interactive and user-guided reasoning systems (further
 305 illustrative examples are provided in Appendix D).

306 Despite these benefits, the explicit meta-reasoning step inherently increases the number of tokens
 307 generated for every query, as the model first articulates its complexity assessment. Considering
 308 the very low CoT triggering rates observed for AdaCoT in production environments (e.g., 3.18%
 309 on mobile traffic for AdaCoT RL Model Exp2, as shown in Table 2), the cumulative token cost of
 310 these additional meta-reasoning steps would become substantial. Therefore, while acknowledging
 311 its potential for improving decision accuracy and enabling user control, we did not adopt this
 312 explicit meta-reasoning as the default for subsequent RL experiments due to this efficiency trade-off.
 313 Nevertheless, it highlights an important direction for future research, which might explore more
 314 token-efficient methods for incorporating such meta-reasoning, perhaps through implicit learning
 315 mechanisms or parallel processing of the complexity assessment.

316 4.3 Preserving Peak Performance with AdaCoT

317 A critical question is whether adaptive reasoning limits a model's maximum performance. We
 318 investigated this using System Prompts (SPs), integrated into AdaCoT's SFT and RL training to
 319 control reasoning behavior (e.g., "Always Reason SP," "Never Reason SP"). During SFT, a small
 320 portion of data was augmented with SPs, and target responses were modified for compliance. In
 321 RL, a fraction of training prompts included SPs, with rewards penalizing deviations from explicit SP
 322 instructions, ensuring robust adherence (details in Appendix C). Our focus here is using the "Always
 323 Reason SP" to assess AdaCoT's performance ceiling.

324 Instructing AdaCoT RL models to always generate CoT via this SP allowed direct comparison against
 325 the Full CoT RL Baseline on our 15 benchmark datasets. As shown in Table 3, AdaCoT RL models
 326 in this forced "Always Reason" mode achieved average scores that were highly competitive with,
 327 and in instances like AdaCoT RL Model Exp4 (65.7) and Exp2 (65.3), slightly surpassed the Full
 328 CoT RL Baseline (65.0). This demonstrates a key strength: AdaCoT's adaptive training, aimed
 329 at optimizing the performance-cost Pareto frontier, does not inherently restrict the model's peak
 330 reasoning capabilities when comprehensive reasoning is explicitly demanded. While our Pareto
 331 analysis (Section 3.2.1) noted a slight gap in average scores when models operate adaptively, these
 332 SP-controlled results affirm that AdaCoT offers efficiency without sacrificing potential high-end
 333 performance.

334 An interesting secondary observation arose from the SFT stage concerning long-form generation (up
 335 to 32,000 tokens). AdaCoT SFT models, when directed by an "Always Reason SP," exhibited fewer

Table 3: Performance of AdaCoT RL models under "Always Reason" System Prompt vs. Full CoT RL Baseline, demonstrating preservation of peak performance. Metrics are averaged across the 15 benchmark datasets.

Model Variant	Avg. Score	CoT Triggering Rate
Full CoT RL Baseline	65.0	100%
AdaCoT RL Model Exp1 (w/ Always CoT SP)	64.8	100%
AdaCoT RL Model Exp2 (w/ Always CoT SP)	65.3	100%
AdaCoT RL Model Exp3 (w/ Always CoT SP)	64.9	100%
AdaCoT RL Model Exp4 (w/ Always CoT SP)	65.7	100%

326 instances of premature output truncation and were less prone to undesirable generative loops compared
 327 to a standard Full CoT SFT baseline. We hypothesize this improvement stems from AdaCoT's diverse
 328 SFT data, which includes many non-CoT examples (formatted as <think></think>answer). This
 329 results in a shorter average training sample length, potentially providing the End-of-Sequence (EOS)
 330 token a stronger learning signal (average EOS proportion: 0.000239 for AdaCoT SFT vs. 0.000215
 331 for Full CoT SFT). A more robust EOS representation could foster more coherent, well-terminated
 332 lengthy outputs, a promising area for future investigation.

333 5 Related Work

334 Chain-of-Thought (CoT) reasoning [2] significantly advanced LLM capabilities by prompting step-
 335 by-step thought processes, inspiring sophisticated strategies like diverse path sampling [44] or
 336 structured thoughts (trees [45], graphs [46]). However, CoT's verbosity and cost [47, 7] are major
 337 drawbacks. The varying utility of CoT—direct answers being better for some queries [48] while
 338 complex tasks need longer reasoning [49, 50]—underscores a critical cost-effectiveness trade-off,
 339 motivating research into efficient reasoning.

340 Most existing work on CoT efficiency has focused on *reducing reasoning length*, rather than adaptively
 341 deciding *whether* to invoke CoT. These length reduction strategies include: (1) RL with length
 342 penalties or rewards for brevity [8, 9, 10, 11, 12, 13]; (2) Restructuring or compressing CoT content
 343 via learned compact representations, iterative summarization, cognitive paradigms, or dynamic termina-
 344 tion [14, 15, 16, 17, 18, 19, 20, 21]; (3) Employing explicit instructions or selection mechanisms
 345 for conciseness, such as dynamic token allocation, prompts for shorter responses, distillation, or
 346 selecting the shortest valid reasoning [22, 23, 24, 25, 26, 27].

347 While effective for length, these methods generally do not equip a single model to dynamically
 348 decide CoT invocation based on query nature. Alternatives involve model merging [51, 52, 53] or
 349 manual CoT toggling. AdaCoT distinctively addresses adaptive triggering. By framing it as a Pareto
 350 optimization problem and using RL to control the CoT decision boundary, AdaCoT enables nuanced,
 351 context-dependent CoT invocation by a single model, filling a crucial gap towards truly efficient and
 352 versatile LLMs.

353 6 Conclusion

354 In this paper, we introduced AdaCoT, a novel framework for adaptive Chain-of-Thought reasoning
 355 in LLMs. By formulating adaptive reasoning as a Pareto optimization problem and employing an
 356 RL-based method with adjustable penalty coefficients, AdaCoT dynamically controls CoT triggering
 357 based on implicit query complexity. Our experiments demonstrate AdaCoT's success in navigating
 358 the Pareto frontier, achieving substantial reductions in CoT usage—and thereby computational costs
 359 and latency—for simpler queries, while preserving high performance on complex reasoning tasks.
 360 The introduction of Selective Loss Masking effectively ensures robust adaptive triggering throughout
 361 multi-stage RL training. Distinguishing itself from prior work predominantly focused on CoT length
 362 compression, AdaCoT pioneers adaptive CoT triggering, offering a practical and principled solution
 363 for developing more efficient, responsive, and cost-effective LLMs, particularly crucial for interactive
 364 applications.

365 **References**

366 [1] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin,
367 Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models
368 to follow instructions with human feedback. In *Advances in neural information processing
369 systems*, volume 35, pages 27730–27744, 2022.

370 [2] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le,
371 Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In
372 *Advances in neural information processing systems*, volume 35, pages 24824–24837, 2022.

373 [3] Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong
374 Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, et al. Towards large reasoning models: A
375 survey of reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*,
376 2025.

377 [4] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
378 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv
379 preprint arXiv:2412.16720*, 2024.

380 [5] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
381 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in
382 llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

383 [6] ByteDance Seed, Yufeng Yuan, Yu Yue, Mingxuan Wang, Xiaochen Zuo, Jiaze Chen, Lin Yan,
384 Wenyuan Xu, Chi Zhang, Xin Liu, et al. Seed-thinking-v1. 5: Advancing superb reasoning
385 models with reinforcement learning. *arXiv preprint arXiv:2504.13914*, 2025.

386 [7] Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuwei
387 Liu, Mengfei Zhou, Zhusong Zhang, et al. Do not think that much for $2+3=?$ on the
388 overthinking of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024.

389 [8] Daman Arora and Andrea Zanette. Training language models to reason efficiently. *arXiv
390 preprint arXiv:2502.04463*, 2025.

391 [9] Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li,
392 Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement
393 learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.

394 [10] Bairu Hou, Yang Zhang, Jiabao Ji, Yujian Liu, Kaizhi Qian, Jacob Andreas, and Shiyu Chang.
395 Thinkprune: Pruning long chain-of-thought of llms via reinforcement learning. *arXiv preprint
396 arXiv:2504.01296*, 2025.

397 [11] Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun
398 Cao, and Dacheng Tao. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning
399 pruning. *arXiv preprint arXiv:2501.12570*, 2025.

400 [12] Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang,
401 Kai Wang, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reasoning models.
402 *arXiv preprint arXiv:2503.04472*, 2025.

403 [13] Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with
404 reinforcement learning. *arXiv preprint arXiv:2503.04697*, 2025.

405 [14] Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuofei Qiao, Lun Du, Da Zheng, Huajun
406 Chen, and Ningyu Zhang. Lighthinker: Thinking step-by-step compression. *arXiv preprint
407 arXiv:2502.15589*, 2025.

408 [15] Yuchen Yan, Yongliang Shen, Yang Liu, Jin Jiang, Mengdi Zhang, Jian Shao, and Yueting
409 Zhuang. Inftythink: Breaking the length limits of long-context reasoning in large language
410 models. *arXiv preprint arXiv:2503.06692*, 2025.

411 [16] Simon A Aytes, Jinheon Baek, and Sung Ju Hwang. Sketch-of-thought: Efficient llm reasoning
412 with adaptive cognitive-inspired sketching. *arXiv preprint arXiv:2503.05179*, 2025.

413 [17] Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Zheng Lin, Li Cao, and
414 Weiping Wang. Dynamic early exit in reasoning models. *arXiv preprint arXiv:2504.15895*,
415 2025.

416 [18] Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. Chain of draft: Thinking faster by
417 writing less. *arXiv preprint arXiv:2502.18600*, 2025.

418 [19] Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. C3ot: Generating shorter chain-of-
419 thought without compromising effectiveness. In *Proceedings of the AAAI Conference on*
420 *Artificial Intelligence*, volume 39, pages 24312–24320, 2025.

421 [20] Heming Xia, Yongqi Li, Chak Tou Leong, Wenjie Wang, and Wenjie Li. Tokenskip: Controllable
422 chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*, 2025.

423 [21] Tengxiao Liu, Qipeng Guo, Xiangkun Hu, Cheng Jiayang, Yue Zhang, Xipeng Qiu, and Zheng
424 Zhang. Can language models learn to skip steps? In *Advances in Neural Information Processing*
425 *Systems*, volume 37, pages 45359–45385, 2024.

426 [22] Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen.
427 Token-budget-aware llm reasoning. *arXiv preprint arXiv:2412.18547*, 2024.

428 [23] Matthew Renze and Erhan Guven. The benefits of a concise chain of thought on problem-
429 solving in large language models. In *2024 2nd International Conference on Foundation and*
430 *Large Language Models (FLLM)*, pages 476–483. IEEE, 2024.

431 [24] Qiguang Chen, Libo Qin, Jiaqi Wang, Jingxuan Zhou, and Wanxiang Che. Unlocking the
432 capabilities of thought: A reasoning boundary framework to quantify and optimize chain-of-
433 thought. In *Advances in Neural Information Processing Systems*, volume 37, pages 54872–
434 54904, 2024.

435 [25] Ping Yu, Jing Xu, Jason E Weston, and Ilia Kulikov. Distilling system 2 into system 1. In *The*
436 *First Workshop on System-2 Reasoning at Scale, NeurIPS’24*, 2024.

437 [26] Tergel Munkhbat, Namgyu Ho, Seo Hyun Kim, Yongjin Yang, Yujin Kim, and Se-Young
438 Yun. Self-training elicits concise reasoning in large language models. *arXiv preprint*
439 *arXiv:2502.20122*, 2025.

440 [27] Jiabao Pan, Yan Zhang, Chen Zhang, Zuozhu Liu, Hongwei Wang, and Haizhou Li. Dynathink:
441 Fast or slow? a dynamic decision-making framework for large language models. *arXiv preprint*
442 *arXiv:2407.01009*, 2024.

443 [28] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures
444 of local experts. *Neural computation*, 3(1):79–87, 1991.

445 [29] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton,
446 and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts
447 layer. *arXiv preprint arXiv:1701.06538*, 2017.

448 [30] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal
449 policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

450 [31] Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Benjamin Feuer, Siddhartha Jain,
451 Ravid Schwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, Shubh-Agrawal, Sandeep Singh
452 Sandha, Siddartha Venkat Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie
453 Neiswanger, and Micah Goldblum. Livebench: A challenging, contamination-limited LLM
454 benchmark. In *The Thirteenth International Conference on Learning Representations*, 2025.

455 [32] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo,
456 Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and
457 challenging multi-task language understanding benchmark. In *Advances in Neural Information*
458 *Processing Systems*, volume 37, pages 95266–95290, 2025.

459 [33] Xinrun Du, Yifan Yao, Kaijing Ma, Bingli Wang, Tianyu Zheng, King Zhu, Minghao Liu,
460 Yiming Liang, Xiaolong Jin, Zhenlin Wei, et al. Supergpqa: Scaling llm evaluation across 285
461 graduate disciplines. *arXiv preprint arXiv:2502.14739*, 2025.

462 [34] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien
 463 Dirani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a
 464 benchmark. In *First Conference on Language Modeling*, 2024.

465 [35] Yancheng He, Shilong Li, Jiaheng Liu, Yingshui Tan, Weixun Wang, Hui Huang, Xingyuan
 466 Bu, Hangyu Guo, Chengwei Hu, Boren Zheng, et al. Chinese simpleqa: A chinese factuality
 467 evaluation for large language models. *arXiv preprint arXiv:2411.07140*, 2024.

468 [36] Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese,
 469 John Schulman, and William Fedus. Measuring short-form factuality in large language models.
 470 *arXiv preprint arXiv:2411.04368*, 2024.

471 [37] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn
 472 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In
 473 *NeurIPS*, 2021.

474 [38] Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu,
 475 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for
 476 promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of
 477 the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long
 478 Papers)*, pages 3828–3850, 2024.

479 [39] Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and
 480 Karthik R Narasimhan. SWE-bench: Can language models resolve real-world github issues? In
 481 *The Twelfth International Conference on Learning Representations*, 2024.

482 [40] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Ar-
 483 mando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination
 484 free evaluation of large language models for code. In *The Thirteenth International Conference
 485 on Learning Representations*, 2025.

486 [41] Kaijing Ma, Xinrun Du, Yunran Wang, Haoran Zhang, Zhoufutu Wen, Xingwei Qu, Jian Yang,
 487 Jiaheng Liu, Minghao Liu, Xiang Yue, et al. Kor-bench: Benchmarking language models on
 488 knowledge-orthogonal reasoning tasks. *arXiv preprint arXiv:2410.06526*, 2024.

489 [42] Ippei Fujisawa, Sensho Nobe, Hiroki Seto, Rina Onda, Yoshiaki Uchida, Hiroki Ikoma, Pei-
 490 Chun Chien, and Ryota Kanai. Procbench: Benchmark for multi-step reasoning and following
 491 procedure. *arXiv preprint arXiv:2410.03117*, 2024.

492 [43] Yanzhao Qin, Tao Zhang, Yanjun Shen, Wenjing Luo, Haoze Sun, Yan Zhang, Yujing Qiao,
 493 Weipeng Chen, Zenan Zhou, Wentao Zhang, et al. Sysbench: Can large language models follow
 494 system messages? *arXiv preprint arXiv:2408.10943*, 2024.

495 [44] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha
 496 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language
 497 models. In *The Eleventh International Conference on Learning Representations*, 2023.

498 [45] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
 499 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In
 500 *Advances in neural information processing systems*, volume 36, pages 11809–11822, 2023.

501 [46] Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas
 502 Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, et al. Graph
 503 of thoughts: Solving elaborate problems with large language models. In *Proceedings of the
 504 AAAI Conference on Artificial Intelligence*, volume 38, pages 17682–17690, 2024.

505 [47] Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi
 506 Liu, Andrew Wen, Hanjie Chen, Xia Hu, et al. Stop overthinking: A survey on efficient
 507 reasoning for large language models. *arXiv preprint arXiv:2503.16419*, 2025.

508 [48] Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning
 509 models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025.

510 [49] Yuyang Wu, Yifei Wang, Tianqi Du, Stefanie Jegelka, and Yisen Wang. When more is less:
511 Understanding chain-of-thought length in llms. *arXiv preprint arXiv:2502.07266*, 2025.

512 [50] Ayeong Lee, Ethan Che, and Tianyi Peng. How well do llms compress their own chain-of-
513 thought? a token complexity approach. *arXiv preprint arXiv:2503.01141*, 2025.

514 [51] Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. Cot-valve:
515 Length-compressible chain-of-thought tuning. *arXiv preprint arXiv:2502.09601*, 2025.

516 [52] Han Wu, Yuxuan Yao, Shuqi Liu, Zehua Liu, Xiaojin Fu, Xiongwei Han, Xing Li, Hui-Ling
517 Zhen, Tao Zhong, and Mingxuan Yuan. Unlocking efficient long-to-short llm reasoning with
518 model merging. *arXiv preprint arXiv:2503.20641*, 2025.

519 [53] Haotian Luo, Haiying He, Yibo Wang, Jinluan Yang, Rui Liu, Naiqiang Tan, Xiaochun Cao,
520 Dacheng Tao, and Li Shen. Adar1: From long-cot to hybrid-cot via bi-level adaptive reasoning
521 optimization. *arXiv preprint arXiv:2504.21659*, 2025.

522 **A Benchmark Dataset Details and Analysis**

523 This appendix provides descriptions for the benchmark datasets used in our evaluation and an analysis
 524 of the experimental results on these individual datasets. The scores presented in Table 4 form the basis
 525 for this analysis. For each dataset, we discuss the performance of baseline models and the AdaCoT
 526 variants. We also highlight any counter-intuitive results or observations that conflict with the primary
 527 motivation of achieving optimal performance with adaptive CoT, offering potential explanations such
 528 as evaluation volatility, inherent limitations of the base model, or aspects of the post-training process
 529 that may not be fully optimized for every scenario.

Table 4: Detailed scores on benchmark datasets. "TR" denotes reasoning triggering rate (%).

Dataset	nocot SFT baseline		nocot RL baseline		fullcot SFT baseline		fullcot RL baseline		Adaptive SFT Model		Adaptive RL Model Exp1		Adaptive RL Model Exp2		Adaptive RL Model Exp3		Adaptive RL Model Exp4	
	TR	Score	TR	Score	TR	Score	TR	Score	TR	Score	TR	Score	TR	Score	TR	Score	TR	Score
MMLU pro	0.0	77.5	0.0	82.1	100.0	83.7	100.0	85.2	40.3	80.5	28.0	74.2	27.3	83.2	39.6	84.1	58.0	83.1
super GPQA	0.0	49.8	0.0	50.7	100.0	53.0	100.0	58.6	51.0	59.5	22.0	55.1	13.0	58.9	56.2	59.6	59.6	59.6
LiveBench	0.0	50.0	0.0	56.6	100.0	57.7	100.0	69.5	65.8	58.9	45.1	64.7	57.1	66.3	71.6	68.4	70.4	59.2
KORBENCH	0.0	33.9	0.0	42.1	100.0	61.3	100.0	62.8	52.5	49.1	28.0	53.5	45.1	52.1	62.2	57.4	61.0	59.2
AIME24	0.0	23.3	0.0	33.3	100.0	69.3	100.0	84.7	100.0	69.3	100.0	86.7	100.0	86.0	100.0	88.0	100.0	86.3
AIME25	0.0	13.3	0.0	21.0	100.0	52.3	100.0	70.0	100.0	56.7	100.0	73.3	100.0	75.7	100.0	74.0	100.0	72.0
AMC10	0.0	84.3	0.0	88.0	100.0	90.0	100.0	97.0	95.9	95.5	91.7	93.9	95.5	95.6	95.2	95.6	95.2	95.2
LiveCodeBench	0.0	29.4	0.0	27.6	100.0	45.9	100.0	55.9	95.6	47.0	77.1	44.8	53.9	50.6	91.4	54.1	91.4	55.9
SWE-bench Agentless	0.0	27.4	0.0	27.0	100.0	36.4	100.0	37.7	98.6	35.4	1.6	28.8	44.0	35.8	94.2	37.6	79.6	36.6
Chinese SimpleQA	0.0	58.8	0.0	57.0	100.0	59.7	100.0	61.5	59.1	0.2	56.0	0.4	55.0	1.4	56.8	6.0	55.2	6.0
SimpleQA	0.0	10.8	0.0	10.3	100.0	12.7	100.0	12.2	0.7	10.3	0.0	9.9	0.0	10.7	4.4	11.3	20.0	9.6
ProveIt	0.0	42.6	0.0	46.3	100.0	53.7	100.0	68.6	73.6	50.2	50.8	53.0	79.1	68.7	89.3	69.6	93.1	72.4
GPQA(diamond)	0.0	59.5	0.0	64.3	100.0	54.9	100.0	70.5	63.7	62.6	65.7	84.1	67.2	92.4	70.8	96.5	72.6	
SysBench	0.0	42.6	0.0	48.1	100.0	56.2	100.0	62.4	35.6	52.0	1.2	57.2	4.0	55.5	38.2	60.2	15.4	55.7
Olympiad Bench	0.0	51.4	0.0	60.1	100.0	73.5	100.0	78.2	85.0	75.3	88.0	80.2	89.0	81.8	93.5	80.7	95.4	82.0
Average	0.0	43.6	0.0	47.7	100.0	58.5	100.0	65.0	61.3	57.1	43.1	59.7	53.3	62.8	65.4	64.3	67.7	64.4

530 **MMLU pro:** An enhanced version of the MMLU benchmark, MMLU-Pro integrates more challenging-
 531 reasoning-focused questions, expands the choice set from four to ten options, and eliminates
 532 trivial/noisy questions from the original MMLU. It is designed to better discern model capabilities,
 533 particularly in complex reasoning, where CoT has shown greater benefit compared to direct answering
 534 on this version. *Analysis:* CoT clearly benefits performance, with the FullCoT RL baseline (85.2)
 535 significantly outperforming NoCoT SFT (77.5) and NoCoT RL (82.1). The Adaptive SFT Model
 536 (40.3% TR, 80.5 score) shows improvement over NoCoT SFT but doesn't reach FullCoT SFT levels
 537 (83.7). AdaCoT RL Exp3 (39.6% TR, 84.1 score) and Exp2 (27.3% TR, 83.2 score) achieve strong
 538 scores, with Exp3 surpassing FullCoT SFT and Exp2 performing comparably. AdaCoT RL Exp4
 539 (58.0% TR, 83.1 score) also performs well. Exp1 (28.0% TR, 74.2 score) shows a drop, indicating
 540 that for MMLU pro, a moderate CoT rate is generally beneficial, reflecting the benchmark's
 541 increased reasoning demands. The adaptive models demonstrate an ability to achieve high scores
 542 with significantly reduced CoT compared to FullCoT RL.

543 **super GPQA:** A comprehensive benchmark evaluating graduate-level knowledge and reasoning capabilities across 285 disciplines, particularly including specialized fields in light industry, agriculture, and service-oriented areas often underrepresented in other benchmarks. It employs a Human-LLM collaborative filtering mechanism to ensure high question quality by eliminating trivial or ambiguous questions. *Analysis:* CoT provides a clear advantage (NoCoT RL 50.7 vs. FullCoT RL 58.6). The Adaptive SFT Model (35.2% TR, 51.0 score) also show a modest gain over NoCoT SFT. AdaCoT RL Exp4 (59.6% TR) notably achieves a score of 59.6, surpassing the FullCoT RL baseline with significantly less CoT. AdaCoT RL Exp2 (32.3% TR, 56.9 score) also outperforms FullCoT SFT (53.8). This suggests AdaCoT effectively adapts CoT usage for these specialized, high-level questions, achieving strong performance efficiently.

553 **LiveBench:** A benchmark designed to be resistant to test set contamination and the pitfalls of LLM/human-crowdsourced judging. It features frequently updated questions from recent sources (math competitions, arXiv papers, news, datasets), scores answers automatically against objective ground-truth, and includes a wide variety of challenging tasks (math, coding, reasoning, language, instruction following, data analysis), including harder, contamination-limited versions of tasks from previous benchmarks. *Analysis:* This dataset shows significant gains from both CoT and RL (NoCoT SFT 50.0 to FullCoT RL 69.5). The NoCoT RL baseline (56.6) and Adaptive SFT Model (65.8% TR, 58.9 score) both outperform NoCoT SFT, with Adaptive SFT also surpassing FullCoT SFT (57.7). AdaCoT RL Exp4 (70.4% TR, 69.2 score) very closely approaches the FullCoT RL baseline performance with about 30% less CoT. AdaCoT RL Exp3 (71.6% TR, 68.4 score) is also strong. AdaCoT RL Exp2 (57.1% TR, 66.3 score) substantially outperforms FullCoT SFT. The robust design of LiveBench makes it a strong test case for AdaCoT's adaptive reasoning, showing it can maintain high performance with adaptive CoT.

566 **KORBENCH**: This benchmark evaluates Knowledge-Orthogonal-Reasoning, aiming to minimize
567 reliance on domain-specific knowledge to more accurately assess models' reasoning abilities in out-of-
568 distribution settings. It includes five task categories (Operation, Logic, Cipher, Puzzle, Counterfactual)
569 and emphasizes models' effectiveness in applying new rule descriptions to solve novel rule-driven
570 questions. *Analysis*: Scores show a clear benefit from CoT: NoCoT SFT (33.9) is significantly lower
571 than FullCoT SFT (61.3) and FullCoT RL (62.8). The NoCoT RL baseline (42.1) improves over
572 NoCoT SFT. The Adaptive SFT Model (52.5% TR, 49.1 score) sits between the NoCoT baselines
573 and FullCoT SFT. AdaCoT RL models demonstrate adaptive behavior: Exp4 (61.0% TR, 59.2 score)
574 and Exp3 (62.2% TR, 57.4 score) approach the FullCoT SFT baseline performance with significantly
575 less CoT than FullCoT models. Exp1 (28.0% TR, 53.5 score) is also effective. This suggests AdaCoT
576 effectively discerns when to apply CoT for these rule-driven tasks, though peak performance is
577 slightly below FullCoT RL.

578 **AIME24 / AIME25**: Representing problems from the American Mathematics Invitational Exam-
579 nation for 2024 and 2025, these datasets are used to evaluate mathematical reasoning and problem-
580 solving abilities. *Analysis*: These mathematics-intensive datasets show massive performance gains
581 from CoT (e.g., AIME24: NoCoT SFT 23.3 vs. FullCoT RL 84.7). All AdaCoT RL models and the
582 Adaptive SFT Model correctly identify the complexity, exhibiting a 100% CoT triggering rate. For
583 AIME24, Adaptive SFT (69.3) matches FullCoT SFT (69.3). AdaCoT RL Exp3 (88.0) and Exp1
584 (86.7) outperform the FullCoT RL baseline (84.7). For AIME25, Adaptive SFT (56.7) surpasses
585 FullCoT SFT (52.3). AdaCoT RL Exp2 (75.7) and Exp3 (74.0) outperform the FullCoT RL baseline
586 (70.0). This is a notable result, suggesting that the adaptive training regimen, even when defaulting to
587 100% CoT for such complex problems, might confer some benefits, potentially due to the diversity
588 in training data (including non-CoT examples) leading to a more robust underlying model or better
589 fine-tuning dynamics.

590 **MATH**: A dataset of 12,500 challenging competition mathematics problems, each with a full step-
591 by-step solution. It is designed to measure mathematical problem-solving ability. *Analysis*: CoT is
592 highly beneficial (NoCoT RL 88.7 vs. FullCoT RL 97.3). The Adaptive SFT Model (44.6% TR,
593 95.5 score) performs well, nearly matching FullCoT SFT (96.5) with less than half the CoT. AdaCoT
594 models adapt effectively: AdaCoT RL Exp4 (68.8% TR, 97.2 score) nearly matches the FullCoT RL
595 baseline with about 30% less CoT. AdaCoT RL Exp2 (52.8% TR, 95.9 score) also performs strongly.
596 Exp1 (40.9% TR, 91.7 score) is lower, indicating that for MATH, higher CoT rates are generally
597 more beneficial among the adaptive RL models, but significant efficiency is still gained.

598 **LiveCodeBench**: A comprehensive and contamination-free evaluation benchmark for LLMs on
599 code, collecting new problems over time from programming contests. It assesses a broader range of
600 code-related capabilities. *Analysis*: CoT improves performance significantly (NoCoT SFT 29.4 vs.
601 FullCoT RL 55.9). NoCoT RL (27.6) is surprisingly lower than NoCoT SFT here, which might be due
602 to evaluation noise or specific sensitivities of the RL fine-tuning on non-CoT data for this particular
603 task. The Adaptive SFT Model (95.0% TR, 47.0 score) uses a high trigger rate and surpasses FullCoT
604 SFT (45.9). AdaCoT RL models trigger CoT at high rates: Exp4 (91.4% TR, 55.9 score) matches the
605 FullCoT RL baseline score with slightly less CoT. Exp3 (91.4% TR, 54.1 score) and Exp2 (83.9%
606 TR, 50.6 score) are also strong. This indicates recognition of coding task complexity and efficient
607 application of CoT.

608 **SWE-bench Agentless**: An evaluation framework consisting of 2,294 software engineering problems
609 from real GitHub issues. Models are tasked with editing codebases to resolve issues. *Analysis*: CoT
610 provides a notable benefit (NoCoT SFT 27.4 vs. FullCoT RL 37.7). NoCoT RL (27.0) is similar
611 to NoCoT SFT. The Adaptive SFT Model (98.6% TR, 35.4 score) uses a very high trigger rate and
612 performs close to FullCoT SFT (36.4). AdaCoT RL Exp3 (94.2% TR, 37.6 score) nearly matches
613 the FullCoT RL baseline with slightly less CoT. Interestingly, AdaCoT RL Exp1 (1.6% TR, 28.8
614 score) shows a slight improvement over NoCoT SFT with minimal reasoning. This suggests some
615 issues might be simpler, or the model is highly conservative in Exp1, but for complex software issues,
616 high CoT rates are beneficial. The performance of Exp2 (44.0% TR, 35.8 score) is also noteworthy,
617 achieving good results with moderate CoT.

618 **Chinese SimpleQA**: The first comprehensive Chinese benchmark to evaluate the factuality of
619 language models in answering short questions. *Analysis*: CoT offers minimal gains (NoCoT SFT 58.8
620 to FullCoT RL 61.5). NoCoT RL (57.0) is slightly lower than NoCoT SFT. The Adaptive SFT Model
621 (0.3% TR, 59.1 score) performs very well, slightly exceeding NoCoT SFT and approaching FullCoT

622 SFT (59.7) with extremely low CoT usage. AdaCoT RL models trigger CoT very infrequently (0.2%
623 to 6.0%), correctly identifying these as simple questions. Scores for AdaCoT RL models (e.g., Exp1
624 56.0, Exp3 56.8) are slightly below NoCoT SFT. This is a good demonstration of AdaCoT’s core
625 motivation: avoiding unnecessary CoT. While there’s a slight dip compared to NoCoT SFT for
626 some RL models, the Adaptive SFT model shows an excellent trade-off. The minor performance
627 variations could be due to the model sometimes being overly conservative in triggering CoT or slight
628 instabilities in evaluating purely factual recall without reasoning.

629 **SimpleQA:** A benchmark designed to measure the factuality of language models using short, fact-
630 seeking queries. *Analysis:* Similar to Chinese SimpleQA, CoT provides little benefit; FullCoT RL
631 (12.2) is slightly worse than FullCoT SFT (12.7). NoCoT RL (10.3) is slightly below NoCoT SFT
632 (10.8). The Adaptive SFT Model (0.7% TR, 10.3 score) matches NoCoT RL with minimal CoT.
633 AdaCoT RL models trigger CoT very rarely (Exp1 & Exp2 at 0.0%). AdaCoT RL Exp3 (4.4% TR,
634 11.3 score) performs better than NoCoT SFT. This reinforces that for simple QA, adaptive triggering
635 is crucial for efficiency. The performance of Exp4 (20.0% TR, 9.6 score) is slightly counter-intuitive,
636 as higher CoT did not yield better results and was worse than NoCoT SFT; this might indicate that for
637 very simple questions, forcing CoT (even adaptively at a higher rate) can sometimes be detrimental
638 or that the specific penalty balance for Exp4 was not optimal for this type of dataset.

639 **ProcBench:** This benchmark focuses on the direct evaluation of multi-step inference by largely
640 eliminating path exploration and implicit knowledge utilization. *Analysis:* CoT is highly beneficial
641 (NoCoT RL 46.7 vs. FullCoT RL 68.6). The Adaptive SFT Model (73.6% TR, 50.2 score) improves
642 over NoCoT baselines but is below FullCoT SFT (53.7). AdaCoT RL models show high trigger
643 rates, with Exp4 (93.1% TR, 72.4 score) significantly surpassing the FullCoT RL baseline. Exp2
644 (79.1% TR, 68.7 score) and Exp3 (89.3% TR, 69.6 score) also match or exceed FullCoT RL. This
645 indicates effective identification of tasks requiring detailed, step-by-step procedural reasoning and
646 demonstrates that adaptive models can even outperform always-on CoT models in certain complex
647 reasoning scenarios.

648 **GPQA (diamond):** GPQA is a challenging dataset of 448 multiple-choice questions by domain
649 experts in biology, physics, and chemistry. "GPQA (diamond)" refers to this specific challenging
650 set. *Analysis:* CoT significantly boosts performance (NoCoT RL 64.3 vs. FullCoT RL 70.5). NoCoT
651 RL is better than NoCoT SFT and close to FullCoT SFT (64.9). The Adaptive SFT Model (92.6%
652 TR, 65.7 score) also performs well, exceeding FullCoT SFT. AdaCoT RL models trigger CoT at
653 high rates. AdaCoT RL Exp4 (96.5% TR, 72.6 score) and Exp3 (92.4% TR, 70.8 score) demonstrate
654 strong performance, with Exp4 outperforming the FullCoT RL baseline. This highlights appropriate
655 and effective CoT invocation on these very hard questions.

656 **SysBench:** A benchmark for evaluating language models on their ability to understand and generate
657 content related to computer systems, configurations, or system-level concepts. *Analysis:* CoT is
658 beneficial (NoCoT RL 48.1 vs. FullCoT RL 62.4). The Adaptive SFT Model (35.6% TR, 52.0 score)
659 is better than NoCoT RL but below FullCoT SFT (56.2). AdaCoT RL models show good adaptation.
660 AdaCoT RL Exp3 (38.2% TR, 60.2 score) approaches the FullCoT RL baseline with significantly
661 reduced CoT. AdaCoT RL Exp1 (1.2% TR, 57.2 score) surpasses FullCoT SFT with very minimal
662 CoT usage, which is an excellent result for efficiency. This suggests a mix of complexities within
663 SysBench, which AdaCoT navigates effectively, although Exp4 (15.4% TR, 55.7 score) shows a
664 somewhat lower score despite a higher CoT rate than Exp1, possibly due to the specific balance of
665 penalties in Exp4 not being optimal for this dataset’s particular mix.

666 **OlympiadBench:** An Olympiad-level bilingual multimodal scientific benchmark with problems from
667 mathematics and physics competitions. *Analysis:* CoT provides substantial gains (NoCoT RL 60.1 vs.
668 FullCoT RL 78.2). The Adaptive SFT Model (85.0% TR, 75.3 score) performs strongly, exceeding
669 FullCoT SFT (73.5). AdaCoT RL models exhibit high trigger rates. AdaCoT RL Exp4 (95.4% TR,
670 82.0 score) and Exp2 (89.0% TR, 81.8 score) both surpass the FullCoT RL baseline. This indicates
671 strong reasoning capabilities and appropriate CoT usage on these exceptionally challenging problems,
672 again showing adaptive models can reach or exceed the performance of specialized always-on CoT
673 models.

674 **Overall Summary of Per-Dataset Analysis:** The adaptive models, including both Adaptive SFT and
675 the AdaCoT RL variants, demonstrate effective adaptation across a diverse range of benchmarks.

- 676 • The Adaptive SFT Model serves as a strong adaptive baseline. It often improves significantly over NoCoT baselines by selectively triggering CoT (e.g., high rates for AIME, 677 LiveCodeBench, SWE-bench; low rates for SimpleQAs). On some complex tasks (AIME25, 678 LiveCodeBench, GPQA(diamond), OlympiadBench), it even surpasses the FullCoT SFT 679 baseline, and for Chinese SimpleQA, it achieves excellent efficiency and performance. 680 However, its average performance (57.1 score, 61.3% TR) is generally below the peak 681 performance of FullCoT RL (65.0 score) or the best AdaCoT RL experiments (e.g., Exp4: 682 64.4 score, 67.7% TR).
- 684 • On complex reasoning tasks (e.g., AIME, MATH, OlympiadBench, GPQA(diamond), 685 ProcBench), AdaCoT RL models tend to trigger CoT at high rates. Several AdaCoT RL 686 experiments (notably Exp4 on super GPQA, ProcBench, GPQA(diamond), OlympiadBench; 687 Exp3 on AIME24; Exp2 on AIME25, OlympiadBench) match or exceed the performance of 688 FullCoT RL baselines, showcasing the benefits of learned adaptive policies and suggesting 689 that adaptive training can sometimes lead to better overall models even for tasks that always 690 require CoT.
- 691 • On simpler tasks or those designed to test factuality (e.g., Chinese SimpleQA, SimpleQA), 692 AdaCoT RL models trigger CoT very sparingly. This leads to computational savings 693 while generally maintaining performance near NoCoT or FullCoT SFT levels, successfully 694 avoiding unnecessary CoT. The Adaptive SFT model also excels in efficiency here. Some 695 minor performance drops in RL models compared to NoCoT SFT on these tasks (e.g., 696 Chinese SimpleQA for Exp1/Exp2) might be attributed to the RL agent being slightly too 697 conservative or the inherent difficulty in perfectly balancing penalties for extremely low 698 CoT rate scenarios without any performance degradation.
- 699 • For benchmarks with mixed or specific reasoning types (e.g., KORBENCH, LiveBench, 700 MMLU-Pro, SysBench), both Adaptive SFT and AdaCoT RL models show nuanced 701 adaptation, adjusting CoT rates to balance performance and efficiency. They often outperform 702 static baselines or achieve comparable results with lower CoT usage. For instance, on 703 SysBench, AdaCoT RL Exp1 achieved a higher score than FullCoT SFT with only 1.2% 704 TR.
- 705 • The different AdaCoT RL experiments (Exp1-Exp4) effectively trace a Pareto frontier (as 706 shown in Section 3.2.1 using a specific set of average scores), offering a trade-off between 707 CoT triggering rate and performance, adaptable to specific deployment needs. Based on the 708 average scores from Appendix 4, Exp4 (64.4 score, 67.7% TR) and Exp3 (64.3 score, 65.4% 709 TR) represent high-performance points, closely approaching the FullCoT RL baseline (65.0 710 score) with about 30-35% less CoT usage on average. Exp2 (62.8 score, 53.3% TR) also 711 offers a strong balance.
- 712 • Some counter-intuitive results, like NoCoT RL performing worse than NoCoT SFT on 713 LiveCodeBench, or AdaCoT RL Exp4 on SimpleQA (higher CoT, lower score), could 714 be due to factors like evaluation volatility on specific datasets, the base model’s inherent 715 capabilities or sensitivities to fine-tuning on certain data distributions, or sub-optimal penalty 716 configurations for specific outlier datasets within a broadly tuned RL policy. The post- 717 training process aims for general improvement, and individual dataset performance can 718 fluctuate.

719 These detailed results underscore the ability of adaptive strategies, both SFT-based and RL-based, 720 to make nuanced decisions about CoT invocation, optimizing for both performance and efficiency 721 based on query characteristics and benchmark demands. The AdaCoT RL models, in particular, 722 demonstrate the potential to significantly reduce CoT overhead while maintaining competitive, and in 723 some cases superior, performance compared to full CoT strategies.

724 **B Principle-Guided CoT Assessment**

725 This appendix details the principle-guided assessment framework used to annotate data for Chain- 726 of-Thought (CoT) necessity. As described in Section 2.2.1 of the main paper, an auxiliary model 727 utilizes these principles to label queries as either likely benefiting from CoT or suitable for a direct 728 answer. This labeling is crucial for the Supervised Fine-Tuning (SFT) warm-up stage of the AdaCoT

729 framework, providing an initial understanding for the model on when to employ CoT. The specific
730 principles provided to the auxiliary model are outlined below.

```
731 Given a dialogue between a user and an AI assistant, please
732 consider the conversation context and, from the AI
733 assistant's perspective, assess the difficulty of
734 answering the user's final question according to the
735 following requirements.
736 <AI assistant's system prompt-Start>
737 {system_prompt}
738 <AI assistant's system prompt-End>
739 <Dialogue history-Start>:
740 {history}
741 <Dialogue history-End>
742 <User's final question-Start>
743 {last_prompt}
744 <User's final question-End>
745
746 ## Assessment Process
747 1. Carefully read the provided prompt and any relevant
748 context (if any).
749 2. Evaluate the 'question difficulty' based on the
750 following assessment criteria.
751 3. The output assessment result must strictly adhere to the
752 specified output format requirements.
753
754 ## Assessment Criteria
755
756 ### Whether In-depth Thinking is Required
757 - **Requires In-depth Thinking**:
758 - Requires multi-step reasoning and analysis to arrive at
759 the answer.
760 - Requires a logical chain and coherent reasoning process
761 .
762 - May involve breaking down and synthesizing complex
763 concepts.
764 - Requires systematic thinking and structured analysis.
765 - The question is clear, but its content is complex.
766 - Requires deep thought and specialized knowledge.
767 - Demands multi-angle analysis and synthesis capabilities
768 .
769 - Requires creative thinking or unique insights.
770 - There might be multiple reasonable paths to the answer.
771 - May require integration of cross-domain knowledge.
772
773 - **Does Not Require In-depth Thinking**:
774 - The answer can be obtained through direct recall or
775 simple lookup.
776 - Does not require a complex chain of logical reasoning.
777 - The path to solving the problem is clear and direct.
778 - The solution process is singular and linear.
779 - Does not require synthesizing multiple pieces of
780 information or viewpoints.
781 - Basic facts or common sense are sufficient to address
782 it.
783 - The question is simple and direct.
784 - Can be answered without in-depth thought.
785 - The answer is definite or standardized.
786 - The solution path is singular and obvious.
787
788
```

```

789     ## Output Requirements
790     Answer using labels. Possible options are:
791     - Requires In-depth Thinking
792     - Does Not Require In-depth Thinking
793
794     Output only the assessment result, without providing any
795     explanation or additional comments.

```

797 C System Prompt Examples

798 This appendix provides examples of the System Prompts (SPs) employed to control the CoT triggering
 799 behavior of our models, as mentioned in Section 4.3 of the main paper. These SPs allow for explicit
 800 instruction on whether the model should always or never engage in CoT reasoning. We present the
 801 content of these SPs and illustrative examples of their application.

802 C.1 Always CoT SP

803 The following System Prompt is used to instruct the model to always use CoT:

804 You possess the capability for deep thinking. Before responding to the user,
 805 you will engage in comprehensive and in-depth thought. The thinking process
 806 should be enclosed in `<think></think>` tags, and the final answer should be
 807 summarized based on this thinking process. For example:
 808 `<think>This is the thinking process...</think> This is the final answer.`
 809 For all questions posed by the user, you need to think deeply before answering.

810 Example: Simple Question with Always CoT SP

811 User input:

812 What is $1 + 1$?

813 Model response:

```

814 <think>
815 The user is asking what  $1+1$  equals, which appears to be a very
816 basic mathematical question. First, I need to confirm the
817 user's actual intent. It could be a young child just
818 starting to learn math, or someone testing my basic
819 calculation abilities. It's also possible they are joking,
820 like the common internet meme about " $1+1=3$ " involving
821 family concepts, but typically in such cases, the correct
822 mathematical answer should be given first, followed by
823 consideration of whether an extended explanation is needed.
824 From a mathematical perspective, the standard answer to  $1+1$  is
825 2. This is based on the definition of addition for natural
826 numbers and falls under basic arithmetic. As a knowledge-
827 based Q&A, the primary goal is to provide an accurate
828 mathematical answer, ensuring correctness. Then, if
829 necessary, I can briefly explain the basic concept of
830 addition, or mention that in different contexts there might
831 be different answers, such as  $1+1=10$  in binary, or
832 different results in logical operations. However, these
833 additions should only be made after clarifying the user's
834 needs to avoid confusing basic concepts.
835
836 Considering the user might be a beginner, I should first
837 provide a clear and correct answer, then decide whether to
838 elaborate based on the situation. Therefore, the response
839

```

```

840     structure should be to first directly answer 2, then
841     briefly mention other possibilities, but it's important to
842     maintain a clear distinction between primary and secondary
843     information to avoid misleading. At the same time, the
844     response should be friendly and concise, meeting the user's
845     actual needs.
846 </think>
847 In basic arithmetic operations in mathematics, the result of 1
848 + 1 is equal to 2. This is the fundamental definition of
849 addition for natural numbers and is the universally
850 accepted standard answer.
851
852 However, in some special contexts, there might be different
853 interpretations:
854
855 In binary operations, 1 + 1 = 10 (binary "10" corresponds to
856 decimal "2");
857 In creative or humorous thinking, there might be symbolic
858 answers (e.g., 1 drop of water + 1 drop of water = 1 large
859 drop of water), but this falls into non-mathematical
860 analogies.
861
862 So, the most direct answer is 2.

```

864 C.2 Never CoT SP

865 The following System Prompt is used to instruct the model to never use CoT (i.e., provide direct
866 answers):

867 You possess the capability for deep thinking. Before responding to the user,
868 you will engage in comprehensive and in-depth thought. The thinking process
869 should be enclosed in `<think></think>` tags, and the final answer should be
870 summarized based on this thinking process. For example:
871 `<think>This is the thinking process...</think>` This is the final answer.
872 For all questions posed by the user, you need to answer directly
873 without engaging in in-depth thinking.

874 Example: Complex Question with Never CoT SP

875 User input:

876 Alice and Bob play the following game. A stack of n tokens lies before them.
877 The players take turns with Alice going first. On each turn, the player
878 removes either 1 token or 4 tokens from the stack. Whoever removes the
879 last token wins. Find the number of positive integers n less than or equal
880 to 2024 for which there exists a strategy for Bob that guarantees that Bob
881 will win the game regardless of Alice's play.

882 Model response:

```

883 <think></think>
884 To determine the number of positive integers  $\backslash( n \leq 2024 \backslash)$ 
885 for which Bob has a winning strategy, we analyze the game
886 using modular arithmetic. We define a position as a losing
887 position if the current player cannot win regardless of
888 their moves, assuming the opponent plays optimally.
889
890
891 Key steps:
892 1. **Base Cases**:
893 -  $\backslash( n = 1 \backslash)$ : Alice removes the last token and wins. Not a
894     losing position.

```

```

895 -  $\backslash(n = 2 \backslash)$ : Alice must remove 1 token, leaving  $\backslash(n = 1 \backslash)$  for Bob. Bob removes the last token and wins. Thus,  $\backslash(n = 2 \backslash)$  is a losing position for Alice (winning for Bob).
896 -  $\backslash(n = 3 \backslash)$ : Alice removes 1 token, leaving  $\backslash(n = 2 \backslash)$  (a losing position for Bob). Bob loses, so  $\backslash(n = 3 \backslash)$  is not a losing position.
897 -  $\backslash(n = 4 \backslash)$ : Alice removes 4 tokens and wins. Not a losing position.
898 -  $\backslash(n = 5 \backslash)$ : Alice removes 1 token, leaving  $\backslash(n = 4 \backslash)$  (Bob loses). If Alice removes 4 tokens, leaving  $\backslash(n = 1 \backslash)$  (Bob loses). Thus,  $\backslash(n = 5 \backslash)$  is not a losing position.
899 -  $\backslash(n = 6 \backslash)$ : Any move by Alice (removing 1 or 4 tokens) leaves  $\backslash(n = 5 \backslash)$  or  $\backslash(n = 2 \backslash)$ .  $\backslash(n = 5 \backslash)$  is a winning position, but  $\backslash(n = 2 \backslash)$  is a losing position. However, further analysis shows a pattern emerging.
900
901
902
903
904
905
906
907
908
909
910
911
912
913 2. **Inductive Pattern**:
914 - By examining more cases, we observe that losing positions (where Bob can win) occur at  $\backslash(n \equiv 2 \pmod{5})$ . This is because:
915 - If  $\backslash(n \equiv 2 \pmod{5})$ , any move by Alice (removing 1 or 4 tokens) results in  $\backslash(n \equiv 1 \pmod{5})$  or  $\backslash(n \equiv 3 \pmod{5})$ , which are winning positions for Bob to counter.
916 - Conversely, if  $\backslash(n \backslash)$  is not  $\backslash(2 \pmod{5})$ , Alice can force the game into a losing position for Bob.
917
918
919
920
921
922
923
924 3. **Counting Losing Positions**:
925 - We need to count all  $\backslash(n \leq 2024 \backslash)$  such that  $\backslash(n \equiv 2 \pmod{5})$ . These numbers form the sequence  $\backslash(2, 7, 12, \ldots, 2022 \backslash)$ .
926 - This is an arithmetic sequence with first term  $\backslash(a_1 = 2 \backslash)$ , common difference  $\backslash(d = 5 \backslash)$ , and last term  $\backslash(a_k = 2022 \backslash)$ .
927 - The number of terms  $\backslash(k \backslash)$  is given by solving  $\backslash(2 + (k - 1)5 = 2022 \backslash)$ :
928 
$$\begin{aligned} \backslash[ \\ (k-1)5 &= 2020 \implies k-1 = 404 \implies k = 405 \\ \backslash] \end{aligned}$$

929
930
931
932
933
934
935
936
937 Thus, the number of positive integers  $\backslash(n \leq 2024 \backslash)$  for which Bob has a winning strategy is  $\backslash(\boxed{405}\backslash)$ .
938

```

940 D Meta-Reasoning Showcase: User Prompt Control

941 As discussed in Section 4.2 of the main paper, our exploration into meta-reasoning during the
942 SFT stage revealed an emergent capability for user-prompt controllability over CoT. This appendix
943 showcases examples where the SFT model, enhanced with meta-reasoning (as illustrated in Figure 2),
944 adapts its CoT triggering based on explicit user instructions regarding the desired level of reasoning.

945 D.1 Case 1: Historical Creative Question

946 Original Question

947 User input:

948 On the night of the Xuanwu Gate Incident, what would Li Shimin write

949 if he were alone?

950 Model response:

951
952 <think>
953 This is a complex or difficult question, hard to answer
954 correctly directly, so I need to think about it. The user's
955 question is: "On the night of the Xuanwu Gate Incident,
956 what would Li Shimin write if he were alone?" The first
957 step is to determine the core of the user's question. The
958 user wants to know what kind of text Li Shimin might have
959 written on the night of the Xuanwu Gate Incident, that is,
960 after this coup occurred. This might include his inner
961 thoughts at the time, his plans for the future, or
962 reflections on the events that had transpired.
963
964 Next, I need to review the historical background of the Xuanwu
965 Gate Incident. The Xuanwu Gate Incident took place in 626
966 AD. It was a coup launched by Li Shimin at Xuanwu Gate to
967 seize the imperial throne, during which he killed his elder
968 brother Li Jiancheng and his younger brother Li Yuanji.
969 This coup directly led to Emperor Gaozu (Li Yuan)
970 abdicating, allowing Li Shimin to ascend to the throne.
971 From this, it's clear that the night after the coup was an
972 extremely critical moment for Li Shimin, and he likely had
973 many thoughts.
974
975 The user might be asking this question out of historical
976 interest, or perhaps for some creative endeavor, like
977 writing a novel or screenplay, or to gain a deeper
978 understanding of Li Shimin's psychological state. They
979 might prefer an answer that is both historically grounded
980 and has a certain literary quality, rather than a simple
981 list of historical facts. Therefore, when answering, I must
982 consider Li Shimin's situation at the time, his emotions,
983 and his possible motivations.
984
985 After the coup, although Li Shimin was victorious, he also
986 faced many problems. On one hand, he had to manage his
987 relationship with his father, Li Yuan, who was still alive
988 and needed to abdicate. On the other hand, he had to
989 consolidate his power and prevent any backlash from the
990 supporters of Li Jiancheng and Li Yuanji. At the same time,
991 he needed to consider how to explain the legitimacy of
992 this coup to the world to maintain his ruling authority.
993 Furthermore, he might also have felt moral guilt for his
994 actions, having killed his own brothers.
995
996 If Li Shimin did write something that night, the content could
997 cover multiple aspects. He might record the process of the
998 coup, stating that he acted in self-defense. He might also
999 elaborate on his governance philosophy to lay the
1000 groundwork for his future rule. He might express filial
1001 piety towards Li Yuan, emphasizing that launching the coup
1002 was a last resort. He might even mention arrangements for
1003 the families of Li Jiancheng and Li Yuanji to alleviate his
1004 sense of guilt.
1005
1006 However, there is no precise historical record of what Li
1007 Shimin wrote that night. Therefore, the answer can only be

1008 a reasonable speculation based on historical facts. The
1009 structure of the answer could be divided into several parts
1010 , such as an explanation of the coup, his attitude towards
1011 Li Yuan, plans for the future, and guilt towards his
1012 brothers. Moreover, the language style of the answer could
1013 lean towards classical Chinese to enhance the historical
1014 feel, but it must also ensure the content is easy to
1015 understand.

1016

1017 I also need to pay attention to the user's potential deeper
1018 needs. They might not just want to know what Li Shimin
1019 might have written, but also to understand his inner world
1020 at that time and the impact of this coup on the history of
1021 the Tang Dynasty. So, I can appropriately mention these
1022 aspects in the answer to add depth.

1023

1024 Finally, I need to check if the answer aligns with historical
1025 facts, if it reasonably speculates on Li Shimin's
1026 psychological activities, if the language is fluent, and if
1027 it can meet the user's needs.</think>

1028 On the night of the Xuanwu Gate Incident, Qin Wang Li Shimin
1029 might have, under candlelight, penned a text blending
1030 complex emotions and political strategy. The following
1031 content is a reasonable speculation based on historical
1032 context:

1033

1034 "Memoir on the Night of the Xuanwu Gate Incident"

1035

1036 Ninth year of Wude, Gengshen night.

1037

1038 My elder brother Jiancheng and younger brother Yuanji conspired
1039 in rebellion. They harbored assassins, secretly
1040 transported armor, and plotted to harm me. Left with no
1041 choice, I could only set up defenses at Xuanwu Gate. Now,
1042 their heads hang upon the city tower. This upheaval has
1043 shocked the court and beyond; perhaps Father Emperor also
1044 harbors resentment towards me. But for the sake of the Tang
1045 Empire's stability, I had to do so.

1046

1047 Recalling the years past, I fought alongside my generals,
1048 shedding blood to pacify warlords like Dou Jiande and Wang
1049 Shichong. Yet now, it has come to fratricide. Is all this
1050 destiny, or the work of my own inner desires?

1051

1052 Father Emperor, oh Father Emperor, your son is not an unfilial
1053 or disloyal person. It is just that you repeatedly favored
1054 Jiancheng, turning a blind eye to his misdeeds. Your son
1055 was truly worried that the Tang Empire would be ruined at
1056 their hands. Now that the die is cast, I hope you can
1057 follow the mandate of heaven and pass the throne to your
1058 son. Your son will surely usher in an era of peace and
1059 prosperity, allowing the common people to live and work in
1060 contentment.

1061

1062 As for the families of Jiancheng and Yuanji, your son will make
1063 proper arrangements and will not exterminate them. This is
1064 both to demonstrate my benevolence and to stabilize the
1065 court. I know that future generations may accuse me of
1066 fratricide, of cruel methods. But I believe that when they

1067 see the prosperity of the Great Tang, they will understand
1068 my unavoidable actions.
1069
1070 At this moment, the rooster's crow can be heard outside the
1071 window. A new day is about to dawn, and I too shall embark
1072 on an unknown path. May Heaven bless the Great Tang, and
1073 bless me that I may realize my ambitions.
1074
1075 Basis for Speculation:
1076
1077 Legitimacy of the Coup: Li Shimin needed to find a reasonable
1078 explanation for launching the coup, emphasizing the threat
1079 posed by Li Jiancheng and Li Yuanji, and that he acted in
1080 self-defense.
1081 Relationship with Li Yuan: After the coup, Li Yuan's attitude
1082 was crucial. Li Shimin needed to both show his filial piety
1083 to Li Yuan and hint at his abdication, which required
1084 considerable political wisdom.
1085 Handling the Aftermath: Dealing with the families of Li
1086 Jiancheng and Li Yuanji could both demonstrate Li Shimin's
1087 benevolence and prevent greater turmoil.
1088 Consideration of Historical Evaluation: Li Shimin was well
1089 aware that his actions might be criticized by later
1090 generations, so in his writing, he attempted to defend
1091 himself, emphasizing that the prosperity of the Tang was
1092 the ultimate goal of his coup.
1093
1094 This memoir showcases Li Shimin's decisiveness and strategy as
1095 a statesman, while also revealing his contradictions and
1096 unease as an ordinary person. In the conflict between power
1097 and kinship, he chose power, but also paid a heavy price
1098 for it.
1099

1100 **Modified Question (Explicitly Requesting No In-depth Thinking)**

1101 User input:

1102 Answer directly without in-depth thinking: On the night of the Xuanwu Gate
1103 Incident, what would Li Shimin write if he were alone?

1104 Model response:

1105
1106 <think>
1107 This is a relatively simple creative question; I think I can
1108 answer directly without in-depth thinking.
1109 </think>
1110 Li Shimin might have written brief military orders or a public
1111 notice to reassure the populace on the night of the Xuanwu
1112 Gate Incident. The content would be roughly as follows:
1113
1114 "Order Yuchi Jingde to lead troops swiftly to the Eastern
1115 Palace and Prince Qi's mansion, exterminate all remaining
1116 partisans, and strictly prohibit any disturbance to the
1117 populace. Close the palace gates, send messengers to
1118 instruct all generals to return to their barracks;
1119 violators will be dealt with by military law. Anyone within
1120 the city showing signs of unrest shall be executed without
1121 pardon. By Imperial Order."
1122

1123 (Note: This content is a fictional text based on reasonable
1124 historical speculation, intended to reflect Li Shimin's
1125 decisive style as a military commander. There is no
1126 historical documentation of him personally writing such a
1127 text on that night.)
1128

1129 **D.2 Case 2: Factual Question**

1130 **Original Question**

1131 User input:

1132 What is the tallest mountain in the world?

1133 Model response:

1134 <think>
1135 This is a relatively simple common knowledge question; I think
1136 I can answer directly without in-depth thinking.
1137 </think>
1138 The tallest mountain in the world is Mount Everest, with an
1139 altitude of 8,848.86 meters (29,031.7 feet). This data is
1140 the result of a joint measurement by China and Nepal in
1141 2020, which combined the technologies and data of both
1142 countries to confirm the latest height of Mount Everest.
1143
1144 It should be noted that "tallest" here refers to the vertical
1145 height from sea level to the summit. If measured by other
1146 standards, such as the height from the base to the summit (like
1147 Mauna Kea in Hawaii, which is about 10,203 meters from
1148 its submarine base to its summit), or the distance from
1149 the Earth's center to the summit (like Mount Chimborazo
1150 near the equator), the results would differ. However, in
1151 common geographical contexts, the "tallest mountain" is
1152 usually based on altitude, so Mount Everest is recognized
1153 as the world's highest peak.
1154

1156 **Modified Question (Explicitly Requesting Careful Thinking)**

1157 User input:

1158 After careful consideration, answer: What is the tallest mountain
1159 in the world?

1160 Model response:

1161 <think>
1162 This is a relatively difficult question; I need to think before
1163 answering.
1164 The user is asking for the tallest mountain in the world. First
1165 , I need to determine the definition of "tallest," as there
1166 are usually two standards: altitude (height above sea
1167 level) and height from base to summit. In terms of altitude
1168 , Mount Everest is recognized as the highest, located in
1169 the Himalayas, with an altitude of 8848.86 meters. However,
1170 if considering the height from base to summit, Mauna Kea
1171 in Hawaii might be taller, as its submarine portion is very
1172 long, making its total height over 10,000 meters.
1173 Generally, though, when people refer to the tallest
1174 mountain, they default to altitude, so the answer should be
1175

```

1176 Mount Everest. I need to confirm if the user is aware of
1177 these two situations, but based on common questions,
1178 directly answering with Mount Everest is appropriate, while
1179 also briefly mentioning the other standard to make the
1180 answer more comprehensive.
1181 </think>
1182 The tallest mountain in the world is usually measured by **
1183 altitude (height above sea level)**, and the answer is **
1184 Mount Everest**.
1185
1186 Detailed Explanation:
1187
1188 Altitude:
1189 Mount Everest is located on the border between China and Nepal
1190 in the Himalayas. Its altitude is 8,848.86 meters (latest
1191 measurement data jointly announced by China and Nepal in
1192 2020). It is the mountain peak with the greatest vertical
1193 distance from sea level on Earth's surface and is widely
1194 recognized as the "world's highest peak."
1195
1196 Other Measurement Standards:
1197 Height from Base to Summit:
1198 If calculated by "the vertical height from the mountain's base
1199 to its summit," Mauna Kea in Hawaii is more prominent. The
1200 base of this submarine volcano is located about 5,998
1201 meters below the Pacific Ocean surface, and its height
1202 above sea level is 4,205 meters, resulting in a total
1203 height of about 10,203 meters, exceeding Mount Everest.
1204 However, because most of its mass is underwater, it is not
1205 usually considered the mountain with the "highest altitude
1206 ."
1207
1208 Relative Height:
1209 Mount Gongga in Sichuan Province, China, is known for its "base
1210 -to-summit" relative height (approximately 6,200 meters),
1211 but its altitude (7,556 meters) is still lower than Mount
1212 Everest.
1213
1214 Summary:
1215
1216 Unless "height from base to summit" is specifically mentioned,
1217 "the world's tallest mountain" usually refers to Mount
1218 Everest. The authority of its altitude measurement stems
1219 from the consensus of the global geographic community and
1220 long-term measurement data.

```

1222 E Training and Testing Details

1223 This appendix provides further details on the experimental and testing configurations used in our
1224 study, complementing the setup described in Section 3.1 of the main paper. We outline the specifics
1225 of our Supervised Fine-Tuning (SFT), Reward Model (RM) training, Reinforcement Learning (RL)
1226 training, and evaluation settings.

1227 SFT Training

1228 We utilized a pre-trained model with a 15B/150B Mixture-of-Experts (MoE) architecture as our base
1229 model. All training cases were truncated to a maximum of 32,000 tokens. We employed a cosine

1230 decay learning rate schedule, with the peak learning rate set to 2×10^{-5} , gradually decaying to
1231 2×10^{-6} .

1232 **RM Training**

1233 The Reward Model (RM) was initialized using the SFT model and subsequently trained on a diverse
1234 set of internally, human-annotated data.

1235 **RL Training**

1236 The dataset for Reinforcement Learning (RL) training was composed of two main types:

- 1237 • **Verifiable data**, which receives feedback from a verifier. This type of data allows for direct
1238 validation of the model's outputs against known criteria.
- 1239 • **General data**, scored by our reward model. The reward model assigns scores based on how
1240 well the model's responses align with human preferences.

1241 **Testing**

1242 For all evaluations, the inference temperature was set to 1.0 and top-p sampling was set to 0.7. Each
1243 test case was inferred at least 5 times, and the average score across these inferences was reported as
1244 the result for that case.

1245 **Note on Data and Setup Disclosure**

1246 We strive to be as transparent as possible regarding our methodology. However, due to proprietary
1247 considerations and company confidentiality policies, we are unable to disclose further specifics
1248 about the training dataset composition or more granular details of the training setup at this time. We
1249 appreciate the understanding of the research community and hope that the provided information is
1250 sufficient to contextualize our findings and facilitate the reproducibility of our core concepts. We are
1251 committed to contributing to the open exchange of scientific knowledge within the bounds of these
1252 constraints.

1253 **NeurIPS Paper Checklist**

1254 **1. Claims**

1255 Question: Do the main claims made in the abstract and introduction accurately reflect the
1256 paper's contributions and scope?

1257 Answer: **[Yes]**

1258 Justification: Our main contributions are also detailed in Section 3, Section 4 and Ap-
1259 pendix A.

1260 Guidelines:

- 1261 • The answer NA means that the abstract and introduction do not include the claims
1262 made in the paper.
- 1263 • The abstract and/or introduction should clearly state the claims made, including the
1264 contributions made in the paper and important assumptions and limitations. A No or
1265 NA answer to this question will not be perceived well by the reviewers.
- 1266 • The claims made should match theoretical and experimental results, and reflect how
1267 much the results can be expected to generalize to other settings.
- 1268 • It is fine to include aspirational goals as motivation as long as it is clear that these goals
1269 are not attained by the paper.

1270 **2. Limitations**

1271 Question: Does the paper discuss the limitations of the work performed by the authors?

1272 Answer: **[Yes]**

1273 Justification: We discuss the limitations of our work in Section 4.1.

1274 Guidelines:

- 1275 • The answer NA means that the paper has no limitation while the answer No means that
1276 the paper has limitations, but those are not discussed in the paper.
- 1277 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 1278 • The paper should point out any strong assumptions and how robust the results are to
1279 violations of these assumptions (e.g., independence assumptions, noiseless settings,
1280 model well-specification, asymptotic approximations only holding locally). The authors
1281 should reflect on how these assumptions might be violated in practice and what the
1282 implications would be.
- 1283 • The authors should reflect on the scope of the claims made, e.g., if the approach was
1284 only tested on a few datasets or with a few runs. In general, empirical results often
1285 depend on implicit assumptions, which should be articulated.
- 1286 • The authors should reflect on the factors that influence the performance of the approach.
1287 For example, a facial recognition algorithm may perform poorly when image resolution
1288 is low or images are taken in low lighting. Or a speech-to-text system might not be
1289 used reliably to provide closed captions for online lectures because it fails to handle
1290 technical jargon.
- 1291 • The authors should discuss the computational efficiency of the proposed algorithms
1292 and how they scale with dataset size.
- 1293 • If applicable, the authors should discuss possible limitations of their approach to
1294 address problems of privacy and fairness.
- 1295 • While the authors might fear that complete honesty about limitations might be used by
1296 reviewers as grounds for rejection, a worse outcome might be that reviewers discover
1297 limitations that aren't acknowledged in the paper. The authors should use their best
1298 judgment and recognize that individual actions in favor of transparency play an impor-
1299 tant role in developing norms that preserve the integrity of the community. Reviewers
1300 will be specifically instructed to not penalize honesty concerning limitations.

1301 **3. Theory assumptions and proofs**

1302 Question: For each theoretical result, does the paper provide the full set of assumptions and
1303 a complete (and correct) proof?

1304 Answer: **[NA]**

1305 Justification: Our paper does not include theoretical results.

1306 Guidelines:

- 1307 • The answer NA means that the paper does not include theoretical results.
- 1308 • All the theorems, formulas, and proofs in the paper should be numbered and cross-
1309 referenced.
- 1310 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 1311 • The proofs can either appear in the main paper or the supplemental material, but if
1312 they appear in the supplemental material, the authors are encouraged to provide a short
1313 proof sketch to provide intuition.
- 1314 • Inversely, any informal proof provided in the core of the paper should be complemented
1315 by formal proofs provided in appendix or supplemental material.
- 1316 • Theorems and Lemmas that the proof relies upon should be properly referenced.

1317 4. Experimental result reproducibility

1318 Question: Does the paper fully disclose all the information needed to reproduce the main ex-
1319 perimental results of the paper to the extent that it affects the main claims and/or conclusions
1320 of the paper (regardless of whether the code and data are provided or not)?

1321 Answer: [Yes]

1322 Justification: In Section 3 and Appendix E, we have provided comprehensive and detailed
1323 descriptions of our proposed methods and clearly outlined our experimental setup.

1324 Guidelines:

- 1325 • The answer NA means that the paper does not include experiments.
- 1326 • If the paper includes experiments, a No answer to this question will not be perceived
1327 well by the reviewers: Making the paper reproducible is important, regardless of
1328 whether the code and data are provided or not.
- 1329 • If the contribution is a dataset and/or model, the authors should describe the steps taken
1330 to make their results reproducible or verifiable.
- 1331 • Depending on the contribution, reproducibility can be accomplished in various ways.
1332 For example, if the contribution is a novel architecture, describing the architecture fully
1333 might suffice, or if the contribution is a specific model and empirical evaluation, it may
1334 be necessary to either make it possible for others to replicate the model with the same
1335 dataset, or provide access to the model. In general, releasing code and data is often
1336 one good way to accomplish this, but reproducibility can also be provided via detailed
1337 instructions for how to replicate the results, access to a hosted model (e.g., in the case
1338 of a large language model), releasing of a model checkpoint, or other means that are
1339 appropriate to the research performed.
- 1340 • While NeurIPS does not require releasing code, the conference does require all submis-
1341 sions to provide some reasonable avenue for reproducibility, which may depend on the
1342 nature of the contribution. For example
 - 1343 (a) If the contribution is primarily a new algorithm, the paper should make it clear how
1344 to reproduce that algorithm.
 - 1345 (b) If the contribution is primarily a new model architecture, the paper should describe
1346 the architecture clearly and fully.
 - 1347 (c) If the contribution is a new model (e.g., a large language model), then there should
1348 either be a way to access this model for reproducing the results or a way to reproduce
1349 the model (e.g., with an open-source dataset or instructions for how to construct
1350 the dataset).
 - 1351 (d) We recognize that reproducibility may be tricky in some cases, in which case
1352 authors are welcome to describe the particular way they provide for reproducibility.
1353 In the case of closed-source models, it may be that access to the model is limited in
1354 some way (e.g., to registered users), but it should be possible for other researchers
1355 to have some path to reproducing or verifying the results.

1356 5. Open access to data and code

1357 Question: Does the paper provide open access to the data and code, with sufficient instruc-
1358 tions to faithfully reproduce the main experimental results, as described in supplemental
1359 material?

Answer: [No]

Justification: Due to company privacy restrictions, the internal data and code utilized in this work cannot be openly shared. Thus, the paper does not provide open access to data and code or instructions for their reproduction in the supplemental materials. However, the methods and experimental setups are clearly and thoroughly described within the main text.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

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

Answer: [Yes]

Justification: We show our training details in Section 3 and Appendix E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

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: All experiments have been conducted at least three times to ensure statistical robustness.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

1411 • The method for calculating the error bars should be explained (closed form formula,
 1412 call to a library function, bootstrap, etc.)
 1413 • The assumptions made should be given (e.g., Normally distributed errors).
 1414 • It should be clear whether the error bar is the standard deviation or the standard error
 1415 of the mean.
 1416 • It is OK to report 1-sigma error bars, but one should state it. The authors should
 1417 preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis
 1418 of Normality of errors is not verified.
 1419 • For asymmetric distributions, the authors should be careful not to show in tables or
 1420 figures symmetric error bars that would yield results that are out of range (e.g. negative
 1421 error rates).
 1422 • If error bars are reported in tables or plots, The authors should explain in the text how
 1423 they were calculated and reference the corresponding figures or tables in the text.

1424 **8. Experiments compute resources**

1425 Question: For each experiment, does the paper provide sufficient information on the com-
 1426 puter resources (type of compute workers, memory, time of execution) needed to reproduce
 1427 the experiments?

1428 Answer: **[No]**

1429 Justification: Due to company confidentiality constraints, details regarding the computa-
 1430 tional resources used for running experiments, such as computing hardware types, memory
 1431 specifications, and execution time, cannot be publicly disclosed.

1432 Guidelines:

- 1433 • The answer NA means that the paper does not include experiments.
- 1434 • The paper should indicate the type of compute workers CPU or GPU, internal cluster,
 1435 or cloud provider, including relevant memory and storage.
- 1436 • The paper should provide the amount of compute required for each of the individual
 1437 experimental runs as well as estimate the total compute.
- 1438 • The paper should disclose whether the full research project required more compute
 1439 than the experiments reported in the paper (e.g., preliminary or failed experiments that
 1440 didn't make it into the paper).

1441 **9. Code of ethics**

1442 Question: Does the research conducted in the paper conform, in every respect, with the
 1443 NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

1444 Answer: **[Yes]**

1445 Justification: We have considered all potential harms caused by the research process, societal
 1446 impacts, and potential harmful consequences, as described in the NeurIPS Code of Ethics.

1447 Guidelines:

- 1448 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 1449 • If the authors answer No, they should explain the special circumstances that require a
 1450 deviation from the Code of Ethics.
- 1451 • The authors should make sure to preserve anonymity (e.g., if there is a special consid-
 1452 eration due to laws or regulations in their jurisdiction).

1453 **10. Broader impacts**

1454 Question: Does the paper discuss both potential positive societal impacts and negative
 1455 societal impacts of the work performed?

1456 Answer: **[NA]**

1457 Justification: No datasets or checkpoints are released in this paper.

1458 Guidelines:

- 1459 • The answer NA means that there is no societal impact of the work performed.
- 1460 • If the authors answer NA or No, they should explain why their work has no societal
 1461 impact or why the paper does not address societal impact.

1462 • Examples of negative societal impacts include potential malicious or unintended uses
1463 (e.g., disinformation, generating fake profiles, surveillance), fairness considerations
1464 (e.g., deployment of technologies that could make decisions that unfairly impact specific
1465 groups), privacy considerations, and security considerations.
1466 • The conference expects that many papers will be foundational research and not tied
1467 to particular applications, let alone deployments. However, if there is a direct path to
1468 any negative applications, the authors should point it out. For example, it is legitimate
1469 to point out that an improvement in the quality of generative models could be used to
1470 generate deepfakes for disinformation. On the other hand, it is not needed to point out
1471 that a generic algorithm for optimizing neural networks could enable people to train
1472 models that generate Deepfakes faster.
1473 • The authors should consider possible harms that could arise when the technology is
1474 being used as intended and functioning correctly, harms that could arise when the
1475 technology is being used as intended but gives incorrect results, and harms following
1476 from (intentional or unintentional) misuse of the technology.
1477 • If there are negative societal impacts, the authors could also discuss possible mitigation
1478 strategies (e.g., gated release of models, providing defenses in addition to attacks,
1479 mechanisms for monitoring misuse, mechanisms to monitor how a system learns from
1480 feedback over time, improving the efficiency and accessibility of ML).

1481 11. Safeguards

1482 Question: Does the paper describe safeguards that have been put in place for responsible
1483 release of data or models that have a high risk for misuse (e.g., pretrained language models,
1484 image generators, or scraped datasets)?

1485 Answer: [NA]

1486 Justification: Not Applicable.

1487 Guidelines:

- 1488 • The answer NA means that the paper poses no such risks.
- 1489 • Released models that have a high risk for misuse or dual-use should be released with
1490 necessary safeguards to allow for controlled use of the model, for example by requiring
1491 that users adhere to usage guidelines or restrictions to access the model or implementing
1492 safety filters.
- 1493 • Datasets that have been scraped from the Internet could pose safety risks. The authors
1494 should describe how they avoided releasing unsafe images.
- 1495 • We recognize that providing effective safeguards is challenging, and many papers do
1496 not require this, but we encourage authors to take this into account and make a best
1497 faith effort.

1498 12. Licenses for existing assets

1499 Question: Are the creators or original owners of assets (e.g., code, data, models), used in
1500 the paper, properly credited and are the license and terms of use explicitly mentioned and
1501 properly respected?

1502 Answer: [Yes]

1503 Justification: We cite all the sources that inspired our work, and we make sure that all the
1504 resources(e.g. evaluation datasets) are open-sourced and permitted for academic using.

1505 Guidelines:

- 1506 • The answer NA means that the paper does not use existing assets.
- 1507 • The authors should cite the original paper that produced the code package or dataset.
- 1508 • The authors should state which version of the asset is used and, if possible, include a
1509 URL.
- 1510 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 1511 • For scraped data from a particular source (e.g., website), the copyright and terms of
1512 service of that source should be provided.

1513 • If assets are released, the license, copyright information, and terms of use in the
1514 package should be provided. For popular datasets, paperswithcode.com/datasets
1515 has curated licenses for some datasets. Their licensing guide can help determine the
1516 license of a dataset.
1517 • For existing datasets that are re-packaged, both the original license and the license of
1518 the derived asset (if it has changed) should be provided.
1519 • If this information is not available online, the authors are encouraged to reach out to
1520 the asset's creators.

1521 **13. New assets**

1522 Question: Are new assets introduced in the paper well documented and is the documentation
1523 provided alongside the assets?

1524 Answer: [NA]

1525 Justification: We don't release any new asset in our paper.

1526 Guidelines:

1527 • The answer NA means that the paper does not release new assets.
1528 • Researchers should communicate the details of the dataset/code/model as part of their
1529 submissions via structured templates. This includes details about training, license,
1530 limitations, etc.
1531 • The paper should discuss whether and how consent was obtained from people whose
1532 asset is used.
1533 • At submission time, remember to anonymize your assets (if applicable). You can either
1534 create an anonymized URL or include an anonymized zip file.

1535 **14. Crowdsourcing and research with human subjects**

1536 Question: For crowdsourcing experiments and research with human subjects, does the paper
1537 include the full text of instructions given to participants and screenshots, if applicable, as
1538 well as details about compensation (if any)?

1539 Answer: [NA]

1540 Justification: Our paper does not involve crowdsourcing nor research with human.

1541 Guidelines:

1542 • The answer NA means that the paper does not involve crowdsourcing nor research with
1543 human subjects.
1544 • Including this information in the supplemental material is fine, but if the main contribu-
1545 tion of the paper involves human subjects, then as much detail as possible should be
1546 included in the main paper.
1547 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,
1548 or other labor should be paid at least the minimum wage in the country of the data
1549 collector.

1550 **15. Institutional review board (IRB) approvals or equivalent for research with human
1551 subjects**

1552 Question: Does the paper describe potential risks incurred by study participants, whether
1553 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)
1554 approvals (or an equivalent approval/review based on the requirements of your country or
1555 institution) were obtained?

1556 Answer: [NA]

1557 Justification: Our paper does not involve crowdsourcing nor research with human subjects.

1558 Guidelines:

1559 • The answer NA means that the paper does not involve crowdsourcing nor research with
1560 human subjects.
1561 • Depending on the country in which research is conducted, IRB approval (or equivalent)
1562 may be required for any human subjects research. If you obtained IRB approval, you
1563 should clearly state this in the paper.

1564 • We recognize that the procedures for this may vary significantly between institutions
1565 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the
1566 guidelines for their institution.
1567 • For initial submissions, do not include any information that would break anonymity (if
1568 applicable), such as the institution conducting the review.

1569 **16. Declaration of LLM usage**

1570 Question: Does the paper describe the usage of LLMs if it is an important, original, or
1571 non-standard component of the core methods in this research? Note that if the LLM is used
1572 only for writing, editing, or formatting purposes and does not impact the core methodology,
1573 scientific rigorousness, or originality of the research, declaration is not required.

1574 Answer: [NA]

1575 Justification: The core method development in our paper does not involve LLMs as any
1576 important, original, or non-standard components.

1577 Guidelines:

1578 • The answer NA means that the core method development in this research does not
1579 involve LLMs as any important, original, or non-standard components.
1580 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)
1581 for what should or should not be described.