

# 000 BEYOND TEMPLATES: DYNAMIC ADAPTATION OF REASONING 001 DEMONSTRATIONS VIA FEASIBILITY-AWARE EXPLORATION 002

003 **Anonymous authors**

004 Paper under double-blind review

## 005 ABSTRACT

006 Large language models (LLMs) have shown remarkable reasoning capabilities, yet aligning such  
007 abilities to small language models (SLMs) remains a challenge due to distributional mismatches and  
008 limited model capacity. Existing reasoning datasets, typically designed for powerful LLMs, often lead  
009 to degraded performance when directly applied to weaker models. In this work, we introduce Dynamic  
010 Adaptation of Reasoning Trajectories (DART), a novel data adaptation framework that bridges  
011 the capability gap between expert reasoning trajectories and diverse SLMs. Instead of uniformly  
012 imitating expert steps, DART employs a *selective imitation strategy* guided by step-wise adaptability  
013 estimation via solution simulation. When expert steps surpass the student’s capacity—signaled by  
014 an *imitation gap*—the student autonomously explores alternative reasoning paths, constrained by  
015 outcome consistency. We validate DART across multiple reasoning benchmarks and model scales,  
016 demonstrating that it significantly improves generalization and data efficiency over static fine-tuning.  
017 Our method enhances supervision quality by aligning training signals with the student’s reasoning  
018 capabilities, offering a scalable solution for reasoning alignment in resource-constrained models.  
019

## 020 1 INTRODUCTION

021 Large language models (LLMs) have recently achieved remarkable performance in complex reasoning tasks such as  
022 mathematics and programming (OpenAI, 2024; Shao et al., 2024). A key insight from recent work (Zhou et al., 2024;  
023 Yue et al., 2024; Ye et al., 2025) is that small, high-quality instruction datasets are surprisingly effective at eliciting  
024 sophisticated reasoning abilities in large models. This discovery challenges traditional beliefs (Li et al., 2024; Yu et al.,  
025 2024) that complex cognitive skills necessarily require massive supervised fine-tuning, opening promising avenues for  
026 data-efficient model alignment.

027 Despite the remarkable effectiveness of small, high-quality instruction datasets in eliciting sophisticated reasoning,  
028 mainstream approaches (Zhou et al., 2024; Ye et al., 2025; Muennighoff et al., 2025) remain reliant on **static, pre-  
029 collected** reasoning datasets. While effective in controlled environments, these datasets struggle to generalize across  
030 heterogeneous pretraining distributions, particularly for small language models (SLMs) with diverse training data and  
031 limited reasoning capabilities (Xu et al., 2024; Yeo et al., 2025). Disparities in model scale, reasoning proficiency, and  
032 training history exacerbate distributional mismatches, significantly hindering the activation of reasoning skills.

033 To address these challenges, we introduce **Dynamic Adaptation of Reasoning Trajectories (DART)**, a novel data  
034 adaptation framework designed to bridge the distribution gap between static reasoning datasets and diverse SLMs.  
035 Instead of enforcing uniform imitation of expert demonstrations, DART introduces a *selective imitation* strategy guided  
036 by *imitation feasibility estimate*. For each step provided by the expert, DART dynamically assesses the likelihood that  
037 the student model can successfully complete the reasoning process conditioned on adopting that step. When imitation is  
038 deemed infeasible, the student autonomously explores alternative trajectories while maintaining the consistency of the  
039 outcome with the objective of the original task. This approach enables DART to flexibly adapt high-quality reasoning  
040 datasets to heterogeneous model populations, significantly improving reasoning elicitation under distribution shift.

041 In summary, our contributions are as follows.

- 042 • We identify the critical limitations of applying static curated reasoning datasets to diverse small language  
043 models and propose **DART**, a novel framework for adapted reasoning data guided by imitation feasibility.
- 044 • We introduce a Monte Carlo simulation-based method to estimate the feasibility of imitation per step, allowing  
045 selective supervision tailored to the student model capabilities.
- 046 • We develop an autonomous exploration mechanism that allows models to recover from infeasible supervision  
047 points, generating outcome-consistent alternative reasoning paths.

054

- Through extensive experiments across different model scales and benchmarks, we demonstrate that DART

055 substantially improves reasoning performance over static fine-tuning, achieving superior data efficiency and

056 generalization.

057

058

## 2 PRELIMINARIES AND LIMITATIONS OF SUPERVISED IMITATION ON EXPERT 059 TRAJECTORIES

060

061

### 2.1 PROBLEM DEFINITION: REASONING CAPABILITY ELICITATION VIA MINIMAL DEMONSTRATIONS

062

063 We define the reasoning elicitation problem in the context of large language models (LLMs) with latent pre-trained  
064 knowledge. Let  $\mathcal{Q}$  denote the space of reasoning problems,  $\mathcal{A}$  the space of answers, and  $\mathcal{R}$  the space of reasoning  
065 chains, where each  $r \in \mathcal{R}$  is a sequence of logical steps  $r = \{s_1, s_2, \dots, s_n\}$ .

066

067 The goal is to learn a reasoning function:

068

$$f : \mathcal{Q} \rightarrow \mathcal{R} \times \mathcal{A} \quad (1)$$

069

070 so that, given a question  $q \in \mathcal{Q}$ , the model generates a logically valid reasoning chain  $r \in \mathcal{R}$  and a verifiable final  
071 answer  $a \in \mathcal{A}$ .

072 Prior work (e.g., (Ye et al., 2025), (Muennighoff et al., 2025)) suggests that reasoning competence in large language  
073 models (LLMs) can be elicited not by scale alone, but a small set of carefully crafted demonstrations that expose  
074 the underlying cognitive structure of reasoning. This paradigm assumes that latent reasoning skills embedded within  
075 pretrained models can be activated through appropriately designed prompts in the form of explicit multi-step exemplars.

076 Let  $\mathcal{D} = \{(q_i, r_i, a_i)\}_{i=1}^N$  represent a compact yet high-quality dataset ( $N \ll |\mathcal{Q}|$ ), where each tuple contains a question  
077  $q_i$ , a structured reasoning chain  $r_i$ , and its corresponding answer  $a_i$ . Each  $r_i$  serves as a **cognitive template**—an  
078 interpretable, step-wise reasoning demonstration designed to guide the model through logical steps with intermediate  
079 verification. Instead of introducing new knowledge, these templates activate the model’s latent reasoning capabilities by  
080 leveraging structured prompting (Wei et al., 2022; Zhou et al., 2024; Ye et al., 2025).

081

082

### 2.2 LIMITATIONS OF SUPERVISED IMITATION ON EXPERT DEMONSTRATIONS

083

084 Despite its pedagogical appeal, supervised imitation over expert demonstrations exhibits critical limitations when  
085 applied to LLMs with diverse capacity levels.

086 This paradigm (Wei et al., 2022; Ye et al., 2025) assumes that the model possesses sufficient latent competence to  
087 internalize and reproduce the reasoning trajectory in each template. In practice, this assumption frequently fails. A  
088 template  $r_i$  may (i) over-challenge the model by invoking reasoning procedures not encoded in its weights, or (ii)  
089 misalign with the model’s inductive biases, causing representational mismatch. We define a reasoning failure event  $\mathcal{F}$   
090 as the inability of the model to emulate the intended behavior given an input-template pair:

091

$$\mathcal{F}(f; q, r, a) = \mathbb{I}[f(q) \not\approx (r, a)] \quad (2)$$

092

093 where  $\mathbb{I}[\cdot]$  is the indicator function. Such failures may arise from superficial imitation, incomplete reasoning chains, or  
094 insufficient justification for the final answer.

095 Compounding this challenge is the substantial cost associated with constructing template datasets  $\mathcal{D}$  that satisfy the  
096 Cognitive Template Demonstration criterion. Such templates demand meticulous logical decomposition, intermediate  
097 verification, and fine-grained pedagogical design. Furthermore, a template crafted for a specific model often fails  
098 to generalize to others due to differences in scale, pretraining corpus, or architectural inductive biases, resulting in  
099 distributional shifts. As highlighted in prior work on imitation learning (Pomerleau, 1991; Ross et al., 2011), relying on  
100 static datasets for training can lead to a distribution mismatch between the output sequences encountered during training  
101 and those generated auto-regressively by the student at inference time, undermining generalization and robustness.

102

103 **The Need for Imitation Feasibility-Aware Adaptation.** These limitations highlight the inadequacy of static demon-  
104 strations in addressing the diversity of model behaviors. We argue for a dynamic grounding mechanism that aligns  
105 template presentation with the target model’s internal capacity and abstraction level. Rather than treating  $\mathcal{D}$  as fixed  
106 input, the elicitation process should adaptively align the demonstrated reasoning path with the model’s own preferred or  
107 accessible inference trajectories, potentially reformulating how the reasoning unfolds to match internal representations.  
This motivates our central question:

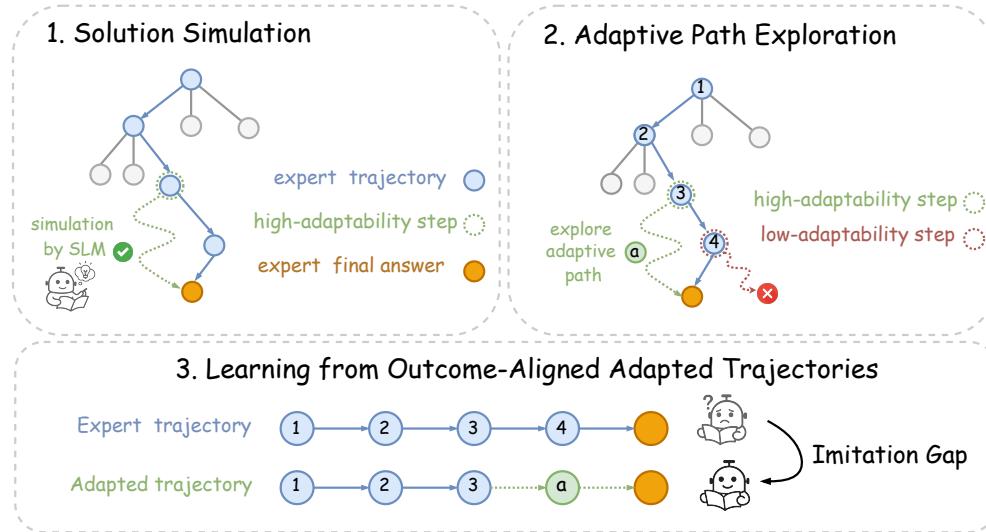


Figure 1: Overview of the DART framework.

126  
127  
128  
129 *Can we design a dynamic adaptation mechanism that reliably anchors cognitive templates in model-  
130 specific latent space, enabling scalable and robust reasoning?*

131 In the following section, we instantiate this motivation via our proposed framework — **Dynamic Adaptation of**  
132 **Reasoning Trajectories (DART)**.

### 134 3 METHODOLOGY

135  
136 In this section, we propose **Dynamic Adaptation of Reasoning Trajectories (DART)**, a capability-aware adaptation  
137 framework designed to align expert-level reasoning data with the capacity of small language models (SLMs). Instead  
138 of statically mimicking expert trajectories from the elicitation template set, DART introduces a selective imitation  
139 mechanism that dynamically adapts supervision signals based on the model’s reasoning proficiency. The framework  
140 comprises three key components: (1) step-wise adaptability estimation via solution simulation (Section 3.1), (2)  
141 imitation gap detection and adaptive path exploration (Section 3.2), and (3) learning from outcome-aligned adapted  
142 trajectories (Section 3.3). Figure 1 provides an overview of the pipeline, while Algorithm 1 formalizes the procedure,  
143 detailing how these components are operationalized to generate and train on adapted trajectories.

#### 144 3.1 STEP-WISE ADAPTABILITY ESTIMATION VIA SOLUTION SIMULATION

145 To determine whether a given expert step is suitable for imitation, we introduce the concept of **adaptability**: the  
146 likelihood that a student model can reach the correct answer when conditioned on that step. This evaluation is conducted  
147 via solution simulation—akin to Monte Carlo Tree Search (Kocsis & Szepesvári, 2006; Silver et al., 2016; Świechowski  
148 et al., 2023)—by rolling out multiple completions from partially constructed trajectories that incorporate the candidate  
149 step.

150 Let  $s_{<t} = \{s_0, s_1, \dots, s_{t-1}\}$  be the prefix of expert steps, and  $s_t$  the candidate step under evaluation. The adaptability  
151 score  $f_t$  is computed as:

$$152 \quad f_t = Q(s_{<t}, s_t) = \frac{1}{N_{\text{sim}}} \sum_{i=1}^{N_{\text{sim}}} \mathbb{I}(a_i^{\text{final}} = a^*) \quad (3)$$

153  
154 where  $N_{\text{sim}}$  denotes the total number of *rollouts* performed for each candidate step  $s_t$ , with each rollout simulating a  
155 complete reasoning trajectory conditioned on the prefix  $s_{<t}$  and the adoption of step  $s_t$ .

156 Empirically observed patterns (see Section 5.1) suggest that adaptability tends to rise in the early stages of expert  
157 trajectories, but drops sharply beyond a certain point. This non-monotonic behavior motivates our definition of the

---

162 **Algorithm 1** DART: Dynamic Adaptation of Reasoning Trajectories

---

163 **Require:** Expert trajectory  $\tau_{\text{expert}} = \{s_0, s_1, \dots, s_T\}$ , student model  $\pi_{\text{student}}$ , ground-truth answer  $a^*$ , adaptation  
 164 simulation count  $N_{\text{sim}}$   
 165 1: Initialize prefix  $\tau_{\text{prefix}} \leftarrow \emptyset$ ; adaptability scores  $\mathcal{F} \leftarrow []$   
 166 2: **for**  $t = 0$  to  $T$  **do**  
 167 3:   Compute adaptability score  $f_t \leftarrow Q(s_{<t}, s_t)$  ▷ See Eq. equation 3  
 168 4:   Append  $f_t$  to  $\mathcal{F}$   
 169 5: Find  $t_{\text{peak}}$  where  $f_t$  attains its local maximum;  
 170  
 171    $t_{\text{gap}} \leftarrow \min \{t > t_{\text{peak}} \mid f_t < f_{t_{\text{peak}}} - \epsilon\}$  where  $\epsilon > 0$  defines a significant drop threshold  
 172  
 173 6: Truncate expert prefix:  $\tau_{\text{prefix}} \leftarrow \{s_0, \dots, s_{t_{\text{gap}}-1}\}$   
 174 7: Initialize adapted trajectory:  $\tau_{\text{adapt}} \leftarrow \tau_{\text{prefix}}$   
 175 8: **while** not terminal **do**  
 176 9:   Sample next step  $s' \sim \pi_{\text{student}}(\cdot \mid \tau_{\text{adapt}})$   
 177 10:   Append  $s'$  to  $\tau_{\text{adapt}}$   
 178 11: **if**  $\mathcal{O}_{\text{adapt}}(\tau_{\text{adapt}}) = \mathcal{O}_{\text{expert}}(\tau_{\text{expert}})$  **then** ▷ See Eq. equation 6  
 179 12:   Retain  $\tau_{\text{adapt}}$  for distillation  
 180 13: **else**  
 181 14:   Discard  $\tau_{\text{adapt}}$ 

---

182  
 183 **imitation gap**, a regime in which continued imitation becomes counterproductive due to the increasing complexity of  
 184 the remaining expert steps.  
 185

186 3.2 ADAPTIVE PATH EXPLORATION  
 187

188 To avoid overfitting to brittle expert demonstrations, we monitor the *adaptability score* throughout the trajectory and  
 189 halt imitation once a significant drop is detected (see Equation 3). Motivated by the need to overcome low-adaptability  
 190 segments that may hinder generalization, DART transitions to autonomous rollout beyond the gap, generating a  
 191 continuation from the last high-adaptability prefix:

$$\tau_{\text{adapt}} = (s_0, s_1, \dots, s_{t-1}, s'_t, s'_{t+1}, \dots, s'_T), \quad (4)$$

192 where  $s'_t, \dots, s'_T$  are student-generated reasoning steps. Inspired by outcome-based learning strategies (DeepSeek-AI  
 193 et al., 2025), we do not constrain this trajectory to mimic the expert’s form. Instead, we enforce an *outcome consistency*  
 194 constraint to ensure semantic alignment, as described in Eq. equation 5, as we observe that process supervision (Lightman  
 195 et al., 2024; Zhang et al., 2025), such as via a Process Reward Model (PRM), often encounters inherent ambiguities and  
 196 standardization challenges in practice.  
 197

$$C(\tau_{\text{adapt}}, \tau_{\text{expert}}) = \begin{cases} 1, & \text{if } \mathcal{O}(\tau_{\text{adapt}}) = \mathcal{O}(\tau_{\text{expert}}), \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

200 Here,  $C \in \{0, 1\}$  denotes task-level agreement, with  $\mathcal{O}(\cdot)$  representing the final answer obtained by executing a  
 201 reasoning path. Specifically,  $\mathcal{O}(\tau_{\text{expert}})$  refers to the outcome of the expert demonstration, while  $\mathcal{O}(\tau_{\text{adapt}})$  captures the  
 202 result of the student’s adapted trajectory. The constraint  $\mathcal{O}(\tau_{\text{adapt}}) = \mathcal{O}(\tau_{\text{expert}})$  ensures that, although the reasoning  
 203 paths may differ, their semantic outcomes are equivalent. This outcome consistency criterion allows the student to  
 204 depart from brittle expert traces while preserving task correctness.  
 205

206 This strategy empowers the student model to develop its own reasoning strategies beyond segments with low adapt-  
 207 ability, guided solely by the correctness of the final outcome. By anchoring supervision at the outcome level rather  
 208 than mimicking intermediate steps, we alleviate the brittleness of process-level imitation. This encourages robust  
 209 generalization, reduces reliance on ambiguous or inconsistent expert demonstrations, and aligns with the broader goal  
 210 of enabling flexible yet goal-directed reasoning.  
 211

212 3.3 LEARNING FROM OUTCOME-ALIGNED ADAPTED TRAJECTORIES  
 213

214 To effectively activate the student model’s own reasoning ability, we apply a standard cross-entropy loss on the outcome-  
 215 aligned adapted trajectories generated during autonomous exploration. This training objective encourages the model to  
 reinforce reasoning patterns that are not only aligned with the task goal but also feasible under its own capacity.  
 216

216 Table 1: Main results (%) on LIMO and Math-QwQ-32B across adaptation strategies and model sizes. **Static** overfits to  
 217 noisy data, while **Adaptation-Full** improves results through exploration and filtering of low-adaptability segments.  
 218

219 <b>Dataset</b>	220 <b>Method</b>	221 <b>GSM8K</b>	222 <b>GaoKao</b>	223 <b>Olympiad</b> 224 <b>Bench</b>	225 <b>College</b> 226 <b>Math</b>	227 <b>MMLU</b> 228 <b>STEM</b>	229 <b>Avg.</b>
222 <b>Qwen2.5-0.5B-Instruct</b>							
223      -	224      No-Tuning	225      49.1	226      30.4	227      9.3	228      28.9	229      36.7	230      30.9
231      Math-QwQ-32B	232      Static	233      39.8 <b>-9.3</b>	234      20.5 <b>-9.9</b>	235      5.9 <b>-3.4</b>	236      17.3 <b>-11.6</b>	237      27.9 <b>-8.8</b>	238      22.3 <b>-8.6</b>
239      Math-QwQ-32B	240 <b>Adaptation-Full</b>	241      49.6 <b>+0.5</b>	242      30.9 <b>+0.5</b>	243      9.3 <b>+0.0</b>	244      27.5 <b>-1.4</b>	245      37.5 <b>+0.8</b>	246      31.0 <b>+0.1</b>
247      LIMO	248      Static	249      49.6 <b>+0.5</b>	250      26.8 <b>-3.6</b>	251      7.7 <b>-1.6</b>	252      27.3 <b>-1.6</b>	253      32.9 <b>-3.8</b>	254      28.9 <b>-2.0</b>
255      LIMO	256 <b>Adaptation-Full</b>	257      52.2 <b>+3.1</b>	258      32.5 <b>+2.1</b>	259      9.8 <b>+0.5</b>	260      29.1 <b>+0.2</b>	261      37.2 <b>+0.5</b>	262      32.2 <b>+1.3</b>
263 <b>Qwen2.5-3B-Instruct</b>							
264      -	265      No-Tuning	266      87.0	267      56.6	268      27.3	269      39.9	270      47.6	271      51.7
272      Math-QwQ-32B	273      Static	274      82.0 <b>-5.0</b>	275      46.2 <b>-10.4</b>	276      20.1 <b>-7.2</b>	277      35.7 <b>-4.2</b>	278      51.3 <b>+3.7</b>	279      47.1 <b>-4.6</b>
280      Math-QwQ-32B	281 <b>Adaptation-Full</b>	282      86.6 <b>-0.4</b>	283      57.9 <b>+1.3</b>	284      29.0 <b>+1.7</b>	285      44.4 <b>+4.5</b>	286      50.6 <b>+3.0</b>	287      53.7 <b>+2.0</b>
288      LIMO	289      Static	290      85.4 <b>-1.6</b>	291      53.8 <b>-2.8</b>	292      25.2 <b>-2.1</b>	293      41.8 <b>+1.9</b>	294      54.8 <b>+7.2</b>	295      52.2 <b>+0.5</b>
296      LIMO	297 <b>Adaptation-Full</b>	298      87.2 <b>+0.2</b>	299      59.5 <b>+2.9</b>	300      30.4 <b>+3.1</b>	301      43.9 <b>+4.0</b>	302      62.7 <b>+15.1</b>	303      56.7 <b>+5.0</b>
304 <b>LLaMA-3B-Instruct</b>							
305      -	306      No-Tuning	307      38.4	308      21.3	309      10.7	310      16.0	311      48.0	312      26.9
313      Math-QwQ-32B	314      Static	315      67.7 <b>+29.3</b>	316      30.4 <b>+9.1</b>	317      10.2 <b>-0.5</b>	318      21.1 <b>+5.1</b>	319      39.2 <b>-8.8</b>	320      33.7 <b>+6.8</b>
321      Math-QwQ-32B	322 <b>Adaptation-Full</b>	323      72.3 <b>+33.9</b>	324      34.0 <b>+12.7</b>	325      10.2 <b>-0.5</b>	326      21.8 <b>+5.8</b>	327      38.3 <b>-9.7</b>	328      35.3 <b>+8.4</b>
329      LIMO	330      Static	331      28.1 <b>-10.3</b>	332      15.1 <b>-6.2</b>	333      3.6 <b>-7.1</b>	334      11.4 <b>-4.6</b>	335      48.6 <b>+0.6</b>	336      21.4 <b>-5.5</b>
337      LIMO	338 <b>Adaptation-Full</b>	339      47.2 <b>+8.8</b>	340      24.9 <b>+3.6</b>	341      6.4 <b>-4.3</b>	342      18.1 <b>+2.1</b>	343      47.2 <b>-0.8</b>	344      28.8 <b>+1.9</b>

247 Training proceeds by distilling the adapted trajectory  $\tau_{\text{adapt}}$  using a standard cross-entropy loss (Kim & Rush, 2016;  
 248 Bengio et al., 2003):

$$249 \quad L_{\text{DART}} = - \sum_{t=1}^T \mathbb{E}_{(s_{<t}, a_t) \sim \tau_{\text{adapt}}} [\log \pi_{\text{student}}(a_t \mid s_{<t})] \quad (6)$$

252 Here,  $s_{<t} = \{s_0, \dots, s_{t-1}\}$  denotes the contextual prefix consisting of all prior reasoning steps up to time  $t$ , and  $a_t$  is  
 253 the corresponding next-step decision. This loss encourages the student model  $\pi_{\text{student}}$  to maximize the likelihood of  
 254 producing  $a_t$  when conditioned on its own reasoning history.

255 By learning from outcome-aligned yet model-compatible trajectories, DART provides high-quality supervision that  
 256 reflects the student’s actual competence. This approach decouples the training signal from rigid trajectory matching,  
 257 improving both robustness and scalability across models with varying capacity.

## 259      4 EXPERIMENTS

261 We evaluate DART across a series of mathematical reasoning benchmarks to assess its effectiveness in adapting expert  
 262 data to student models of varying capacities.

### 264      4.1 EXPERIMENTAL SETUP

266 **Adaptation Datasets.** We conduct adaptation experiments using two datasets. (1) **LIMO** dataset (Ye et al., 2025),  
 267 a curated set of 817 high-quality math reasoning examples with multi-step CoT demonstrations tailored. We use the  
 268 official filtered release<sup>1</sup>. (2) The **Math-QwQ-32B** dataset is a synthetic dataset derived from the MATH benchmark

269 <sup>1</sup><https://huggingface.co/GAIR/LIMO>

(Hendrycks et al., 2021b), where the Qwen/QwQ-32B-Preview model<sup>2</sup> generates long-form Chain-of-Thought (CoT) solutions for 5,383 problems in the training subset.

**Adaptation Strategies.** We evaluate three adaptation strategies to disentangle the effects of selective imitation and adaptive exploration. (1) *No-Tuning* denotes direct zero-shot evaluation. (2) *Static* reflects standard offline supervised fine-tuning on the full set of expert trajectories, without any filtering or adaptability mechanism. (3) The *Adaptation-Full* strategy represents the complete DART pipeline, integrating imitation gap detection with outcome-consistent student exploration. This approach empowers the model to autonomously explore alternative reasoning paths when expert imitation becomes unreliable. If the model can't find a suitable alternative path, it discards that expert example. We evaluate our method on Qwen2.5-Instruct models at 0.5B<sup>3</sup>, 1.5B<sup>4</sup>, 3B<sup>5</sup>, and LLaMA-3B-Instruct<sup>6</sup> scales, covering a diverse range of SLMs.

**Benchmark Tasks.** We evaluate DART on seven diverse benchmarks encompassing a broad spectrum of mathematical reasoning. These include GSM8K (Cobbe et al., 2021), covering grade-school to competition-level problems. To assess linguistic and cultural generalization, we incorporate GaoKao 2023 En (Liao et al., 2024), a Chinese national exam benchmark. OlympiadBench (He et al., 2024) features high-difficulty, compositional problems from international math competitions. College Math (Tang et al., 2024) probes undergraduate-level topics in calculus, algebra, and discrete math. MMLU-STEM (Hendrycks et al., 2021a) evaluates STEM-focused reasoning breadth. Overall adaptation is quantified by the arithmetic mean (Avg.) across all benchmarks.

**Training and Model Selection** To ensure the reliability of experimental results, we conducted systematic training and model selection for all models. For both the *Static* and *Adaptation-Full* strategies, we trained models at the 0.5B, 1.5B, and 3B scales for 15 epochs, saving a model checkpoint at the end of each epoch, resulting in 15 checkpoints per model. These checkpoints were evaluated on the validation sets, and the model with the best performance was selected as the final model. The training parameter settings were consistent with LIMO (Ye et al., 2025). All experiments are conducted on the same NVIDIA A100 GPU infrastructure. Additional implementation details, including configurations and setups, are provided in Appendix A.

## 4.2 MAIN RESULTS

Table 1 reports performance across five mathematical reasoning benchmarks, demonstrating the effectiveness of the proposed DART framework in aligning expert reasoning with the capabilities of small language models (SLMs).

**Static Results** Static, which rigidly imitates expert trajectories without adaptation, exhibits clear limitations. On Qwen2.5-0.5B, it decreases accuracy by **8.6 points** on *Math-QwQ-32B* and by **2.0 points** on *LIMO*. On Qwen2.5-3B, Static reduces accuracy on *Math-QwQ-32B* by **4.6 points**. A similar degradation is observed on LLaMA-3B, where accuracy on *LIMO* drops by **5.5 points**. These results are obtained under our careful training and model selection protocol (see Sec. 4.1), ensuring that the observed degradation is not caused by insufficient training but rather reflects the inherent limitations of the Static strategy. These findings indicate that imitating expert demonstrations without adaptation not only constrains small models but can also undermine the performance of larger ones.

**Adaptation-Full Results** Adaptation-Full shows robust improvements over No-Tuning, consistently enhancing performance across datasets and model scales. For example, on the Qwen2.5-0.5B model, Adaptation-Full improves the average accuracy on *LIMO* by +1.3 points, while on the larger Qwen2.5-3B, it yields gains of +2.0 and +5.0 points on *Math-QwQ-32B* and *LIMO*, respectively. The effect is even more pronounced on LLaMA-3B, where Adaptation-Full boosts *Math-QwQ-32B* by +8.4 points and *LIMO* by +1.9 points. On average, Adaptation-Full achieves a **+4.9 point** improvement over No-Tuning, demonstrating its effectiveness in aligning reasoning trajectories with model capacity while maintaining stability across different architectures.

## 5 ANALYSIS

We further analyze the internal mechanisms of DART, aiming to understand why selective imitation and autonomous exploration improve reasoning capabilities.

<sup>2</sup><https://huggingface.co/Qwen/QwQ-32B-Preview>

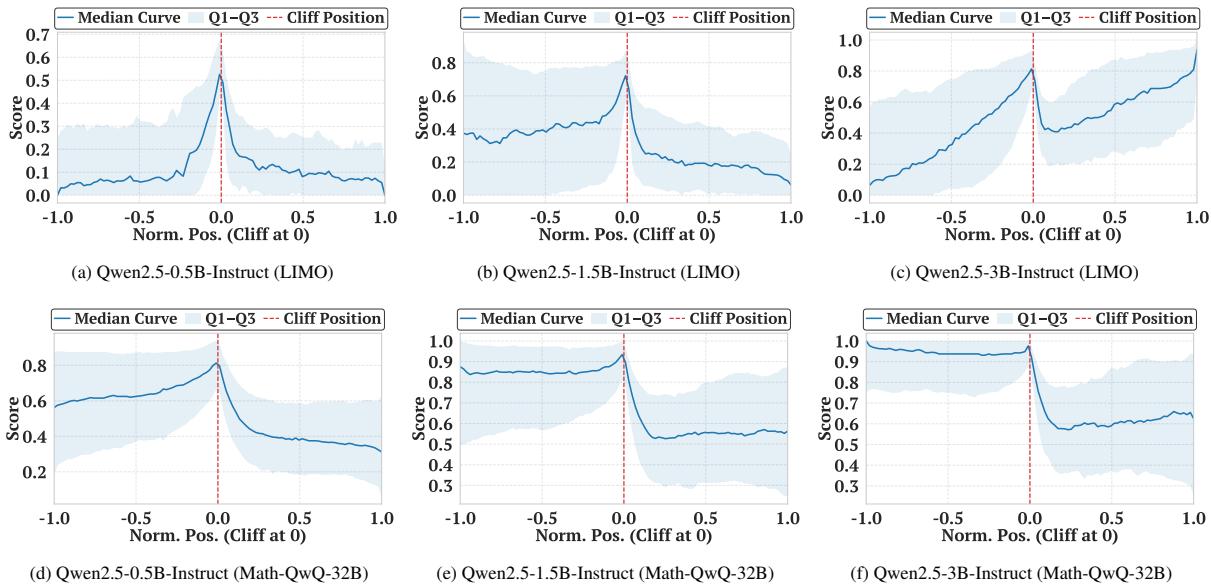
<sup>3</sup><https://huggingface.co/Qwen/Qwen2.5-0.5B-Instruct>

<sup>4</sup><https://huggingface.co/Qwen/Qwen2.5-1.5B-Instruct>

<sup>5</sup><https://huggingface.co/Qwen/Qwen2.5-3B-Instruct>

<sup>6</sup><https://huggingface.co/meta-llama/LLaMA-3B-Instruct>

324  
 325 Figure 2: Step-wise adaptability scores across expert trajectories for Qwen2.5-Instruct student models of varying  
 326 sizes (0.5B, 1.5B, 3B parameters) under LIMO (top row) and Math-QwQ-32B dataset (bottom row) supervision. The  
 327 emergence of the **Imitation Gap** is evident: initial steps yield positive adaptation, but continued step-by-step imitation  
 328 can become harmful.



### 5.1 STEP-WISE ADAPTABILITY REVEALS THE EMERGENCE OF THE IMITATION GAP

To empirically validate the *imitation gap* hypothesis introduced in Section 3, we estimate the step-wise adaptability scores of Qwen2.5-Instruct student models across three parameter scales (0.5B, 1.5B, 3B) on two reasoning datasets (LIMO and Math-QwQ-32B). Each adaptability score quantifies the model’s average probability of reaching the correct final answer when conditioned on imitating a given intermediate step from the expert trajectory. To remove any confounding effect of unequal trajectory lengths, we length-normalize every trace and register the detected cliff at  $x = 0$ . The curve shows the median adaptability score, and the shaded area the interquartile range (Q1–Q3).

As shown in Figure 2, these curves reveal a consistent behavioral pattern: early in the reasoning path, student models exhibit increasing adaptability as they benefit from following expert steps. However, beyond a certain point, adaptability scores sharply decline—signaling that the student has encountered steps that exceed its reasoning capacity, leading to degraded rollout completions and a collapse in trajectory success.

This non-monotonic pattern reveals the **imitation gap**—a critical region where student models falter due to misalignment between their capabilities and the expert’s step distribution. This misalignment arises from distributional discrepancies, where expert trajectories include reasoning patterns outside the student’s abstraction space. Consequently, continued imitation in this zone not only fails to benefit learning but actively impairs performance. This phenomenon underscores our central argument: effective reasoning supervision must be dynamically aligned with model-specific capabilities, as realized in our DART framework. To elucidate the imitation gap’s impact on adaptability score decline and reasoning performance, we present a case study on the LIMO dataset, pinpointing its onset in a complex reasoning task through Qwen2.5-3B-Instruct’s step-wise adaptability scores (see Table 2).

### 5.2 IMPACT OF SEARCH RESTRICTION ON ADAPTATION STRATEGIES

To evaluate the impact of adaptation strategies without autonomous search, we designed two variants: **Adaptation-First** and **Adaptation-Gap**. **Adaptation-First** halts imitation once a feasible solution state is detected, whereas **Adaptation-Gap** monitors adaptability scores and terminates imitation when sharp declines occur, as described in Section 5.1. Table 3 presents the evaluation results on Math-QwQ-32B for 1.5B and 3B Qwen2.5-Instruct models. Both strategies exhibit performance degradation compared to **Adaptation-Full**, highlighting the critical role of autonomous search for recovery in complex reasoning paths. Notably, **Adaptation-Gap** consistently outperforms **Adaptation-First** across all benchmarks, with significant accuracy gains in the average performance (**35.8%** vs. **30.3%** for 1.5B and

378 Table 2: Case Study on the LIMO Dataset: Solutions are generated by Qwen2.5-3B-Instruct. Step-wise adaptability  
 379 scores are reported for each reasoning step, with steps subsequent to the Imitation Gap (starting from Step 151)  
 380 highlighted in gray.

382 Q: The vertices of a regular nonagon (9-sided polygon) are to be labeled with the digits 1 through 9 in such a way that	383 the sum of the numbers on every three consecutive vertices is a multiple of 3. Two acceptable arrangements are	384 considered to be indistinguishable if one can be obtained from the other by rotating the nonagon in the plane. Find the	385 number of distinguishable acceptable arrangements.	Scores
386	387	388	389	386
387 Step 145: Therefore, except in the case of $d=1$ , which would require all digits being same, but which is impossible with	388 unique digits, otherwise, no fix. Therefore, fixed labelings only for identity, which are all our 1296 acceptable labelings.	389	390	0.750
389 Step 146: Therefore, the number of distinct arrangements is 1296 divided by 9, which is 144.	390	391	392	0.938
390 Step 147: So, is 144 the answer? But wait, let me cross-verify.	391	392	393	0.875
391 Step 148: Wait, though: in our prior thought process, we had 6 residue assignments, each giving 216 colorings. So $6 \times$	392	393	394	0.875
392 $216 = 1296$ .	393	394	395	1.000
393 Step 149: If we consider the condition that rotating must not create overcounts, then perhaps we had overcounted by	394	395	396	0.938
394 factor 9, so we have to divide by 9. $1296 / 9 = 144$ .	395	396	397	0.938
395 Step 150: Alternatively, to think in terms of necklace counting: the number of distinct necklaces with beads labeled 1-9,	396	397	398	0.438
396 under rotation, with a constraint that any three consecutive numbers sum to multiple of 3. So, the formula, via Burnside,	397	398	399	0.688
397 gives 144. Does that make sense?	398	399	400	0.938
398 Step 151: But, another route. Let me think, maybe I have miscalculated for fixed labelings. Wait, notice that the fixed	399	400	401	0.438
399 labelings under rotation by $k=3$ positions?	400	401	402	0.688
400 Step 152: Wait, because $d = \gcd(3,9) = 3$ , so the cycles decompose into 3 cycles, each of length 3. So, fixed labelings	401	402	403	0.938
401 here would have to assign the same digit to each cycle, but since labels must be unique, is it possible?	402	403	404	0.938
402 Step 153: Wait, meaning that for each cycle (1,4,7), (2,5,8), (3,6,9); all three digits in cycle (1,4,7) must be the same,	403	404	405	0.938
403 similarly for others, which is impossible with unique digits.	404	405	406	0.938

405 Table 3: Accuracy (%) on the Math-QwQ-32B dataset for Qwen2.5-1.5B-Instruct and Qwen2.5-3B-Instruct models  
 406 under different adaptation strategies. **Adaptation-First** performs early stopping at feasible solution states, while  
 407 **Adaptation-Gap** selectively truncates imitation paths based on adaptability declines. **Adaptation-Full** integrates  
 408 autonomous search, achieving the highest performance across benchmarks. **Bold** values indicate the best results in each  
 409 group.

411 Model	412 Method	GSM8K	GaoKao	Olympiad Bench	College Math	MMLU STEM	Avg.
414 Qwen2.5-1.5B-Instruct	Adaptation-First	41.2	33.0	10.7	24.9	41.7	30.3
	Adaptation-Gap	60.0	34.5	13.6	29.7	41.3	35.8
	<b>Adaptation-Full</b>	<b>74.2</b>	<b>48.6</b>	<b>19.6</b>	<b>39.4</b>	<b>57.7</b>	<b>47.9</b>
417 Qwen2.5-3B-Instruct	Adaptation-First	36.7	32.7	12.9	26.1	46.8	31.0
	Adaptation-Gap	77.9	43.4	18.4	36.0	44.0	43.9
	<b>Adaptation-Full</b>	<b>86.6</b>	<b>57.9</b>	<b>29.0</b>	<b>44.4</b>	<b>50.6</b>	<b>53.7</b>

421 **43.9% vs. 31.0%** for 3B). This improvement stems from its capacity-aware truncation, which effectively filters out  
 422 low-adaptability segments, preventing error propagation and enhancing stability.

### 424 5.3 CAPACITY-ALIGNED LEXICAL DYNAMICS UNDER ADAPTATION

425 To investigate how DART reshapes student model behavior at different scales, we analyze keyword frequency changes  
 426 between static and adapted dataset. Table 4 lists the top 20 tokens with the largest shifts in the first sentence of each  
 427 reasoning step for the Qwen2.5-Instruct series 0.5B, 1.5B, and 3B models.

428 Adaptation reduces exploratory terms like *but*, *wait*, and *alternatively*, while amplifying goal-oriented expressions such  
 429 as *step*, *solve*, *find*, and *need*. In the 1.5B model, *but* and *wait* drop by 0.36% and 0.20% percentage points, while  
 430 *find* and *need* rise by 0.13% and 0.14% points. This shift reflects a transition from hesitant exploration to decisive,  
 431 solution-driven reasoning. These changes reduce uncertainty and digression—traits often seen in expert trajectories

432 Table 4: Top 20 Keyword Frequency Changes Across Model Sizes  
433

434 <b>Keyword</b>	435 <b>0.5B (%)</b>			436 <b>1.5B (%)</b>			437 <b>3B (%)</b>		
	438 Static	439 Adapted	440 $\Delta$	441 Static	442 Adapted	443 $\Delta$	444 Static	445 Adapted	446 $\Delta$
but	2.73	2.59	-0.14	2.73	2.37	-0.36	2.73	2.27	-0.46
alternatively	0.86	0.79	-0.07	0.86	0.72	-0.14	0.86	0.71	-0.15
wait	2.30	2.23	-0.07	2.30	2.10	-0.20	2.30	2.00	-0.30
therefore	1.55	1.50	-0.05	1.55	1.43	-0.13	1.55	1.40	-0.15
check	0.51	0.47	-0.04	0.51	0.40	-0.11	0.51	0.33	-0.18
another	0.29	0.26	-0.03	0.29	0.20	-0.09	0.29	0.17	-0.12
then	0.97	0.94	-0.03	-	-	-	-	-	-
pi	0.11	0.09	-0.02	-	-	-	-	-	-
perhaps	0.55	0.53	-0.02	-	-	-	-	-	-
length	0.22	0.24	+0.02	-	-	-	-	-	-
step	0.27	0.30	+0.02	0.27	0.37	+0.10	0.27	0.41	+0.14
now	0.43	0.46	+0.03	0.43	0.49	+0.06	0.43	0.50	+0.07
first	0.89	0.92	+0.03	0.89	0.96	+0.07	0.89	0.96	+0.07
since	0.80	0.83	+0.03	-	-	-	-	-	-
have	0.64	0.67	+0.03	0.64	0.69	+0.05	-	-	-
let	1.83	1.86	+0.04	1.83	1.88	+0.05	-	-	-
need	0.54	0.58	+0.04	0.54	0.68	+0.14	0.54	0.69	+0.16
find	0.37	0.42	+0.04	0.37	0.50	+0.13	0.37	0.54	+0.16
newline	-	-	-	0.00	0.05	+0.05	-	-	-
equation	-	-	-	0.76	0.81	+0.05	0.76	0.85	+0.09

452 but burdensome for smaller models. In static supervision, such expressions appear frequently, straining low-capacity  
453 models and widening the *Imitation Gap* (Sec. 3.1), where expert strategies exceed model capabilities.

454 DART bridges this gap by replacing brittle reasoning paths with model-originated decision traces. This adaptation  
455 maintains task objectives while restructuring execution to fit model capacity, leading to stable and efficient reasoning.

## 458 6 RELATED WORK

459 **Chain-of-Thought Reasoning** Early work on chain-of-thought reasoning (CoT) (Wei et al., 2022) primarily focused  
460 on *short CoT*, where models generate concise reasoning paths to solve problems. Recent advances (Chen et al.,  
461 2025) have shifted towards *long CoT prompting*, encouraging more elaborate reasoning chains that enable systematic  
462 exploration of multiple paths (*branching*) and backtracking when errors are detected. While techniques like knowledge  
463 distillation (Hinton et al., 2015; Luo et al., 2025) and reinforcement learning (Hou et al., 2025) have been used to  
464 equip large language models (LLMs) with long CoT capabilities, these efforts remain largely confined to models with  
465 substantial parameter sizes. In contrast, our work specifically addresses the unique challenges associated with training  
466 smaller-scale models for complex reasoning tasks.

467 **Data-Efficient Reasoning Elicitation** A related line of work investigates how minimal supervision can elicit latent  
468 reasoning abilities in pretrained models (Ye et al., 2025; Muennighoff et al., 2025). These methods rely on a few  
469 carefully designed *cognitive templates*, to guide reasoning, but often assume that models possess the necessary prior  
470 knowledge. This assumption makes the templates brittle when cognitive demands exceed model capacity. To address  
471 this limitation, we propose a feasibility-aware adaptation framework that dynamically adjusts supervision to model  
472 ability, enabling robust reasoning across diverse capacity profiles.

## 474 7 CONCLUSION

475 We propose Dynamic Adaptation of Reasoning Trajectories (DART), a data adaptation framework designed to improve  
476 reasoning elicitation for small language models. By introducing adaptability-based selective imitation and outcome-  
477 consistent exploration, our method aims to better align expert demonstrations with model capabilities. Experimental  
478 results across several benchmarks show that DART can improve reasoning performance compared to static fine-tuning.  
479 We hope this work provides a step toward more flexible and model-aware data alignment strategies for reasoning tasks.

## 482 8 LIMITATIONS AND FUTURE WORK

483 Our framework is effective for structured reasoning tasks with verifiable outcomes. However, its extension to open-  
484 ended tasks with inherent output uncertainty remains limited, suggesting the need for refined supervision mechanisms  
485 and evaluation metrics to ensure outcome consistency.

486 REFERENCES  
487

488 489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539

Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. A neural probabilistic c model. *J. Mach. Learn. Res.*, 3(null):1137–1155, March 2003. ISSN 1532-4435.

Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models, 2025. URL <https://arxiv.org/abs/2503.09567>.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. CoRR, abs/2110.14168.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>.

Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3828–3850, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.211. URL <https://aclanthology.org/2024.acl-long.211>.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In J. Vanschoren and S. Yeung (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021b. URL [https://datasets-benchmarks-proceedings.neurips.cc/paper\\_files/paper/2021/file/be83ab3ecd0db773eb2dc1b0a17836a1-Paper-round2.pdf](https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/be83ab3ecd0db773eb2dc1b0a17836a1-Paper-round2.pdf).

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015. URL <https://arxiv.org/abs/1503.02531>.

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

Yoon Kim and Alexander M. Rush. Sequence-level knowledge distillation. In Jian Su, Kevin Duh, and Xavier Carreras (eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1317–1327, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1139. URL <https://aclanthology.org/D16-1139/>.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pp. 282–293. Springer, 2006.

Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 2024.

Wendi Li and Yixuan Li. Process reward model with q-value rankings. In *The Thirteenth International Conference on Learning Representations*.

Minpeng Liao, Chengxi Li, Wei Luo, Jing Wu, and Kai Fan. MARIO: math reasoning with code interpreter output - a reproducible pipeline. In Lun-Wei Ku, André Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, pp. 905–924. Association for Computational Linguistics, 2024.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=v8L0pN6EOi>.

540 Yijia Luo, Yulin Song, Xingyao Zhang, Jiaheng Liu, Weixun Wang, GengRu Chen, Wenbo Su, and Bo Zheng.  
 541 Deconstructing long chain-of-thought: A structured reasoning optimization framework for long cot distillation, 2025.  
 542 URL <https://arxiv.org/abs/2503.16385>.

543 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer,  
 544 Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling, 2025. URL <https://arxiv.org/abs/2501.19393>.

545 OpenAI. Openai o1 system card, 2024. URL <https://arxiv.org/abs/2412.16720>.

546

547 Dean A. Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3  
 548 (1):88–97, 1991. doi: 10.1162/neco.1991.3.1.88.

549 Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction  
 550 to no-regret online learning. In Geoffrey Gordon, David Dunson, and Miroslav Dudík (eds.), *Proceedings of  
 551 the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of  
 552 Machine Learning Research*, pp. 627–635, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR. URL <https://proceedings.mlr.press/v15/ross11a.html>.

553

554 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K.  
 555 Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models,  
 556 2024. URL <https://arxiv.org/abs/2402.03300>.

557

558 David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser,  
 559 Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks  
 560 and tree search. *nature*, 529(7587):484–489, 2016.

561

562 Maciej Świechowski, Konrad Godlewski, Bartosz Sawicki, and Jacek Mańdziuk. Monte carlo tree search: A review of  
 563 recent modifications and applications. *Artificial Intelligence Review*, 56(3):2497–2562, 2023.

564

565 Zhengyang Tang, Xingxing Zhang, Benyou Wang, and Furu Wei. Mathscales: Scaling instruction tuning for mathematical  
 566 reasoning. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024. URL <https://arxiv.org/abs/2403.02884>.

567

568 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou.  
 569 Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal,  
 570 D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp.  
 571 24824–24837. Curran Associates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/9d5609613524ecf4f15af0f7b31abca4-Paper-Conference.pdf).

572

573 Zhangchen Xu, Fengqing Jiang, Luyao Niu, Bill Yuchen Lin, and Radha Poovendran. Stronger models are not stronger  
 574 teachers for instruction tuning, 2024. URL <https://arxiv.org/abs/2411.07133>.

575

576 Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for reasoning, 2025.  
 577 URL <https://arxiv.org/abs/2502.03387>.

578

579 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-of-thought reasoning  
 580 in llms, 2025. URL <https://arxiv.org/abs/2502.03373>.

581

582 Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T. Kwok, Zhenguo Li, Adrian  
 583 Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models, 2024.  
 584 URL <https://arxiv.org/abs/2309.12284>.

585

586 Xiang Yue, Tuney Zheng, Ge Zhang, and Wenhui Chen. Mammoth2: Scaling instructions from the web, 2024. URL  
 587 <https://arxiv.org/abs/2405.03548>.

588

589 Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou,  
 590 and Junyang Lin. The lessons of developing process reward models in mathematical reasoning, 2025. URL  
 591 <https://arxiv.org/abs/2501.07301>.

592

593 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili  
 594 Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36, 2024.

594

## APPENDIX

595

596

## A IMPLEMENTATION DETAILS

597

598

In this section, we provide a detailed account of the experimental configurations and setups to ensure the transparency and reproducibility of our research. We introduce the prompt designs used in our experiments. Furthermore, we elaborate on the parameter configurations for simulation experiments and adaptive path exploration. These configurations are designed to balance computational efficiency and response diversity, ensuring the stability and adaptability of the model across various tasks.

603

604

## A.1 EXPERIMENT PROMPTS

605

606

607

In our simulation experiments, we employed a structured prompting approach to guide the language model through multi-step reasoning tasks. The primary simulation prompt used in our study is defined as follows:

608

## Simulation Prompt

609

```
Problem: data["question"]
Existing reasoning path: data["answer"]
Guidelines for continuing the reasoning:
1. Understand the existing path: Carefully analyze the existing reasoning
path and understand the logic and basis of each step.
2. Identify the next step: Based on the last step of the existing path,
determine the possible directions for the next step of reasoning.
3. Reason step-by-step: Start from the last step of the existing path and
proceed with the reasoning step-by-step.
4. Final conclusion: When the reasoning is complete, put your final
answer within boxed{}.
Continue reasoning step by step, and put your final answer within \boxed{}.
```

621

This prompt encourages the model to decompose the problem into intermediate steps and to clearly indicate the final answer using LaTeX-style boxed notation. This formatting ensures consistency across outputs and facilitates automated evaluation of results.

624

625

In addition to standard simulation prompting, we introduce a dedicated exploration prompt tailored for the adaptive trajectory rollout described in Section 3.2. This prompt is activated once a low-adaptability segment is detected and aims to continue reasoning beyond the imitation gap. It conditions the model on the prefix of high-adaptability reasoning steps and allows for autonomous continuation constrained only by outcome correctness:

629

## Exploration Prompt

630

```
Problem: {data["question"]}
Existing reasoning (read-only; cite only key points when anchoring, do NOT
restate the whole text): {data["answer"]}
[Guidelines (strictly follow)]
1. Role & Boundaries
- Continue only from the last step of the existing reasoning; do not
restate or rewrite prior content.
- If new symbols/variables are needed, first define their meaning and
domain in one sentence, then use them.
2. Anchoring & Continuation
- Use one line to anchor the key equation/state of the "last step" (key
points only; do not restate the full text).
- If you can determine the next step number from previous steps, continue
that numbering; if not, do not number--start reasoning directly.
3. Explore
- Following your own reasoning style and anchored to what has been
established, carry the reasoning forward from here.
Final conclusion: When the reasoning is complete, put your final answer
within \boxed{}.
```

648 This exploration prompt encourages the model to develop its own reasoning path from the last trustworthy segment,  
 649 fostering flexible generalization while maintaining semantic alignment with the expert outcome.  
 650

## 651 A.2 PARAMETER CONFIGURATION FOR SIMULATION

652 The simulation procedure in Algorithm 3.1 adopts stochastic decoding to explore alternative reasoning paths beyond  
 653 expert demonstrations. We sample  $N = 4$  candidate continuations per step, corresponding to the adaptation simulation  
 654 count  $N_{\text{sim}}$ .  
 655

656 Each trajectory is generated with a maximum length of `MAX_NEW_TOKENS = 4000`. To promote determinism while  
 657 retaining minimal stochasticity, we set the sampling temperature to `TEMPERATURE = 0.1`. Decoding is performed in  
 658 batches of `BATCH_SIZE = 32` to enable efficient parallel inference under hardware constraints. These settings ensure  
 659 stable simulation rollouts with low-variance outputs, suitable for evaluating adaptability under controlled decoding  
 660 conditions.  
 661

662 Regarding the computational cost of our simulation, we employ SGLang as the inference deployment framework  
 663 for small-scale models. Given that our approach primarily focuses on adapting reasoning templates to small-scale  
 664 models, it encounters challenges stemming from distributional mismatches and the limited capacity of small language  
 665 models (SLMs) compared to large language models (LLMs). Consequently, we are able to deploy our model on a  
 666 single GPU. To enhance simulation efficiency, we utilize a distributed rollout engine, with multiple SGLang workers  
 667 managed by an SGLang router to achieve load balancing. Within our code framework, for a 3B model with an estimated  
 668 four simulations per step, processing the LIMO dataset on a single node equipped with eight A100 GPUs requires  
 669 approximately six hours.  
 670

## 671 A.3 PARAMETER CONFIGURATION FOR ADAPTIVE PATH EXPLORATION

672 To support the adaptive rollout mechanism described in Section 3.2, we configured the `EXPLORE` phase with carefully  
 673 selected hyperparameters to balance computational efficiency and response diversity. The sampling procedure was  
 674 executed with a candidate beam size of `NUM_SAMPLES = 8`, meaning that at each decision step, eight reasoning  
 675 continuations were generated for evaluation based on the adaptability score.  
 676

677 We set the maximum generation length to `MAX_NEW_TOKENS = 2000` to allow sufficient space for multi-step  
 678 reasoning without premature truncation. A temperature of `TEMPERATURE = 0.7` was employed to introduce  
 679 moderate randomness in token sampling, facilitating the exploration of alternative reasoning paths while retaining  
 680 coherence.  
 681

682 Batch inference was performed with a `BATCH_SIZE = 64` to utilize GPU resources efficiently during large-scale  
 683 rollouts. The underlying language model was run using half-precision arithmetic (`DTYPE = float16`), which  
 684 reduced memory footprint and improved throughput without compromising output quality.  
 685

686 Additionally, the maximum number of concurrent sequences handled by the inference engine (VLLM) was set to  
 687 `MAX_NUM_SEQS = 512`, enabling high-throughput parallel generation during exploration. These settings ensured  
 688 scalable, stable, and semantically diverse adaptation rollouts that align with the outcome consistency constraint described  
 689 in Equation 5.  
 690

## 691 B IMPACT OF SEARCH PATH QUALITY ON MODEL PERFORMANCE

692 To investigate the impact of search quality on model performance, we conducted a comparative experiment (see Table 5).  
 693 After completing the adaptation path search, we removed paths exhibiting severe repetition phenomena. As illustrated in  
 694 Figure 3, the proportion of repeated paths during exploration decreases progressively with increasing model parameter  
 695 size, indicating that improvements in the model’s generative capability and contextual memory effectively reduce  
 696 repetition.  
 697

698 We refer to the results after removing such repeated paths as *Adaptation-Cleaned* and systematically evaluated these  
 699 against the complete search results without removing repeated paths, denoted as *Adaptation-Raw*. Experimental results  
 700 demonstrate that filtering out repeated paths leads to significant performance gains, further highlighting the critical role  
 701 of search path quality in overall model performance.  
 702

## 703 C COMPARATIVE ANALYSIS OF TRUNCATION METHODS UNDER SEARCH CONSTRAINTS

704 In our previous section (see Section 5.2), we investigate two truncation methods under different search constraints.  
 705 Specifically, we designed two variants: **Adaptation-First** and **Adaptation-Gap**. The **Adaptation-First** method  
 706

702 Table 5: Comparison of accuracy (%) on the Math-QwQ-32B dataset for 0.5B and 1.5B models under different  
703 adaptation strategies. The table contrasts the performance between *Adaptation-Raw* (without removing repeated paths)  
704 and *Adaptation-Cleaned* (with repeated paths removed). Columns for MATH and Minerva Math are excluded, and the  
705 average is computed over the remaining datasets. **Bold** values indicate the best results.

Model	Method	GSM8K	GaoKao	Olympiad Bench	College Math	MMLU STEM	Avg.
<b>Math-QwQ-32B Dataset</b>							
Qwen2.5-0.5B-Instruct	Adaptation-Raw	47.5	29.1	9.3	26.8	28.2	28.2
	Adaptation-Cleaned	<b>49.6</b>	<b>30.9</b>	<b>9.3</b>	<b>27.5</b>	<b>37.5</b>	<b>31.0</b>
Qwen2.5-1.5B-Instruct	Adaptation-Raw	70.8	44.2	18.8	38.3	44.6	43.3
	Adaptation-Cleaned	<b>74.2</b>	<b>48.6</b>	<b>19.6</b>	<b>39.4</b>	<b>57.7</b>	<b>47.9</b>

**Repetition Ratios by Model Size and Dataset**

Model Size	Limo Dataset (%)	Math-QwQ-32B Dataset (%)
0.5B	52.0%	33.6%
1.5B	24.7%	5.0%
3B	21.3%	3.3%

732 Figure 3: Repetition ratios(%) in search paths across different model sizes and datasets. Smaller models tend to have  
733 higher repetition ratios, particularly on the Limo dataset.

736 halts imitation once a feasible solution state is detected, whereas **Adaptation-Gap** monitors adaptability scores and  
737 terminates imitation when sharp declines occur, as detailed in Section 5.1.

738 We compare the truncation positions of the two methods across different datasets and model sizes. Our analysis indicates  
739 that on more challenging datasets, or when the model capacity is limited (e.g., results on the 0.5B models for both  
740 datasets), the truncation points identified by **Adaptation-First** and **Adaptation-Gap** are largely consistent. This can be  
741 attributed to the complexity of the reasoning cognitive templates in these datasets relative to the model’s capabilities:  
742 once the model identifies a path leading to a feasible solution, continued imitation often ventures into regions that are  
743 difficult to adapt to, typically accompanied by a sharp decline in adaptability scores. Consequently, the truncation  
744 positions under both **Adaptation-First** and **Adaptation-Gap** modes are generally aligned.

745 Conversely, on the Math-Qwen dataset, notable differences in truncation positions emerge. Many models, after reaching  
746 the step at which the final answer can be searched, continue to utilize subsequent adaptable path segments. Thus, the  
747 **Adaptation-Gap** method is able to detect and leverage a greater number of these usable step fragments, resulting in  
748 more substantial performance improvements, as reported in Table 3.

## D PROOF OF EXISTENCE OF IMITATION GAP

752 To rigorously establish the existence of the imitation gap in behavioral cloning for reasoning tasks, we model the  
753 process as a deterministic Markov Decision Process (MDP)  $M = (\mathcal{S}, \mathcal{A}, \mathcal{T}, r, \rho, T)$  (Li & Li), where:

754 


755 - $\mathcal{S}$ : state space of reasoning prefixes including the initial instruction  $x$ ;
- $\mathcal{A}$ : action space of reasoning steps  $s_t$ ;

756 Table 6: Comparison of truncation positions between **Adaptation-First** and **Adaptation-Gap** methods across datasets  
 757 and model sizes. The relative localization difference represents the absolute difference between the relative truncation  
 758 positions of these two methods. Higher differences are highlighted with deeper red.

760 <b>Dataset</b>	761 <b>Model Size</b>	762 <b>First Position</b>	763 <b>Gap Position</b>	764 <b>Relative Localization Difference</b>
765 Limo Dataset	0.5B	0.7901	0.7707	0.0194
	1.5B	0.6785	0.6985	0.0200
	3B	0.5718	0.5983	0.0265
766 Math-QwQ-32B Dataset	0.5B	0.5055	0.5244	0.0189
	1.5B	0.2113	0.3664	0.1551
	3B	0.0840	0.3239	0.2399

770

- 771 •  $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ : deterministic transition appending  $s_t$  to the prefix;
- 772 •  $r : \mathcal{S} \rightarrow \mathbb{R}$ : reward function, with  $r(s_T) = 1$  if the trajectory yields the correct answer  $a^*$ , and 0 otherwise;
- 773 •  $\rho$ : initial distribution over instructions ( $s_0 \sim \rho$ );
- 774 •  $T$ : maximum trajectory length (horizon).

775 Consider an expert trajectory  $\tau_{\text{expert}} = (s_0, s_1, \dots, s_T)$  generated by a strong policy  $\pi_E$ , assumed to produce near-optimal steps. The student policy  $\pi_S$ , trained via behavioral cloning (BC) on expert demonstrations, minimizes the loss  $\mathbb{E}_{\tau \sim d^{\pi_E}} \left[ \sum_{t=1}^T -\log \pi_S(s_t | s_{\leq t}) \right]$ , where  $s_{\leq t} = (s_0, \dots, s_{t-1})$  is the prefix, and  $d^{\pi_E}$  is the expert state distribution.

780 Define the Q-value under  $\pi_S$  for appending the expert action  $s_t$  at prefix  $s_{\leq t}$ :

$$f_t = Q^{\pi_S}(s_{\leq t}, s_t) = \mathbb{E}_{s_{t+1:T} \sim \pi_S(\cdot | s_{\leq t})} [\mathbb{I}(\mathcal{O}(\tau) = a^*)],$$

783 where  $\tau = (s_0, \dots, s_T)$ ,  $\mathcal{O}(\tau)$  extracts the final answer,  $a^*$  is the ground truth,  $\mathbb{I}$  is the indicator function, and  $s_{\leq t} = (s_0, \dots, s_t)$ . Following (Li & Li), we use the sigmoid-transformed Q-value for probability interpretations:

$$f_t^\sigma = \sigma(f_t) = \mathbb{P}^{\pi_S}(\mathcal{O}(\tau) = a^* | s_{\leq t}).$$

787 **Lemma 1** (Existence of Imitation Gap). *There exists a step  $t_{\text{gap}} \in [1, T]$  such that the sequence of  $f_t$  values satisfies  $f_1 < f_2 < \dots < f_{t_{\text{gap}}-1}$ , followed by a sharp drop  $f_{t_{\text{gap}}} \ll f_{t_{\text{gap}}-1}$ .*

790 *Proof.* The proof is structured in three parts, leveraging Q-value rankings from process reward models (Li & Li) and 791 the impact of distribution mismatch on the student policy.

792 **Part 1: Pre-gap monotonic increase.** For  $t < t_{\text{gap}}$ , the prefixes  $s_{\leq t}$  remain aligned with  $d^{\pi_E}$ , as the student policy  $\pi_S$  793 closely approximates  $\pi_E$ . Since the expert actions  $s_t$  are correct, we apply Lemma 3.3 from (Li & Li): for two correct 794 steps  $s_n, s_m$  in a solution  $\tau$  with  $n < m$ , we have:

$$Q^*(s_{\leq n}, s_n) < Q^*(s_{\leq m}, s_m).$$

797 The proof, adapted to our student policy:

$$\begin{aligned} f_n^\sigma - f_m^\sigma &= \mathcal{P}^{\pi_S}(s_m | s_{\leq n}) \mathcal{P}^{\pi_S}(\tau | s_{\leq m}) + \mathcal{P}^{\pi_S}(\overline{s_m} | s_{\leq n}) \mathcal{P}^{\pi_S}(\tau | \overline{s_m}) - \mathcal{P}^{\pi_S}(\tau | s_{\leq m}), \\ &= \mathcal{P}^{\pi_S}(\overline{s_m} | s_{\leq n}) [\mathcal{P}^{\pi_S}(\tau | \overline{s_m}) - \mathcal{P}^{\pi_S}(\tau | s_{\leq m})], \end{aligned}$$

802 where the first equality uses the Q-function definition, and the second uses  $\mathcal{P}^{\pi_S}(s_m | s_{\leq n}) + \mathcal{P}^{\pi_S}(\overline{s_m} | s_{\leq n}) = 1$ . Under 803 Assumption 3.1,  $\mathcal{P}^{\pi_S}(\tau | \overline{s_m}) - \mathcal{P}^{\pi_S}(\tau | s_{\leq m}) < 0$ , since the correct step  $s_m$  has a higher probability of leading to a 804 correct outcome. Thus, for  $n < m$ ,  $f_n^\sigma < f_m^\sigma$ , implying  $f_n < f_m$ . Since  $\pi_S \approx \pi_E$  for early steps, this holds for all 805  $t < t_{\text{gap}}$ , yielding:

$$f_1 < f_2 < \dots < f_{t_{\text{gap}}-1}.$$

807 **Part 2: Distribution mismatch and emergence of non-optimal step.** Due to differences in model capacity (e.g., the 808 student being a smaller model), the expert data distribution  $d^{\pi_E}$  and the student model distribution  $d^{\pi_S}$  are inconsistent. 809 As the number of steps increases, the prefixes  $s_{\leq t}$  grow increasingly complex, becoming likely to fall outside the training 810 distribution of  $\pi_S$ . Consequently, the state observed by  $\pi_S$  at step  $t$  diverges from that of  $\pi_E$ , such that the

810 expert action  $s_t$ , optimal under  $\pi_E$ , is not necessarily optimal under  $\pi_S$ . This distribution mismatch leads to a critical  
 811 step  $t_{\text{gap}}$  where the expert action  $s_{t_{\text{gap}}} = s_E$  is non-optimal for  $\pi_S$ , as it does not maximize the expected reward under  
 812 the student's policy:

$$813 \quad Q^{\pi_S}(s_{<t_{\text{gap}}}, s_E) < \max_{s \in \mathcal{A}} Q^{\pi_S}(s_{<t_{\text{gap}}}, s).$$

815 This non-optimality arises because the OOD prefix  $s_{<t_{\text{gap}}}$  causes  $\pi_S$  to misjudge the value of  $s_E$ , favoring an alternative  
 816 action that aligns better with its biased distribution, analogous to selecting an incorrect step from a correct prefix.

817 **Part 3: Sharp drop behavior.** At  $t_{\text{gap}}$ , appending the non-optimal expert action  $s_{t_{\text{gap}}}$  produces an OOD state  $s_{\leq t_{\text{gap}}}$ ,  
 818 significantly reducing the probability of correct completion. We compare the Q-value of the correct prefix at  $t_{\text{gap}} - 1$  to  
 819 the non-optimal step at  $t_{\text{gap}}$ . For the correct prefix at  $t_{\text{gap}} - 1$ , let  $s_{t_{\text{gap}}-1}$  be correct, and for the non-optimal step  $s_{t_{\text{gap}}}$ ,  
 820 we have:

$$821 \quad f_{t_{\text{gap}}-1}^{\sigma} - \mathcal{V}^{\pi_S}(x) = \mathcal{P}^{\pi_S}(\overline{s_{t_{\text{gap}}-1}}|x) (\mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}-1}) - \mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}-1}})),$$

$$822 \quad f_{t_{\text{gap}}}^{\sigma} - \mathcal{V}^{\pi_S}(x) = \mathcal{P}^{\pi_S}(s_{t_{\text{gap}}}|x) (\mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}}}) - \mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}})) ,$$

823 where  $\mathcal{V}^{\pi_S}(x) = \mathbb{P}^{\pi_S}(\mathcal{O}(\tau) = a^*|x)$ . Under Assumption 3.1,  $\mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}-1}) > \mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}-1}})$ , so the first  
 824 difference is positive, implying  $f_{t_{\text{gap}}-1}^{\sigma} > \mathcal{V}^{\pi_S}(x)$ . For the non-optimal step,  $\mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}}}) < \mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}})$ , and since  
 825  $s_{t_{\text{gap}}}$  is non-optimal due to distribution mismatch,  $\mathcal{P}^{\pi_S}(s_{t_{\text{gap}}}|x) \gg \mathcal{P}^{\pi_S}(\overline{s_{t_{\text{gap}}}}|x)$ , amplifying the negative difference.  
 826 Thus:

$$827 \quad f_{t_{\text{gap}}}^{\sigma} < \mathcal{V}^{\pi_S}(x) < f_{t_{\text{gap}}-1}^{\sigma},$$

828 implying  $f_{t_{\text{gap}}} \ll f_{t_{\text{gap}}-1}$ , as the non-optimal step's Q-value is significantly lower due to the low probability of recovery  
 829 from incorrect branches.

830 The key size relation for the drop is:

$$833 \quad f_{t_{\text{gap}}}^{\sigma} - f_{t_{\text{gap}}-1}^{\sigma} = \mathcal{P}^{\pi_S}(s_{t_{\text{gap}}}|s_{<t_{\text{gap}}}) [\mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}}) - \mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}}})] + \mathcal{P}^{\pi_S}(\overline{s_{t_{\text{gap}}}}|s_{<t_{\text{gap}}}) [\mathcal{P}^{\pi_S}(\tau|\overline{s_{\leq t_{\text{gap}}}}) - \mathcal{P}^{\pi_S}(\tau|s_{\leq t_{\text{gap}}})] < 0,$$

834 where the negative term dominates under Assumption 3.1, ensuring  $f_{t_{\text{gap}}} \ll f_{t_{\text{gap}}-1}$ .  $\square$

## 836 E DESCRIPTION OF LARGE LANGUAGE MODEL USAGE

837 In the preparation of this manuscript, we leveraged a large language model (LLM), specifically Grok developed by  
 838 xAI, to facilitate specific aspects of the writing process. The LLM was employed primarily for linguistic refinement,  
 839 encompassing tasks such as enhancing sentence coherence, improving syntactic clarity, and elevating the overall  
 840 readability of the text, while preserving the integrity of the scientific content, methodologies, and findings. The rationale  
 841 for this approach was to optimize the communicative efficacy of the manuscript, ensuring that intricate technical  
 842 concepts are articulated with precision and accessibility for a diverse academic readership. All outputs generated by  
 843 the LLM were subjected to rigorous scrutiny, validation, and, where necessary, revision by the authors to uphold the  
 844 principles of accuracy, originality, and academic rigor. Notably, the LLM was not utilized for the generation of novel  
 845 intellectual contributions, experimental frameworks, data analyses, or conclusions, which were exclusively derived from  
 846 human expertise. This judicious application of LLMs adheres to established ethical standards for AI-assisted academic  
 847 writing, balancing the enhancement of textual quality with a commitment to transparency and scholarly integrity.