

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 RL OF THOUGHTS: NAVIGATING LLM REASONING WITH INFERENCE-TIME REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

## ABSTRACT

Despite rapid advancements in large language models (LLMs), the token-level autoregressive nature constrains their complex reasoning capabilities. To enhance LLM reasoning, inference-time techniques, including Chain/Tree/Graph-of-Thought(s), successfully improve the performance, as they are fairly cost-effective by guiding reasoning through external logical structures without modifying LLMs' parameters. However, these manually predefined, task-agnostic frameworks are applied uniformly across diverse tasks, lacking adaptability. To improve this, we propose **RL-of-Thoughts (RLoT)**, where we train a lightweight navigator model with reinforcement learning (RL) to generate task-adaptive logical structures at inference time, enhancing LLM reasoning. Specifically, we design five basic logic blocks from the perspective of human cognition. During the reasoning process, the trained RL navigator dynamically selects the suitable logic blocks and combines them into task-specific logical structures according to problem characteristics. Experiments across multiple reasoning benchmarks (AIME, MATH, GPQA, etc.) with multiple LLMs (GPT, Llama, Qwen, and DeepSeek) illustrate that RLoT outperforms established inference-time techniques in most cases and improves up to 13.4% in challenging situations. Remarkably, with less than 3K parameters, our RL navigator is able to make sub-10B LLMs comparable to 100B-scale counterparts. Moreover, the RL navigator demonstrates strong transferability: a model trained on one specific LLM-task pair can effectively generalize to unseen LLMs and tasks. Our code is open-source at <https://anonymous.4open.science/r/RL-LLM-Reasoning-1A30>.

## 1 INTRODUCTION

Recent years have witnessed unprecedented advancements in large language models (LLMs), achieving remarkable success across diverse natural language tasks (Chang et al., 2024), including translation (Xu et al., 2024), semantic analysis (Lan et al., 2024b;a), and information retrieval (Hao et al., 2024). Despite these advancements, the inherent token-level autoregressive nature of LLMs poses a significant limitation for complex reasoning tasks (Zhao et al., 2023), such as solving mathematical problems (Ahn et al., 2024) or answering intricate questions (Zhuang et al., 2023). These tasks require sophisticated logical structures and long-term dependencies that go beyond the scope of simple sequential token prediction, leaving a considerable gap between current LLM capabilities and the demands of advanced reasoning applications.

Plentiful research has been devoted to enhancing LLM reasoning. On one hand, fine-tuning approaches attain substantial improvements on pretrained LLMs (Zhong et al., 2024; DeepSeek-AI et al., 2025; Team et al., 2025). However, these methods demand massive computational resources and large-scale datasets, being costly to implement. On the other hand, inference-time techniques, exemplified by Chain-of-Thought (Wei et al., 2022), Tree-of-Thoughts (Yao et al., 2023), and Graph-of-Thoughts (Besta et al., 2024), offer a lightweight alternative by enhancing reasoning through predefined external logical structures. While cost-effective, their logical structures rely on manual design and are task-agnostic, lacking the adaptability to diverse reasoning tasks.

Addressing such limitations in inference-time techniques presents significant challenges. First, reasoning tasks span various domains, including mathematics, STEM, commonsense, etc., where tasks in each domain exhibit diverse characteristics, making it infeasible to manually design logical

structures specified for each task. Second, complex reasoning tasks often require multiple steps, where the problem-solving status evolves after each step, requiring dynamic adjustments to the logical structure for subsequent reasoning. Therefore, predefined logical structures fail to adapt to the changes, limiting their effectiveness in stepwise reasoning tasks. These challenges highlight the need for more adaptive inference-time techniques to handle reasoning tasks with diversity and dynamics.

Facing these challenges, we introduce **RL-of-Thoughts (RLoT)**, a framework that leverages reinforcement learning (RL) at inference time to enhance the reasoning capabilities of LLMs. Specifically, we model long-sequence reasoning as a Markov Decision Process (MDP) and design five human cognition-inspired basic logical blocks as potential actions for decision-making. Within the MDP framework, we train an RL agent, namely the navigator model, to dynamically select and combine these blocks along the reasoning process, constructing task-specific logical structures and thereby enhancing the LLM’s ability to handle complex reasoning tasks. We conduct experiments across a wide range of reasoning benchmarks, including AIME (Olympic mathematics), MATH (elementary mathematics), GPQA (STEM), StrategyQA (commonsense), etc. The results demonstrate that our RLoT design outperforms various established inference-time techniques in most cases while being compatible with multiple well-known LLMs, such as GPT, Llama, Qwen, and DeepSeek. Remarkably, our RL navigator, which contains less than 3K parameters, is able to enhance the performance of sub-10B LLMs, making them comparable to much larger LLMs with  $10\times$  parameters. Moreover, the RL navigator exhibits strong transferability: a model trained with one specific LLM on one task domain can effectively generalize to unseen LLMs and tasks without fine-tuning.

In summary, the main contributions of this work include:

- We propose RLoT, an inference-time technique that leverages RL to adaptively construct task-specific logical structures, enhancing LLM reasoning.
- We conduct extensive experiments to verify the effectiveness of our method to improve LLM reasoning across various tasks. Compatible with multiple widely known LLMs, RLoT outperforms established inference-time techniques by up to 13.4%.
- We demonstrate the transferability and efficiency of our method, where the trained navigator model can transfer across various LLMs and reasoning tasks. With < 3K parameters, it enhances multiple sub-10B LLMs to be comparable to  $10\times$  larger counterparts.

## 2 PRELIMINARIES

### 2.1 LARGE LANGUAGE MODELS (LLMs)

Large language models (LLMs) are a class of advanced neural networks characterized by parameter scales up to billions, primarily trained through next-token prediction objectives. Given a sequence  $\{w_1, w_2, \dots, w_{t-1}\}$ , the models output  $w_t$  to maximize the observation likelihood in the corpus as:

$$\prod_{t=1}^T P(w_t|w_1, w_2, \dots, w_{t-1}). \quad (1)$$

Recent advancements in LLMs, exemplified by architectures like the GPT series (Brown et al., 2020; Kalyan, 2023; Achiam et al., 2023), the Llama family (Touvron et al., 2023; Dubey et al., 2024), etc, have demonstrated remarkable proficiency across diverse natural language understanding and generation tasks, including semantic parsing, cross-lingual translation (Zhao et al., 2023; Chang et al., 2024). Meanwhile, extensive researches integrate inference-time techniques like Chain-of-Thought (CoT) (Wei et al., 2022) and Tree-of-Thoughts (ToT) (Yao et al., 2023) to enhance the multi-step reasoning capability of LLMs. On the other hand, fine-tuning strategies leverage Outcome Reward Models (ORM) and Process Reward Models (PRM) to optimize the reasoning process through reward-guided learning (Lightman et al., 2023; Wang et al., 2024c; Luo et al., 2024). These approaches address both structural limitations of auto-regressive decoding and the challenge of maintaining logical coherence in complex tasks (Xu et al., 2025).

## 108 2.2 MARKOV DECISION PROCESS (MDP)

109

110 A Markov Decision Process (MDP) provides the core framework for sequential decision-making  
 111 problems. An MDP is mathematically defined by the tuple  $(\mathcal{S}, \rho, \mathcal{A}, P, R)$ , where  $\mathcal{S}$  is the state space,  
 112 and  $\rho \in \Delta(\mathcal{S})$  represents the probability distribution over initial states, with  $\Delta(\mathcal{S})$  being the set of all  
 113 probability distributions over  $\mathcal{S}$ . The action space is denoted by  $\mathcal{A}$ . Given a specific action taken in a  
 114 particular state, the state transition probability function  $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$  and the reward function  
 115  $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  define the likelihood of transitioning between states and the reward associated with  
 116 each action. At each time step  $t$ , the agent chooses an action  $a_t \in \mathcal{A}$  in state  $s_t \in \mathcal{S}$ , receives a  
 117 reward  $r_t$ , and transitions to the next state  $s_{t+1}$ . The agent's objective in an MDP is to maximize the  
 118 total accumulated reward over time, which is the sum of the discounted rewards obtained at each step.  
 119 The cumulative reward at time step  $t$  is expressed as  $G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$ , where  $\gamma$  is the discount  
 120 factor that weighs the significance of future rewards.

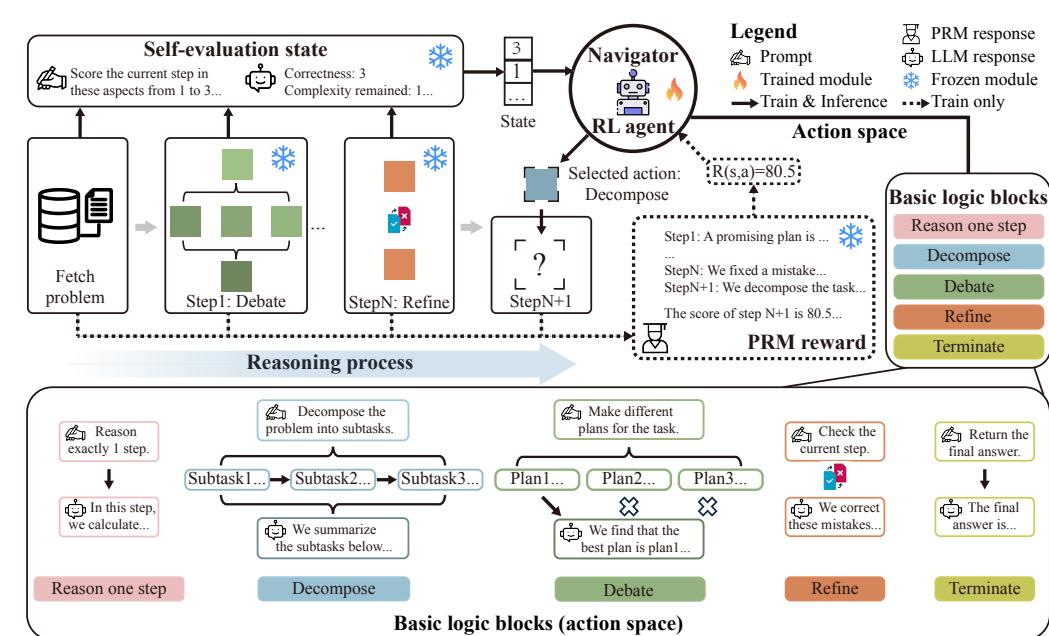
120

## 121 3 METHODS

122

## 123 3.1 OVERVIEW

124



145

146 Figure 1: Framework of RL-of-Thoughts (RLoT). We train an RL agent as the navigator, which  
 147 dynamically selects and combines basic logic blocks along the reasoning process, constructing  
 148 task-specific logical structures for each task and thereby enhancing the LLMs' ability to handle  
 149 complex reasoning tasks.

150

151 In this paper, we propose to enhance LLM reasoning at inference time with RL, namely **RL-of-**  
 152 **Thoughts (RLoT)**. The overall framework of RLoT is illustrated in Figure 1. Specifically, we design the action  
 153 space, state space, and state transition mechanism, making the sequence of decisions within an  
 154 episode correspond to the generation of a logical structure for reasoning (Section 3.2). Within this  
 155 MDP framework, we train an RL agent, referred to as the navigator model, which conducts sequential  
 156 decision-making to construct task-specific logical structures, thereby enhancing LLMs' capability to  
 157 address complex reasoning tasks (Section 3.3).

158

## 159 3.2 LLM REASONING AS MDP

160

161 To leverage the sequential decision-making capability of RL for adaptively designing logical structures  
 at inference time based on problem characteristics, we first model long-sequence reasoning as a

162 specially designed MDP. Within this MDP, the sequence of decisions in an episode corresponds to  
 163 the generation of a logical structure for reasoning. The designs of the state, action, reward, and state  
 164 transition are as follows.

165 **State.** The state space is designed to capture how the current solving status of the task is after steps of  
 166 reasoning, thereby supporting dynamic adjustments to the logical structure for subsequent reasoning.  
 167 We employ a self-evaluation mechanism to extract concise and informative states during the reasoning  
 168 process. Specifically, we prompt the LLM itself to evaluate the current reasoning steps from three  
 169 major and seven detailed aspects, which are listed in Table 1.

170  
 171 **Table 1: Aspects for state self-evaluation.**

172 Major	173 A: Correctness	174 B: Complexity	175 C: Completeness
176 Detailed	A1: Correctness of modeling A2: Clarity for further reasoning A3: Correctness of calculation	B1: Complexity to the final answer B2: Alternative methods in further reasoning	C1: Closeness to the final solution C2: Completeness within the step

177 For each detailed aspect, we require the LLM to assign a score from 1 to 3, which is then aggregated to  
 178 form the state of the MDP. This approach summarizes complex reasoning steps into low-dimensional  
 179 states, offering a comprehensive overview of the changing problem-solving status during the reasoning  
 180 process, facilitating the RL agent in adjusting the strategy for subsequent reasoning accordingly.  
 181 Please refer to detailed self-evaluation prompts in Appendix K.1.

182 **Action.** When addressing difficult and complex problems, humans often employ specific cognitive  
 183 strategies. For example, we break down complex tasks into smaller components and review previous  
 184 steps when encountering anomalies in the solution. As evidenced by previous research, understanding  
 185 and applying these cognitive strategies can significantly enhance the reasoning capabilities of  
 186 LLMs (Wu et al., 2024; Xue et al., 2024).

187 With inspiration from human cognition, we design five “basic logic blocks” that can be flexibly  
 188 combined, constituting the action space in our MDP. By selecting and cascading the blocks, the agent  
 189 thereby constructs flexible logical structures, paving reasoning pathways from the initial problem to  
 190 final solutions. In detail, the basic logic blocks include:

- 192 • **Reason one step:** Perform reasoning for a single step of the current task, which may not  
 193 directly lead to the final answer but contributes to the overall process.
- 194 • **Decompose:** Break the current task into simpler subtasks and execute them sequentially.  
 195 Then, we prompt the LLM to briefly summarize the results of these subtasks as the final  
 196 result of this action.
- 197 • **Debate:** Generate multiple plans or approaches for the task at hand and compare them to  
 198 identify the most promising one. Then, we prompt the LLM to reason one step further based  
 199 on the selected plan.
- 200 • **Refine:** Review and revise the current reasoning step to improve clarity and correctness.
- 201 • **Terminate:** Based on all the previous steps, provide the final answer to the original problem  
 202 and show it in a specified format. This action marks the conclusion of the reasoning process.

203 We illustrate the structures and detailed prompts for each blocks in Figure 1 and Appendix K.2.

204 **Reward.** To evaluate the quality of the intermediate results after the agent selects an action during  
 205 the long-sequence reasoning process, we employ the Process Reward Model (PRM) to score the  
 206 intermediate results and set the PRM score of the intermediate result after each action as the single-  
 207 step reward for this action.

208 **State Transition.** In our MDP design, the state transition is straightforward. During the reasoning  
 209 process, executing a specific action based on the current problem-solving state is to prompt the LLM  
 210 to continue reasoning using the logical structure corresponding to that action. After reasoning, the  
 211 new problem-solving state is obtained through the aforementioned self-evaluation approach.

212 Also, we impose a few simple restrictions on the state transition, ensuring the correctness and  
 213 rationality of the constructed logical structures. First, once the answer is already presented in the

216 response after executing some action, no further actions are permitted except for “Terminate”. Second,  
 217 the “Refine” action is automatically converted to “Reason one step” when it appears as the first action,  
 218 as no refinement to the original problem is needed. Finally, to avoid the reasoning process being too  
 219 long, we limit the maximum number of actions, and after reaching the limitation, the “Terminate”  
 220 action will be automatically executed.  
 221

### 222 3.3 TRAINING OF THE NAVIGATOR MODEL

224 Within the MDP framework outlined above, given a specific LLM type and a kind of reasoning task,  
 225 we train the navigator model. Under this formulation, the training process constitutes a standard RL  
 226 problem within a discrete action space. Consequently, our framework is algorithm-agnostic, allowing  
 227 for the employment of arbitrary off-the-shelf RL algorithms for training. To enhance the learning  
 228 for challenging reasoning tasks, we extract hard questions from the training set of the target task,  
 229 i.e., questions that the LLM cannot answer when directly prompted. Then, we use these problems  
 230 for training the navigator model, from which we randomly select one in each episode and repeat it  
 231 multiple times.

232 We illustrate the training and inference pipeline of RLoT in Figure 1. During training, we hire  
 233 PRM to provide reward signals. The parameters of both the PRM and the LLM are kept fixed from  
 234 pre-trained models, and only parameters of the navigator model, i.e., the RL agent, are updated.  
 235 This significantly reduces the computational cost, making the training process highly efficient. After  
 236 training, the PRM model is no longer required, and the trained navigator model is used directly.  
 237 Given an intermediate reasoning state, the navigator selects an action, which is then used to prompt  
 238 the LLM to continue reasoning using the logical structure associated with the selected action. By  
 239 repeating this, the navigator model is able to guide the LLM in solving challenging reasoning tasks  
 240 with task-specific logical structures.

## 241 4 EXPERIMENTS

### 242 4.1 EXPERIMENTAL SETTINGS

243 **Reasoning Tasks.** We conduct a comprehensive evaluation of our RLoT method across a wide range  
 244 of reasoning tasks, encompassing benchmarks in mathematics, STEM, and commonsense question  
 245 answering. For the mathematics domain, we adopt Olympic-level datasets, the AIME24<sup>1</sup> and  
 246 AMC23<sup>2</sup>, as well as elementary math datasets GSM8K (Cobbe et al., 2021) and MATH (Hendrycks  
 247 et al., 2021c). These are widely recognized as representative benchmarks. In the STEM tasks, we  
 248 test on the MMLU-STEM (Hendrycks et al., 2021b;a) and GPQA (Rein et al., 2023) datasets, which  
 249 span various STEM domains and a range of difficulty levels. To evaluate the commonsense reasoning  
 250 ability, we employ the StrategyQA (Geva et al., 2021) benchmark, which presents challenging multi-  
 251 hop questions across diverse contexts. These benchmarks cover various domains, difficulties, and  
 252 task types, forming a systematic evaluation of the reasoning ability.  
 253

254 **LLMs.** Our RLoT framework is designed to be independent of specific LLMs, allowing it to be  
 255 compatible with any off-the-shelf LLM. To evaluate this, we test our approach using four repre-  
 256 sentative LLMs: Qwen2.5-7B-Instruct, Qwen2.5-14B-Instruct (Yang et al., 2024a), Llama3.1-8B-  
 257 Instruct (Dubey et al., 2024), GPT-4o-mini (Hurst et al., 2024), and DeepSeek-R1-Distill-Qwen-  
 258 7B (DeepSeek-AI et al., 2025). In this paper, we mainly focus on sub-10B LLMs, which are often  
 259 constrained in handling complex reasoning tasks due to their relatively smaller size. We expect that  
 260 our RLoT design can substantially enhance smaller LLMs by adaptively generating task-specific  
 261 logical structures at inference time, thereby making their reasoning capabilities comparable to, or  
 262 even exceeding, those of much larger LLMs. Without the need to modify the LLM’s parameters, this  
 263 approach will be fairly computationally efficient. Note that we abandon the results of Llama3.1-8B-  
 264 Instruct on the Olympic-level datasets since the capability of this base LLM is too limited to solve  
 265 these challenging problems.

266 **Baselines.** We compare RLoT against various baselines designed to enhance LLM reasoning at  
 267 inference time. First, we evaluate single-round question-answering techniques, including direct

268<sup>1</sup><https://huggingface.co/datasets/AI-MO/aimo-validation-aime>

269<sup>2</sup><https://huggingface.co/datasets/AI-MO/aimo-validation-amc>

question answering (Direct QA), zero-shot Chain-of-Thought (Zero-shot CoT), and few-shot Chain-of-Thought (Few-shot CoT) (Wei et al., 2022). For Zero-shot CoT, we employ prompts with “Let’s think step by step”, and for Few-shot CoT, we include specific few-shot examples for each benchmark in the prompts as outlined in prior work (Yang et al., 2024b; Fu et al., 2023; Rein et al., 2023; Wei et al., 2022). Additionally, we consider multi-round techniques, including self-consistent Chain-of-Thought (CoT-SC) (Wang et al., 2023) and Tree-of-Thoughts (ToT) (Yao et al., 2023). Following the original settings, we perform majority voting across four reasoning samples in CoT-SC, and we implement a logical tree with two layers and five nodes per layer in ToT. These multi-round approaches facilitate comparison and voting across diverse reasoning paths, thereby enhancing the reasoning capability of LLMs in complex tasks.

## 4.2 TRAINING DETAILS

In our implementation, the navigator model is a simple three-layer multilayer perceptron (MLP) with the Dueling Network architecture (Wang et al., 2016). The model merely contains 2,566 parameters in total, where the lightweight design ensures efficient training and inference. We employ the Double-Dueling-DQN algorithm (Mnih et al., 2015; Van Hasselt et al., 2016; Wang et al., 2016) for optimization. By integrating Double Q-learning to mitigate value overestimation and the Dueling architecture to separate state and advantage representations, this algorithm significantly improves stability during the training process. We train the navigator model for 3,000 episodes with Qwen2.5-14B-Instruct on the MATH benchmark, where the learning curves presented in Figure 2 and Appendix A.1 indicate strong convergence. Furthermore, in Appendix A.2, we investigate the impact of employing alternative RL algorithms and extending training episodes, where the results consistently demonstrate similar convergence.

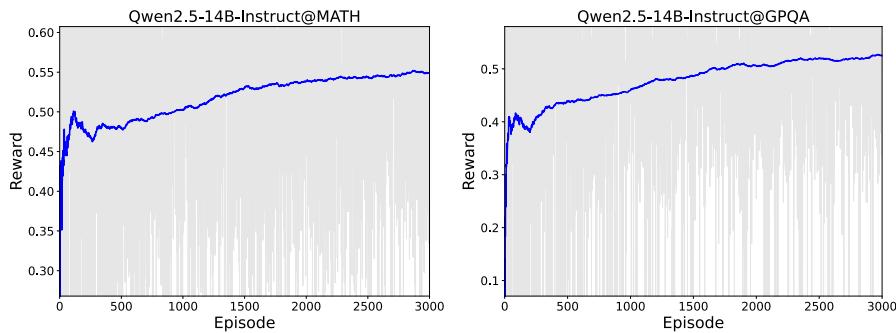


Figure 2: Learning curves during RL training of navigator models.

To obtain reward signals for actions, we utilize the Math-Shepherd as the PRM (Wang et al., 2024c). This model is trained by previous researchers using automatically constructed process-wise supervision data, reducing heavy reliance on manual annotation and thereby achieving remarkable performance. For reproducibility, please refer to more detailed settings and the illustration of the whole pipeline in Appendix J.

## 4.3 OVERALL PERFORMANCE

We use the obtained navigator model to enhance the reasoning of multiple LLMs across different reasoning tasks. Following existing works like CoT-SC, we perform multiple repeated inferences for each task and filter out the trails that do not meet self-consistency, enhancing the robustness of reasoning, where the results are presented in Table 2. For the baseline performance, we prioritize using the results reported in the official technical reports of each LLM (Yang et al., 2024a; Dubey et al., 2024; Hurst et al., 2024), if available. Otherwise, we evaluate the performance through our own experiments following the settings established in previous works (Fu et al., 2023; Yang et al., 2024b; Yao et al., 2023).

The results show that our method performs well, consistently outperforming the established inference-time baselines across almost all reasoning tasks when combined with various LLMs. Notably, our

324  
 325 Table 2: Overall evaluation of RLoT’s capability to enhance multiple LLMs’ reasoning across  
 326 different tasks. The **bold** numbers indicate the best performance in each group of experiments, and  
 327 the underlined numbers indicate the best baseline method.

328 <b>LLM</b>	<b>Method</b>	<b>Olympic math</b>		<b>Elementary math</b>		<b>STEM</b>		<b>Commonsense</b>	<b>Average</b>
		AIME24	AMC23	MATH	GSM8K	GPQA	MMLU-STEM	StrategyQA	
330      Qwen2.5-14B-Instruct	Direct QA	13.33	<u>57.50</u>	78.62	93.93	36.60	85.38	72.34	62.53
	Zero-shot CoT	<u>16.67</u>	<u>62.50</u>	78.56	94.23	38.39	85.63	75.98	<u>64.57</u>
	Few-shot CoT	<u>6.67</u>	<u>55.00</u>	80.00	<b>94.80</b>	<u>45.50</u>	85.06	78.60	63.66
	CoT-SC	<u>6.67</u>	47.50	80.04	<u>94.08</u>	<u>45.54</u>	<u>86.71</u>	80.06	62.94
	ToT	10.00	55.00	79.50	93.78	45.08	<u>86.55</u>	78.17	64.01
	RLoT (ours)	<b>23.33</b>	<b>65.00</b>	<b>80.38</b>	94.16	<b>51.34</b>	<b>88.93</b>	<b>81.22</b>	<b>69.19</b>
334      Qwen2.5-7B-Instruct	Direct QA	10.00	42.50	74.64	91.58	31.25	80.94	68.85	57.11
	Zero-shot CoT	<u>16.67</u>	<u>55.00</u>	74.86	91.58	<u>34.15</u>	81.00	71.17	60.63
	Few-shot CoT	<u>13.33</u>	<u>45.00</u>	75.50	91.60	36.40	81.16	74.38	59.62
	CoT-SC	<u>6.67</u>	50.00	76.36	92.12	38.84	<u>83.25</u>	<u>78.45</u>	60.81
	ToT	13.33	50.00	73.80	91.35	36.60	82.62	<u>70.45</u>	59.74
	RLoT (ours)	<b>23.33</b>	<b>60.00</b>	<b>76.70</b>	<b>92.87</b>	<b>44.64</b>	<b>85.06</b>	<b>79.04</b>	<b>65.95</b>
338      Llama3.1-8B-Instruct	Direct QA			49.12	84.83	32.14	71.58	70.74	61.68
	Zero-shot CoT			49.74	85.37	32.80	72.85	73.07	62.77
	Few-shot CoT			48.52	84.50	33.03	70.66	72.05	61.75
	CoT-SC			51.74	87.04	33.48	73.29	78.89	64.89
	ToT			51.93	86.80	<u>32.24</u>	74.72	<u>71.47</u>	63.43
	RLoT (ours)			<b>56.56</b>	<b>90.07</b>	<b>46.88</b>	<b>80.56</b>	<b>84.42</b>	<b>71.70</b>
342      GPT-4o-mini	Direct QA	13.33	55.00	76.58	93.33	43.08	85.70	77.00	61.17
	Zero-shot CoT	<u>6.67</u>	<u>67.50</u>	<u>76.76</u>	<b>93.93</b>	40.20	85.76	78.17	<u>64.14</u>
	Few-shot CoT	<u>6.67</u>	<u>57.50</u>	75.46	<u>93.48</u>	35.94	85.82	80.06	<u>62.13</u>
	CoT-SC	<u>6.67</u>	45.00	76.84	93.63	<u>46.42</u>	86.55	82.53	62.52
	ToT	<u>6.67</u>	50.00	<u>74.30</u>	93.33	<u>44.42</u>	<u>85.92</u>	<u>76.42</u>	61.58
	RLoT (ours)	<b>20.00</b>	<b>70.00</b>	<b>77.36</b>	93.86	<b>54.02</b>	<b>88.23</b>	<b>82.68</b>	<b>69.45</b>
346      DeepSeek-R1- 347      Distill-Qwen-7B	Direct QA	46.67	60.00	92.27	95.74	54.47	83.61	79.48	73.18
	Zero-shot CoT	53.33	62.50	91.48	95.20	53.13	85.28	78.89	74.26
	Few-shot CoT	56.67	<u>67.50</u>	92.54	94.38	56.47	87.47	79.33	76.34
	CoT-SC	<u>56.67</u>	<u>67.50</u>	<u>95.54</u>	96.13	60.94	89.03	<u>82.82</u>	78.38
	ToT	50.00	55.00	<u>95.18</u>	<u>94.54</u>	55.13	<u>86.55</u>	<u>80.91</u>	<u>73.90</u>
	RLoT (ours)	<b>63.33</b>	<b>77.50</b>	<b>96.56</b>	<b>98.94</b>	<b>67.19</b>	<b>90.77</b>	<b>86.17</b>	<b>82.92</b>

350  
 351  
 352 approach brings about substantial improvements in the GPQA benchmark, which is a challenging task  
 353 where LLMs generally perform poorly. Specifically, when implemented with Llama3.1-8B-Instruct,  
 354 we achieve a 13.4% performance boost. Among the baselines, CoT-SC performs the best across most  
 355 tasks. Meanwhile, we find that despite being more complex in design, ToT performs poorly on many  
 356 reasoning tasks, which is also reported in previous studies (Wu et al., 2024; Zhang et al., 2024a;  
 357 Qi et al., 2024). Furthermore, we report the computational overhead and implementation latency  
 358 of RLoT in Appendix C, and also compare RLoT with more recently proposed test-time scaling  
 359 baselines in Appendix E. By directly generate specific logical structures for each question without  
 360 requiring search-and-trial, our method reached the best performance while maintain a low cost.

#### 363      4.4 PARAMETER SIZE EFFICIENCY

364  
 365 In this section, we demonstrate the parameter size efficiency of our RLoT method in enhancing the  
 366 reasoning capability of sub-10B LLMs, making them comparable to LLMs with several times more  
 367 parameters. Specifically, we select three 10B LLMs, including Qwen2.5-14B-Instruct, Llama3.1-8B-  
 368 Instruct, and GPT-4o-mini, and their respective large-scale counterparts. It is worth noting that the  
 369 parameter size of GPT-4o series models was estimated in previous studies (Abacha et al., 2024).

370      In Appendix D, we present the performance of these models across various reasoning tasks using  
 371      Few-shot CoT, which is the standard technique commonly used in official technical reports (Yang  
 372      et al., 2024a; Dubey et al., 2024; Hurst et al., 2024) and previous studies (Liu et al., 2024; ret, 2024;  
 373      Tran et al., 2024; Kumar et al., 2024; Yu et al., 2024) in evaluation of LLMs. We also show the  
 374      performance of the 10B LLMs after enhancement with RLoT. The results indicate that our RL-based  
 375      navigator, which contains fewer than 3,000 parameters, significantly enhances the performance  
 376      of sub-10B LLMs, making them comparable to much larger counterparts with around 10× more  
 377      parameters. Specifically, RLoT empowers the sub-10B LLMs to be comparable to, compensating  
 378      most of the performance gap, or even surpassing their larger counterparts, demonstrating its efficiency.

378 4.5 TRANSFERABILITY  
379380 To better illustrate the transferability of our navigator model, we conduct further experiments regarding  
381 transferring across different LLMs and reasoning tasks, respectively.  
382383 **Transfer across Different LLMs.** To verify the transferability of RLoT across different LLMs,  
384 we respectively train navigator models with three different LLMs, namely Qwen2.5-14B-Instruct,  
385 Llama3.1-8B-Instruct, and GPT-4o-mini, on the MATH benchmark. Without any fine-tuning, we  
386 cross-test the obtained navigator models to enhance other LLMs on the MATH benchmark and  
387 present the results in Table 3.  
388389 The results indicate that the trained navigator model exhibits strong transferability across different  
390 LLMs. When implementing the navigator model to enhance the reasoning capabilities of a specific  
391 LLM, we find that, regardless of whether the navigator model is trained on the same LLM or a  
392 different one, its performance remains consistent. In all cases, the enhanced LLM outperforms the  
393 best well-known inference-time baseline.  
394395 Table 3: Evaluation of RLoT’s transferability across different LLMs. We train navigator models with  
396 three different LLMs on the MATH benchmark and cross-test the obtained navigator models with  
397 other LLMs. We also list CoT-SC, the best baseline method, for comparison.

Method	Train	Test		
		Qwen2.5-14B-Instruct @MATH	Llama3.1-8B-Instruct @MATH	GPT-4o-mini @MATH
RLoT (ours)	Qwen2.5-14B-Instruct@MATH	80.38	56.56	77.36
	Llama3.1-8B-Instruct@MATH	81.48	53.60	78.14
	GPT-4o-mini@MATH	80.84	56.94	78.08
CoT-SC	-	80.04	51.74	76.84

405 **Transfer across Different Reasoning Tasks.** To verify the transferability of RLoT across different  
406 reasoning tasks, we respectively train navigator models with Qwen2.5-14B-Instruct on three different  
407 benchmarks, namely MATH, GPQA, and StrategyQA. Without any fine-tuning, We cross-test the  
408 obtained navigator models to enhance the reasoning capabilities of Qwen2.5-14B-Instruct on the  
409 other tasks and present the results in Table 4.  
410411 The results indicate that the trained navigator model owns strong transferability across different  
412 reasoning tasks. When utilizing the navigator model to enhance the reasoning capabilities of LLMs  
413 for a specific task, we observe that, regardless of whether the navigator model is trained on the same  
414 task or a different one, its performance remains largely consistent. In most cases, it surpasses the  
415 best-performing inference-time baseline.  
416417 Furthermore, we find that the navigator models trained on mathematical (MATH) and STEM (GPQA)  
418 problems exhibit better transferability to each other. However, the transferability between the  
419 navigator trained on commonsense problems (StrategyQA) and the former two is relatively limited,  
420 which is intuitive given the inherent relations and differences between domains of mathematics,  
421 STEM, and commonsense reasoning.  
422423 Table 4: Evaluation of RLoT’s transferability across different tasks. We train navigator models with  
424 Qwen2.5-14B-Instruct on three different tasks and cross-test the obtained navigator models on other  
425 tasks. We also list CoT-SC, the best baseline method, for comparison.

Method	Train	Test		
		Qwen2.5-14B-Instruct @MATH	Qwen2.5-14B-Instruct @GPQA	Qwen2.5-14B-Instruct @StrategyQA
RLoT (ours)	Qwen2.5-14B-Instruct@MATH	80.38	51.34	81.22
	Qwen2.5-14B-Instruct@GPQA	80.76	53.57	80.64
	Qwen2.5-14B-Instruct@StrategyQA	79.94	50.22	81.37
CoT-SC	-	80.04	45.54	80.06

432  
 433 Table 5: Typical patterns in the task-specific logical structures generated by RLoT (“Reason” is short  
 434 for “Reason one step”).

Task	MATH	GPQA	StrategyQA
<b>Instance number</b>	5000	448	687
<b>Major 2-step pattern</b>	Reason-Refine Reason-Decompose	Reason-Refine Reason-Decide	Reason-Debate Reason-Refine
<b>Major 3-step pattern</b>	Decompose-Refine-Reason Reason-Refine-Decide	Reason-Refine-Decide Reason-Decompose-Refine	Reason-Decompose-Decide Reason-Refine-Decide

#### 442 443 444 4.6 TYPICAL REASONING PATTERNS 445

446  
 447 The experimental results above have demonstrated RLoT’s capability to enhance LLM reasoning  
 448 with task-specific logical structures. In Table 5, we summarize typical patterns observed in the  
 449 logical structures generated by RLoT when solving tasks from the MATH, GPQA, and StrategyQA  
 450 benchmarks using Qwen2.5-14B-Instruct.

451 From the 2-step patterns, we observe that the *Reason-Refine* mode emerges frequently. Particularly  
 452 in MATH and GPQA, which require massive mathematical calculations, this pattern compensates  
 453 for the LLMs’ relatively poor calculation abilities, leading to more reliable results. In the 3-step  
 454 patterns, operations like *Decompose* and *Debate* are frequently employed, which help break down  
 455 challenging problems or facilitate discussions to explore potential solutions. Additionally, the *Refine*  
 456 operation is often used before and after the *Decompose* and *Debate* steps, ensuring correct integration  
 457 with preceding and succeeding reasoning processes. These typical reasoning patterns exhibit strong  
 458 interpretability, further validating the capability of RLoT to generate task-specific logical structures  
 459 that enhance LLM reasoning. For better understanding, we provide an example in Appendix B.1 to  
 460 illustrate how these logical structures empower the correct solution to a problem.

#### 461 462 4.7 ABLATION STUDIES, ANALYSES, AND DISCUSSIONS 463

464  
 465 To better justify our designs, we verify the effectiveness of our logic block designs with ablation  
 466 studies. To examine the impact of each block, we remove each of them and train a navigator, where  
 467 results in Table 6 confirm that all blocks are effective.

470 Table 6: Ablation study results for each logic block.

LLM	Ablation	MATH	GPQA	StrategyQA	Average
Qwen2.5-Instruct-7B	w/o Decompose	75.42	31.92	77.00	61.45
	w/o Debate	74.02	36.61	77.58	62.74
	w/o Refine	75.76	41.29	72.93	63.33
	<b>Full RLoT</b>	<b>76.70</b>	<b>44.64</b>	<b>79.04</b>	<b>66.79</b>
GPT-4o-mini	w/o Decompose	76.08	41.74	79.33	65.72
	w/o Debate	74.68	45.31	78.75	66.25
	w/o Refine	76.44	51.56	74.38	67.46
	<b>Full RLoT</b>	<b>77.36</b>	<b>54.02</b>	<b>82.68</b>	<b>71.35</b>

481  
 482 Also, we conducted a series of analyses, and discussions. We quantify the contribution of each logic  
 483 block in Appendix I. We discuss the reliability of the self-evaluation state design in Appendix F. We  
 484 analyze the necessity of employing RL to train the navigator model in Appendix G. We illustrate the  
 485 role of the process reward model (PRM) in training the navigator in Appendix H.

486 5 RELATED WORKS  
487488 5.1 INFERENCE-TIME REASONING OF LLMs  
489

490 To improve LLMs' reasoning capability, plentiful research has investigated inference-time techniques  
491 without the need for model updates. On the one hand, predefined external logical structures are widely  
492 applied as basic solutions. Most notably, Chain-of-Thought (CoT) (Wei et al., 2022) incorporates  
493 intermediate reasoning steps within the prompt, enhancing the model's abilities in complex tasks. As  
494 a subsequent advancement, CoT with Self-Consistency (CoT-SC) (Wang et al., 2023) further refines  
495 CoT by generating multiple reasoning chains and filtering out those that do not meet self-consistency,  
496 thus increasing the reliability. Moreover, the concept of Tree-of-Thoughts (ToT) (Yao et al., 2023)  
497 and Graph-of-Thoughts (GoT) (Besta et al., 2024) has been introduced, where the reasoning process  
498 is represented as a graph, enabling exploration and backtracking from more promising outcomes.  
499 Despite their success, these methods rely on task-agnostic logical structures that are applied uniformly  
500 across diverse tasks, lacking flexibility.

501 On the other hand, recent researchers have proposed more adaptive inference-time approaches. From  
502 hiring the decompose-analyze-rethink procedure (Xue et al., 2024) to utilizing Monte Carlo Tree  
503 Search to discover more effective logical structures (Wu et al., 2024; Wang et al., 2024a), these  
504 approaches dynamically compose appropriate logical structures for specific tasks to enhance the  
505 reasoning performance. However, the search-and-trail design incurs massive extra computational  
506 cost, limiting the efficiency of reasoning. In contrast, the proposed RLoT method employs RL to  
507 train a navigator agent. With the trained navigator, our design can directly select and combine basic  
508 logic blocks and generate task-specific logical structures, enabling more adaptive reasoning while  
509 eliminating the searching cost.

510 5.2 RL AND LLMs  
511

512 RL has become vital in the development of LLMs (Xu et al., 2025; Hao et al., 2025). One important  
513 direction is aligning LLMs with human preferences, where the key method is Reinforcement Learning  
514 from Human Feedback (RLHF) (Christiano et al., 2017; Ouyang et al., 2022). In RLHF, LLMs are  
515 fine-tuned as actors in RL based on feedback derived from human preferences. Recently, extensive  
516 helpful and harmless LLMs have been created using RLHF (Bai et al., 2022; Casper et al., 2023).  
517 Beyond this, RL is also applied to enhance the reasoning capability of LLMs. In this context, reward  
518 signals are derived from Outcome Reward Models (ORM) (Kazemnejad et al., 2024) or Process  
519 Reward Models (PRM) (Lightman et al., 2023; Wang et al., 2024c; Luo et al., 2024), providing  
520 feedbacks for the LLMs' reasoning process during fine-tuning. By applying RL to LLMs with these  
521 rewards, LLMs can iteratively improve their performance in multi-step reasoning tasks (Havrilla  
522 et al., 2024a;b; Shao et al., 2024; DeepSeek-AI et al., 2025).

523 In summary, existing RL techniques have significantly enhanced LLMs' capabilities by updating  
524 model parameters, yet such fine-tuning demands substantial computational resources when applied  
525 to each pre-trained LLM. In this paper, the proposed RLoT method applies RL at inference time,  
526 leveraging the power of RL to train a lightweight navigator rather than the entire parameters of LLMs,  
527 our method achieves low cost and wide compatibility with various pre-trained LLMs.

528 6 CONCLUSIONS  
529

531 In this paper, we propose RL-of-Thoughts (RLoT), an inference-time technique that utilizes RL to  
532 train a navigator model, which adaptively constructs task-specific logical structures, and thereby  
533 enhances the reasoning capabilities of LLMs. Through extensive experiments across various benchmarks,  
534 we demonstrate the effectiveness of our method in improving the reasoning capability of  
535 various widely known LLMs. Additionally, we show the strong transferability and efficiency of  
536 our approach, where the trained navigator model can effectively transfer across different LLMs  
537 and unseen reasoning tasks. With fewer than 3K parameters, our navigator model enables multiple  
538 sub-10B LLMs to attain performance comparable to larger LLMs with up to 10 times the parameter  
539 size. Our work highlights the potential of RL at inference time in enhancing the reasoning capabilities  
of LLMs, paving the way for more adaptive and efficient LLM reasoning in the future.

540 ETHICS STATEMENT  
541542 This study uses fully open-source or publicly available models and datasets, adhering to their  
543 respective licenses. All resources are properly cited in Section 1 and Section 4. The selected datasets  
544 and models are well-established, representative, and free from bias or discrimination.  
545546 REPRODUCIBILITY STATEMENT  
547548 For Reproducibility, we describe the general experimental settings in Section 4; we list the imple-  
549 mentation details in Appendix J; and our source code are anonymously open source at <https://anonymous.4open.science/r/RL-LLM-Reasoning-1A30>.  
550552 REFERENCES  
553554 Deciphering and enhancing commonsense reasoning in llms from the perspective of intrinsic factual  
555 knowledge retrieval. 2024. <https://openreview.net/forum?id=MbtA7no8Ys>.556 Asma Ben Abacha, Wen-wai Yim, Yujuan Fu, Zhaoyi Sun, Meliha Yetisgen, Fei Xia, and Thomas  
557 Lin. Medec: A benchmark for medical error detection and correction in clinical notes. *arXiv*  
558 *preprint arXiv:2412.19260*, 2024.559 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,  
560 Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.  
561 *arXiv preprint arXiv:2303.08774*, 2023.562 Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. Large language models  
563 for mathematical reasoning: Progresses and challenges. *arXiv preprint arXiv:2402.00157*, 2024.564 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain,  
565 Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with  
566 reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.567 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi,  
568 Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoefer. Graph  
569 of thoughts: Solving elaborate problems with large language models. In *AAAI*, pp. 17682–17690.  
570 AAAI Press, 2024.571 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal,  
572 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel  
573 Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler,  
574 Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott  
575 Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya  
576 Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020.577 Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémie Scheurer, Javier  
578 Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. Open problems  
579 and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint*  
580 *arXiv:2307.15217*, 2023.581 Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan  
582 Yi, Cunxiang Wang, Yidong Wang, et al. A survey on evaluation of large language models. *ACM*  
583 *Transactions on Intelligent Systems and Technology*, 15(3):1–45, 2024.584 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep  
585 reinforcement learning from human preferences. *Advances in neural information processing*  
586 *systems*, 30, 2017.587 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
588 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
589 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,  
590 2021.

594 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,  
 595 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,  
 596 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao  
 597 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,  
 598 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,  
 599 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,  
 600 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang  
 601 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,  
 602 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,  
 603 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,  
 604 Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang,  
 605 Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L.  
 606 Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shengfeng Ye, Shiyu Wang,  
 607 Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng  
 608 Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng  
 609 Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan  
 610 Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang,  
 611 Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen,  
 612 Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li,  
 613 Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang,  
 614 Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan,  
 615 Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia  
 616 He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong  
 617 Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha,  
 618 Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang,  
 619 Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,  
 620 Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen  
 621 Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025.  
 622 URL <https://arxiv.org/abs/2501.12948>.

623 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
 624 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
*arXiv preprint arXiv:2407.21783*, 2024.

625 Yao Fu, Litu Ou, Mingyu Chen, Yuhao Wan, Hao Peng, and Tushar Khot. Chain-of-thought hub:  
 626 A continuous effort to measure large language models' reasoning performance. *arXiv preprint*  
 627 *arXiv:2305.17306*, 2023.

628 Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did Aristotle  
 629 Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. *Transactions*  
 630 *of the Association for Computational Linguistics (TACL)*, 2021.

631 Haixia Han, Jiaqing Liang, Jie Shi, Qianyu He, and Yanghua Xiao. Small language model can  
 632 self-correct. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp.  
 633 18162–18170, 2024.

634 Qianyue Hao, Jingyang Fan, Fengli Xu, Jian Yuan, and Yong Li. Hlm-cite: Hybrid language model  
 635 workflow for text-based scientific citation prediction. *arXiv preprint arXiv:2410.09112*, 2024.

636 Qianyue Hao, Lin Chen, Xiaoqian Qi, Yuan Yuan, Zefang Zong, Hongyi Chen, Keyu Zhao, Shengyuan  
 637 Wang, Yunke Zhang, Jian Yuan, and Yong Li. Reinforcement learning in the era of large language  
 638 models: Challenges and opportunities. *researchgate*, 2025.

639 Alex Havrilla, Yuqing Du, Sharath Chandra Raparthy, Christoforus Nalmpantis, Jane Dwivedi-Yu,  
 640 Maksym Zhuravinskyi, Eric Hambro, Sainbayar Sukhbaatar, and Roberta Raileanu. Teaching large  
 641 language models to reason with reinforcement learning. *arXiv preprint arXiv:2403.04642*, 2024a.

642 Alex Havrilla, Sharath Raparthy, Christoforus Nalmpantis, Jane Dwivedi-Yu, Maksym Zhuravinskyi,  
 643 Eric Hambro, and Roberta Raileanu. Glore: When, where, and how to improve llm reasoning via  
 644 global and local refinements. *arXiv preprint arXiv:2402.10963*, 2024b.

648 Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob  
 649 Steinhardt. Aligning ai with shared human values. *Proceedings of the International Conference on*  
 650 *Learning Representations (ICLR)*, 2021a.

651

652 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob  
 653 Steinhardt. Measuring massive multitask language understanding. *Proceedings of the International*  
 654 *Conference on Learning Representations (ICLR)*, 2021b.

655 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 656 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*  
 657 *preprint arXiv:2103.03874*, 2021c.

658

659 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-  
 660 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*  
 661 *arXiv:2410.21276*, 2024.

662

663 Fangkai Jiao, Geyang Guo, Xingxing Zhang, Nancy F Chen, Shafiq Joty, and Furu Wei. Preference  
 664 optimization for reasoning with pseudo feedback. *arXiv preprint arXiv:2411.16345*, 2024.

665

666 Katikapalli Subramanyam Kalyan. A survey of gpt-3 family large language models including chatgpt  
 667 and gpt-4. *Natural Language Processing Journal*, pp. 100048, 2023.

668

669 Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy,  
 670 Aaron Courville, and Nicolas Le Roux. Vineppo: Unlocking rl potential for llm reasoning through  
 671 refined credit assignment. *arXiv preprint arXiv:2410.01679*, 2024.

672

673 Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vard-  
 674 hamanan, Saiful Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, et al. Dspy: Compiling  
 675 declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*,  
 676 2023.

677

678 Shanu Kumar, Saish Mendke, Karody Lubna Abdul Rahman, Santosh Kurasa, Parag Agrawal, and  
 679 Sandipan Dandapat. Enhancing zero-shot chain of thought prompting via uncertainty-guided  
 680 strategy selection. *arXiv preprint arXiv:2412.00353*, 2024.

681

682 Xiaochong Lan, Yiming Cheng, Li Sheng, Chen Gao, and Yong Li. Depression detection on social  
 683 media with large language models. *arXiv preprint arXiv:2403.10750*, 2024a.

684

685 Xiaochong Lan, Chen Gao, Depeng Jin, and Yong Li. Stance detection with collaborative role-infused  
 686 llm-based agents. In *Proceedings of the International AAAI Conference on Web and Social Media*,  
 687 volume 18, pp. 891–903, 2024b.

688

689 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan  
 690 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. *arXiv preprint*  
 691 *arXiv:2305.20050*, 2023.

692

693 Zihan Liu, Yang Chen, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. Acemath: Advancing  
 694 frontier math reasoning with post-training and reward modeling. *arXiv preprint arXiv:2412.15084*,  
 695 2024.

696

697 Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Harsh Lara, Yunxuan Li, Lei Shu, Yun  
 698 Zhu, Lei Meng, Jiao Sun, et al. Improve mathematical reasoning in language models by automated  
 699 process supervision. *arXiv preprint arXiv:2406.06592*, 2024.

700

701 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri  
 702 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement  
 703 with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.

704

705 Ning Miao, Yee Whye Teh, and Tom Rainforth. Selfcheck: Using llms to zero-shot check their own  
 706 step-by-step reasoning. In *ICLR*. OpenReview.net, 2024.

707

708 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare,  
 709 Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control  
 710 through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

702 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong  
 703 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow  
 704 instructions with human feedback. *Advances in neural information processing systems*, 35:27730–  
 705 27744, 2022.

706 Zhenting Qi, Mingyuan Ma, Jiahang Xu, Li Lyna Zhang, Fan Yang, and Mao Yang. Mutual reasoning  
 707 makes smaller llms stronger problem-solvers. *arXiv preprint arXiv:2408.06195*, 2024.

708

709 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani,  
 710 Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark.  
 711 *arXiv preprint arXiv:2311.12022*, 2023.

712 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 713 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical  
 714 reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

715

716 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun  
 717 Xiao, Chenzhuang Du, Chonghua Liao, Chunling Tang, Congcong Wang, Dehao Zhang, Enming  
 718 Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, Haiqing Guo, Han  
 719 Zhu, Hao Ding, Hao Hu, Hao Yang, Hao Zhang, Haotian Yao, Haotian Zhao, Haoyu Lu, Haoze Li,  
 720 Haozhen Yu, Hongcheng Gao, Huabin Zheng, Huan Yuan, Jia Chen, Jianhang Guo, Jianlin Su,  
 721 Jianzhou Wang, Jie Zhao, Jin Zhang, Jingyuan Liu, Junjie Yan, Junyan Wu, Lidong Shi, Ling Ye,  
 722 Longhui Yu, Mengnan Dong, Neo Zhang, Ningchen Ma, Qiwei Pan, Qucheng Gong, Shaowei Liu,  
 723 Shengling Ma, Shupeng Wei, Sihan Cao, Siying Huang, Tao Jiang, Weihao Gao, Weimin Xiong,  
 724 Weiran He, Weixiao Huang, Wenhai Wu, Wenyang He, Xianghui Wei, Xianqing Jia, Xingzhe Wu,  
 725 Xinran Xu, Xinxing Zu, Xinyu Zhou, Xuehai Pan, Y. Charles, Yang Li, Yangyang Hu, Yangyang  
 726 Liu, Yanru Chen, Yejie Wang, Yibo Liu, Yidao Qin, Yifeng Liu, Ying Yang, Yiping Bao, Yulun Du,  
 727 Yuxin Wu, Yuzhi Wang, Zaida Zhou, Zhaoji Wang, Zhaowei Li, Zhen Zhu, Zheng Zhang, Zhexu  
 728 Wang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Ziyao Xu, and Zonghan Yang. Kimi k1.5: Scaling  
 729 reinforcement learning with llms, 2025. URL <https://arxiv.org/abs/2501.12599>.

730 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
 731 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation  
 732 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

733 Hieu Tran, Zonghai Yao, Junda Wang, Yifan Zhang, Zhichao Yang, and Hong Yu. Rare: Retrieval-  
 734 augmented reasoning enhancement for large language models. *arXiv preprint arXiv:2412.02830*,  
 735 2024.

736 Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-  
 737 learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.

738 Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Dian Yu, Haitao Mi, Jinsong Su, and Dong Yu.  
 739 Litesearch: Efficacious tree search for llm. *arXiv preprint arXiv:2407.00320*, 2024a.

740 Chaojie Wang, Yanchen Deng, Zhiyi Lyu, Liang Zeng, Jujie He, Shuicheng Yan, and Bo An. Q\*: Im-  
 741 proving multi-step reasoning for llms with deliberative planning. *arXiv preprint arXiv:2406.14283*,  
 742 2024b.

743 Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang  
 744 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In  
 745 *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume  
 746 1: Long Papers)*, pp. 9426–9439, 2024c.

747 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha  
 748 Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language  
 749 models. In *ICLR*. OpenReview.net, 2023.

750 Yingxu Wang and Vincent Chiew. On the cognitive process of human problem solving. *Cognitive  
 751 systems research*, 11(1):81–92, 2010.

752

753 Ziyu Wang, Tom Schaul, Matteo Hessel, Hado Hasselt, Marc Lanctot, and Nando Freitas. Dueling  
 754 network architectures for deep reinforcement learning. In *International conference on machine  
 755 learning*, pp. 1995–2003. PMLR, 2016.

756 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
 757 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
 758 *neural information processing systems*, 35:24824–24837, 2022.

759

760 Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, and Jianhua Tao. Beyond  
 761 examples: High-level automated reasoning paradigm in in-context learning via mcts. *arXiv preprint*  
 762 *arXiv:2411.18478*, 2024.

763 Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan,  
 764 Jiahui Gong, Tianjian Ouyang, Fanjin Meng, et al. Towards large reasoning models: A survey of  
 765 reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*, 2025.

766

767 Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton  
 768 Murray, and Young Jin Kim. Contrastive preference optimization: Pushing the boundaries of llm  
 769 performance in machine translation. *arXiv preprint arXiv:2401.08417*, 2024.

770

771 Shangzi Xue, Zhenya Huang, Jiayu Liu, Xin Lin, Yuting Ning, Binbin Jin, Xin Li, and Qi Liu.  
 772 Decompose, analyze and rethink: Solving intricate problems with human-like reasoning cycle. In  
 773 *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

774

775 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,  
 776 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*  
 777 *arXiv:2412.15115*, 2024a.

778

779 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-  
 780 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical  
 781 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024b.

782

783 Cheng Yang, Chufan Shi, Siheng Li, Bo Shui, Yujiu Yang, and Wai Lam. Llm2: Let large language  
 784 models harness system 2 reasoning. *arXiv preprint arXiv:2412.20372*, 2024c.

785

786 Ling Yang, Zhaochen Yu, Tianjun Zhang, Shiyi Cao, Minkai Xu, Wentao Zhang, Joseph E Gonzalez,  
 787 and Bin Cui. Buffer of thoughts: Thought-augmented reasoning with large language models.  
 788 *Advances in Neural Information Processing Systems*, 37:113519–113544, 2024d.

789

790 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan.  
 791 Tree of thoughts: Deliberate problem solving with large language models. In *NeurIPS*, 2023.

792

793 Xiaodong Yu, Ben Zhou, Hao Cheng, and Dan Roth. Reasonagain: Using extractable symbolic  
 794 programs to evaluate mathematical reasoning. *arXiv preprint arXiv:2410.19056*, 2024.

795

796 Di Zhang, Jianbo Wu, Jingdi Lei, Tong Che, Jiatong Li, Tong Xie, Xiaoshui Huang, Shufei Zhang,  
 797 Marco Pavone, Yuqiang Li, et al. Llama-berry: Pairwise optimization for o1-like olympiad-level  
 798 mathematical reasoning. *arXiv preprint arXiv:2410.02884*, 2024a.

799

800 Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xionghui Chen, Jiaqi Chen, Mingchen Zhuge,  
 801 Xin Cheng, Sirui Hong, Jinlin Wang, et al. Aflow: Automating agentic workflow generation. *arXiv*  
 802 *preprint arXiv:2410.10762*, 2024b.

803

804 Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu,  
 805 Jingren Zhou, and Junyang Lin. The lessons of developing process reward models in mathematical  
 806 reasoning. *arXiv preprint arXiv:2501.07301*, 2025.

807

808 Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min,  
 809 Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv*  
 810 *preprint arXiv:2303.18223*, 2023.

811

812 Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu,  
 813 Yanjun Lyu, Peng Shu, Xiaowei Yu, et al. Evaluation of openai o1: Opportunities and challenges  
 814 of agi. *arXiv preprint arXiv:2409.18486*, 2024.

815

816 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,  
 817 Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. Least-to-most prompting enables  
 818 complex reasoning in large language models. In *ICLR*. OpenReview.net, 2023.

810 Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. Toolqa: A dataset for llm  
811 question answering with external tools. *Advances in Neural Information Processing Systems*, 36:  
812 50117–50143, 2023.

813

814

815

816

817

818

819

820

821

822

823

824

825

826

827

828

829

830

831

832

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

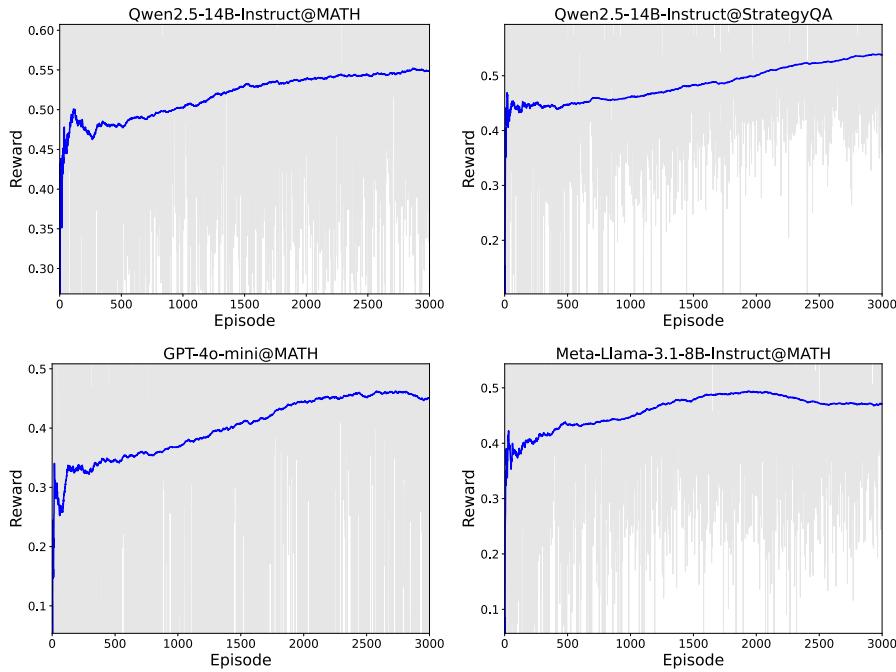
861

862

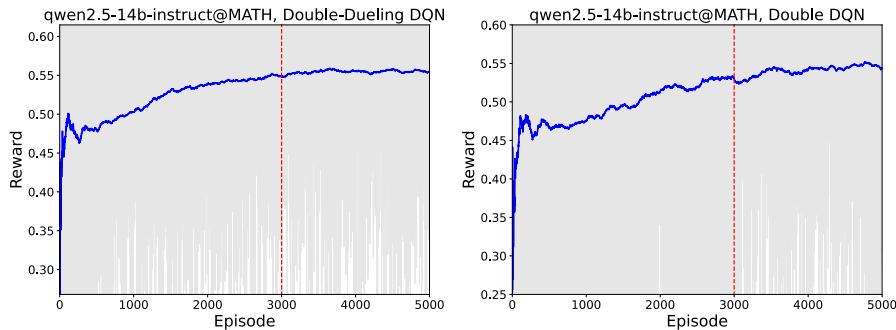
863

864 A TRAINING DETAILS  
865866 A.1 LEARNING CURVES  
867

868 In Figure 3, we show the learning curves during RL training of all navigator models used in experiments  
869 in the main text. We use a sliding window averaging to smooth the reward, and from the  
870 learning curves, we observe good convergence of the RL training.

894 Figure 3: Learning curves during RL training of navigator models.  
895896 A.2 ANALYSIS  
897

898 In Figure 4, we extend the training duration to 5,000 episodes. We observe that the model essentially  
899 reaches convergence around 3,000 episodes, with only marginal performance gains achieved thereafter.  
900 Additionally, we explore training the navigator model using the Double-DQN algorithm. While the  
901 training stability decreases slightly due to the absence of the Dueling network architecture, the model  
902 ultimately converges to a performance level comparable to that of the Double-Dueling-DQN.  
903

915 Figure 4: 5k episodes learning curves during RL training with different algorithms.  
916

917

918 **B CASE STUDY**919 **B.1 CORRECT CASES**920 **Problem:** How many square units are in the region satisfying the inequalities  $y \geq |x|$  and  $y \leq -|x| + 3$ ?921 **Few-shot CoT**

922

923 **Problem:** (Problem).924 Please generate the answer for the problem. Wrap the answer with **boxed**{(answer)}.

925

926 To find the area of the region satisfying the inequalities  $y \geq |x|$  and  $y \leq -|x| + 3$ , we need to analyze the boundaries defined by these inequalities.

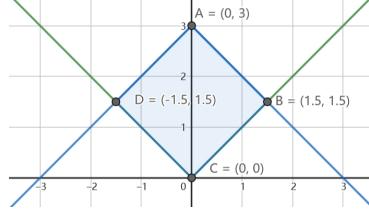
927 To find the intersection points of these lines, we solve the equations:

928 - For  $x \geq 0$ :  $y = x$  and  $y = -x + 3$ .929 - For  $x < 0$ :  $y = -x$  and  $y = x + 3$ .

930 ...

931 **Correct algebraic calculation.**932 The region of interest is a quadrilateral with vertices at  $(0,0)$ ,  $(1.5, 1.5)$ ,  $(-1.5, 1.5)$ , and  $(0, 3)$ . This quadrilateral is symmetric about the y-axis and can be divided into two congruent right triangles.933 **Each triangle has a base of  $\frac{3}{2}$  and a height of  $\frac{3}{2}$ .** The area of one triangle is  $\frac{9}{8}$ . **Wrong geometric understanding.**

934 ...

935 Thus, the area of the region is: **boxed{2.25}**. 937 **Figure 5:** A case study comparing few-shot CoT and RLoT on a representative problem in the MATH dataset. Results are both generated by Qwen2.5-14B-Instruct.938 As shown in Figure 5, we analyze a representative case to illustrate the detailed working procedure of RLoT. The problem is sampled from the test set of MATH (Hendrycks et al., 2021c) benchmark, which involves calculating the area between two functional curves, requiring a combination of algebraic and geometric knowledge. While the Few-shot CoT method attempts to generate the answer in one step, RLoT adopts a task-specific reasoning pathway, which includes *Reason-Refine-Debate*, a typical pattern for solving complex math problems shown in Section 4.6.939 Both Few-shot CoT and RLoT successfully calculate the intersection points of the curves. However, the response of Few-shot CoT fails to accurately identify the shape of the shadowed area, leading to incorrect bases for the triangles and, ultimately, an incorrect result. In contrast, RLoT utilizes the *Refine* action to clarify the intermediate results and correctly identifies the shadowed area as a rhombus. Subsequently, the *Debate* action proposes multiple plans and prioritizes directly using the area formula for a rhombus over dividing it into triangles. By avoiding unnecessary geometric division, RLoT achieves the correct answer.940 **B.2 FAILURE CASES**

941 To provide a comprehensive understanding of the boundaries of RLoT, we also analyze typical failure cases observed during our experiments. Common failures can be categorized into two primary modes: (1) Navigator input errors caused by self-evaluation noise, and (2) Navigator output errors caused by suboptimal action selection.

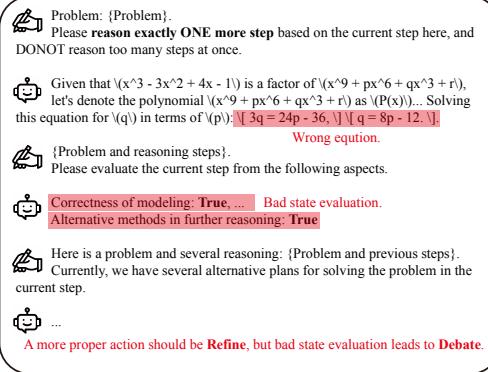
942 As shown in Figure 6, failure happens due to self-evaluation noise. The robustness of the navigator depends partly on the accuracy of the state representation derived from the LLM self-evaluation. In this case, the LLM generates an incorrect equation (wrongly deriving  $3q = 24p - 36$ ) but fails to identify this error during the self-evaluation phase, hallucinating a "True" status for correctness of modeling. This noise in the state input misleads the navigator: instead of selecting the necessary "Refine" action to correct the mistake, the Navigator perceives the current step as valid and selects "Debate" to explore alternative plans.943 **RLoT**944 **Problem:** (Problem).  
945 Please reason exactly ONE more step based on the current step here, and DONOT reason too many steps at once.

946 Certainly, let's start by analyzing the given inequalities...

947 For  $x \geq 0$ ,  $|x| = x$ . For  $x < 0$ ,  $|x| = -x$ . Now we know the points of intersection are  $(\frac{3}{2}, \frac{3}{2})$  and  $(-\frac{3}{2}, \frac{3}{2})$ .948 **Correct algebraic calculation.**949 **Please check and refine** the current thought here, and DONOT conduct further reasoning or calculation.950 **Analysis so far is correct...** The next step would indeed be to calculate the area of the region bounded by these lines. **This region forms a diamond (or rhombus) shape** centered at the origin.951 **Correct geometric understanding.**952 Please propose three **different alternative plans** for solving the problem in the current step.953 **Plan1:** Use the formula for the area of a rhombus.  
954 **Plan2:** Integrate the area between the curves.  
955 **Plan3:** Divide the region into triangles.956 **The most promising plan is Plan1: Given the diagonals of the rhombus, we can now calculate the area.** **Correct method**957 ...  
958 Thus, the area of the region satisfying the inequalities  $|y| \geq |x|$  and  $y \leq -|x| + 3$  is  $(4.5)$  square units.

972  
 973     **Problem:** The polynomial  $x^3 - 3x^2 + 4x - 1$  is a factor of  $x^9 + px^6 + qx^3 + r$ . Enter the ordered triple  $(p, q, r)$ ?

974  
 975     **State self-evaluation error**



988  
 989     Figure 6: A failure case study that is caused by wrong state self-evaluation.

990  
 991

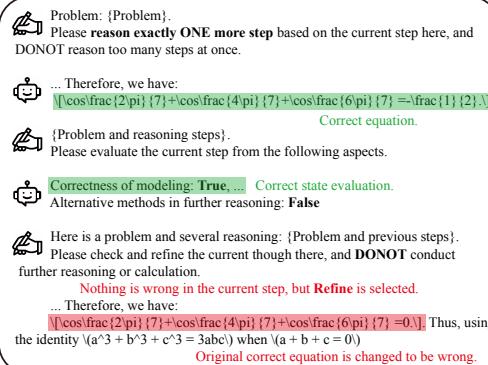
992     As shown in Figure 7, failure happens due to action selection error, where the navigator selects  
 993     a suboptimal action despite receiving accurate state information. In this case, the LLM derives a  
 994     correct intermediate result  $(\cos(2\pi/7) + \dots = -1/2)$  and correctly evaluates it as "True". However,  
 995     the navigator triggers a "Refine" action instead of proceeding with reasoning or termination. This  
 996     unnecessary intervention forces the LLM to modify a correct step, leading to "over-correction" where  
 997     valid reasoning is replaced with an erroneous identity, ultimately causing the solution to fail.

998  
 999

1000     **Problem:** Compute  

$$[\cos^3 \frac{2\pi}{7} + \cos^3 \frac{4\pi}{7} + \cos^3 \frac{8\pi}{7}]$$

1001     **Action selection error**



1016     Figure 7: A failure case study that is caused by wrong action selection.

1017

1018     While these failure modes exist, our extensive quantitative results demonstrate that RLoT achieves  
 1019     performance gains over baselines, indicating that these failure cases represent a minority of the  
 1020     inference trajectories. In our design, the RL training process effectively optimizes the navigator to  
 1021     minimize action selection errors, and the structured self-evaluation prompt is designed to mitigate  
 1022     state noise errors as much as possible.

1023

1024

1025

## 1026 C COMPUTATIONAL OVERHEAD AND LATENCY

1028 As shown in Table 7, RLoT improves the reasoning ability of LLMs with a low cost comparing  
 1029 with complex inference-time techniques. The results show that at inference-time, RLoT’s token  
 1030 consumption only accounts for 78 % of ToT, and the consumption is comparable to CoT-SC. This  
 1031 is because the trained navigator can directly select appropriate logic blocks without searching like  
 1032 ToT, improving the performance and reducing cost. Besides, our navigator with <3k parameters has  
 1033 almost negligible inference cost itself.

1034 Though RL training incurs extra costs, our lightweight navigator makes training efficient, and the  
 1035 trained navigator can transfer across tasks, diluting the cost. Specifically, our navigator is trained for  
 1036 3000 episodes, and each episode consumes an average of 5004 input and 1448 output tokens. Shared  
 1037 across all test questions, it only adds less than 5% input and output tokens per question.

1039 **Table 7: Token consumption per question across tasks. The bold numbers indicate our method.**

Model	Method	AIME24		AMC23		MATH		GSM8K		GPQA		MMLU-STEM		StrategyQA		Average	
		Input	Output	Input	Output	Input	Output	Input	Output	Input	Output	Input	Output	Input	Output	Input	Output
Qwen2.5-14B-Instruct	Direct QA	653	1433	619	1131	46	362	55	183	159	460	82	210	30	259	235	577
	Zero-shot CoT	667	1489	633	1007	58	367	66	206	175	546	87	334	33	376	246	618
	Few-shot CoT	1955	1341	1921	940	466	336	1089	135	1557	589	1114	259	421	223	1218	546
	CoT-SC	7823	5330	7685	4109	2014	1320	4339	543	6295	1945	4448	1039	1586	897	4884	2169
	ToT	5994	10844	3836	8029	4983	6277	3570	4312	6063	7404	4797	5363	5762	6491	5001	6960
Qwen2.5-7B-Instruct	RLoT (ours)	<b>22004</b>	<b>3888</b>	<b>17963</b>	<b>3463</b>	<b>3735</b>	<b>1109</b>	<b>2310</b>	<b>634</b>	<b>5501</b>	<b>1485</b>	<b>2923</b>	<b>791</b>	<b>2348</b>	<b>756</b>	<b>8112</b>	<b>1732</b>
	Direct QA	642	1714	644	1443	44	366	58	182	159	546	66	404	22	279	234	705
	Zero-shot CoT	667	1385	663	1180	47	393	78	222	184	522	82	413	32	401	250	645
	Few-shot CoT	4230	1789	3875	1501	483	581	1074	188	1556	442	1128	221	396	348	1820	724
	CoT-SC	16890	7131	15341	5821	1932	1443	4379	750	6290	1800	4433	883	1612	1391	7268	2746
Llama3.1-8B-Instruct	ToT	5289	9180	3741	7621	6679	7525	5792	5301	5771	6489	5058	4669	3807	2013	5162	6114
	RLoT (ours)	<b>23093</b>	<b>4168</b>	<b>12876</b>	<b>2896</b>	<b>1437</b>	<b>807</b>	<b>1283</b>	<b>476</b>	<b>2112</b>	<b>889</b>	<b>1028</b>	<b>434</b>	<b>1735</b>	<b>606</b>	<b>6223</b>	<b>1468</b>
	Direct QA					44	437	53	180	174	811	71	406	13	201	71	407
	Zero-shot CoT					73	495	64	215	175	661	91	365	20	418	85	431
	Few-shot CoT					510	472	1101	173	1568	769	1127	168	417	272	945	371
GPT-4o-mini	CoT-SC					2049	1885	4325	808	6120	3078	4429	808	1582	1071	3701	1530
	ToT					3474	4307	5212	4780	5615	3794	5331	5999	4320	4868	4790	4750
	RLoT (ours)					<b>4481</b>	<b>1981</b>	<b>2842</b>	<b>1028</b>	<b>9818</b>	<b>3387</b>	<b>4408</b>	<b>1486</b>	<b>3049</b>	<b>1094</b>	<b>4920</b>	<b>1795</b>
	Direct QA	601	1356	571	1184	43	404	54	223	165	397	86	175	13	201	219	563
	Zero-shot CoT	614	1435	585	1192	50	422	68	260	166	525	89	326	16	382	226	649
DeepSeek-R1-Distill-Qwen-7B	Few-shot CoT	1853	1365	1823	1202	513	408	1107	205	1575	505	1119	275	409	254	1200	602
	CoT-SC	7414	5493	7294	5097	2030	1554	4442	823	6305	2027	4360	1005	1601	1014	4778	2430
	ToT	4639	4929	5726	4521	4534	5649	5890	5810	5175	5841	4950	5324	3833	1467	4964	4792
	RLoT (ours)	<b>26529</b>	<b>4778</b>	<b>19513</b>	<b>3973</b>	<b>3483</b>	<b>1575</b>	<b>3069</b>	<b>1048</b>	<b>4372</b>	<b>1550</b>	<b>1882</b>	<b>710</b>	<b>1612</b>	<b>673</b>	<b>8637</b>	<b>2044</b>
	Direct QA	879	11793	597	4686	541	4960	533	1364	808	5462	521	2094	436	1268	616	4518
Besides token consumption, we show the average solving time (min) per problem along with the standard deviation across 3 repeats in Table 8.	Zero-shot CoT	906	12129	640	5140	565	4697	547	1381	822	5443	574	2012	450	1344	643	4592
	Few-shot CoT	2360	7778	2086	5120	1938	3301	3253	1303	4399	5194	2063	1890	1197	1275	2470	3694
	CoT-SC	9240	32005	8556	21859	7749	13071	12976	5123	17562	20553	7883	7820	4791	5094	9822	15075
	ToT	5879	34616	4474	45524	3076	25913	6680	26033	4481	19222	2158	9604	8393	11697	5020	24658
	RLoT (ours)	<b>13556</b>	<b>21870</b>	<b>15684</b>	<b>39195</b>	<b>4003</b>	<b>8450</b>	<b>8347</b>	<b>12389</b>	<b>3734</b>	<b>18787</b>	<b>6732</b>	<b>8613</b>	<b>6445</b>	<b>10546</b>	<b>8357</b>	<b>16081</b>
RLoT plus training cost		21649	8779	16863	10664	3782	2887	3924	3217	5461	5322	3749	2509	3392	2837	8403	5174

1061 Besides token consumption, we show the average solving time (min) per problem along with the  
 1062 standard deviation across 3 repeats in Table 8.

1064 **Table 8: Average solving time (minutes) per problem across various benchmarks. Standard deviations from 3 repeats are shown in parentheses.**

Method	AIME24	AMC23	MATH	GSM8K	GPQA	MMLU-STEM	StrategyQA	Average
Direct QA	0.44(0.02)	0.25(0.01)	0.13(0.00)	0.05(0.00)	0.17(0.01)	0.08(0.00)	0.05(0.00)	0.17(0.01)
Zero-shot CoT	0.44(0.02)	0.25(0.02)	0.13(0.01)	0.06(0.00)	0.17(0.02)	0.08(0.00)	0.06(0.00)	0.17(0.01)
Few-shot CoT	0.52(0.03)	0.42(0.02)	0.17(0.01)	0.18(0.01)	0.33(0.02)	0.17(0.02)	0.10(0.01)	0.27(0.02)
CoT-SC	2.09(0.13)	1.74(0.12)	0.64(0.03)	0.71(0.07)	1.32(0.09)	0.68(0.05)	0.38(0.02)	1.08(0.07)
ToT	1.86(0.07)	1.91(0.15)	1.33(0.12)	1.35(0.11)	1.28(0.06)	0.98(0.10)	0.97(0.11)	1.38(0.10)
RLoT (ours)	2.75(0.21)	2.48(0.22)	0.57(0.02)	0.61(0.05)	0.95(0.05)	0.53(0.03)	0.53(0.04)	1.20(0.09)

1073 Our navigator can directly generate a task-specific logical trajectory, reducing the LLM interaction  
 1074 cost. Also, our lightweight navigator itself brings almost negligible cost.

1080  
1081  
1082  
1083  
1084 D PERFORMANCE COMPARISON WITH LARGER LLMs  
1085  
1086  
1087  
1088  
1089  
1090  
1091  
1092

Table 9: Performance comparison between sub-10B LLMs enhanced by RLoT, and larger LLMs with several times of parameters.

LLM	Size	Method	MATH	GSM8K	GPQA	MMLU-STEM	StrategyQA	Average	Gap
Qwen2.5-Instruct	14B	Few-shot CoT	80.00	94.80	45.50	85.06	78.60	76.79	-4.13
	72B	RLoT (ours)	80.38	94.16	51.34	88.93	81.22	79.21	-1.71
	72B	Few-shot CoT	83.10	95.80	49.00	89.80	86.90	80.92	-
Llama3.1-Instruct	8B	Few-shot CoT	48.52	84.50	33.03	70.66	72.05	61.75	-12.51
	70B	RLoT (ours)	56.56	90.07	46.88	80.56	84.42	71.70	-2.56
	70B	Few-shot CoT	68.00	95.10	46.70	84.81	76.71	74.26	-
GPT-4o	mini (8B)	Few-shot CoT	75.46	93.48	35.94	85.82	80.06	74.15	-4.44
		RLoT (ours)	77.36	93.86	54.02	88.23	82.68	79.23	+0.64
	(200B)	Few-shot CoT	76.60	93.73	53.60	87.90	81.10	78.59	-

1093  
1094 In table 9, we show the performance of the sub-10B LLMs after enhancement with RLoT. The results  
1095 indicate that our RL-based navigator, which contains fewer than 3,000 parameters, significantly  
1096 enhances the performance of sub-10B LLMs, making them comparable to much larger counterparts  
1097 with around 10 $\times$  more parameters. Specifically, our RLoT method empowers the sub-10B LLMs  
1098 to be comparable to, compensating most of the performance gap, or even surpassing their larger  
1099 counterparts, demonstrating remarkable efficiency.1100  
1101  
1102  
1103  
1104  
1105  
1106  
1107  
1108  
1109  
1110  
1111  
1112  
1113  
1114  
1115  
1116  
1117  
1118  
1119  
1120  
1121  
1122  
1123  
1124  
1125  
1126  
1127  
1128  
1129  
1130  
1131  
1132  
1133

1134 **E EXTENDED BASELINE COMPARISONS**  
11351136 **E.1 TEST-TIME SCALING BASELINES**  
11371138 In addition to math reasoning models, we also compare RLoT with multiple models which cover  
1139 different categories and designs. The details of baselines are listed below.1140 **Fixed reasoning patterns.** Baselines from this category adopts a fixed workflow to obtain the answer,  
1141 which

- **Self-Refine** (Madaan et al., 2023): This method adds a simple refine step after every reasoning step to correct potential mistakes.
- **Least-to-Most** (Zhou et al., 2023): Ask LLMs to break down a complex problem into a series of simpler subproblems and then solve them in sequence.
- **SelfCheck** (Miao et al., 2024): Utilize LLMs to check their own outputs and use the results of these checks to improve question-answering performance by conducting weighted voting on multiple solutions to the question. Here we adopt voting with 4 candidates.
- **DeAR** (Xue et al., 2024): The problem is consequently decomposed, analyzed and refined in this workflow. The workflow follows a recurrent manner to a detailed decomposition of original problem.

1153 **Tree search.** Tree-based models divide the reasoning task into sub-steps, and searches for a best path  
1154 to the final answer. Tree-based methods usually requires multiple rounds Q/A, resulting in high  
1155 reasoning cost.  
1156

- **One-step-greedy**: Using PRM model to greedily select the best thought at each step.
- **Basic Monto-carlo Tree Search**: An effective tree-searched method that explore the tree via multiple roll-outs. The MCTS is also supervised by the PRM.
- **Litesearch** (Wang et al., 2024a): An low-cost tree search method that assigns larger search budget to a more promising nodes.
- **Q\*** (Wang et al., 2024b): **Q\*** adopts a value function that estimate the "distance" between the current thought and the correct answer. Then an **A\*** algorithm, which is widely applied in shortest path problem, is used to finish tree-search.
- **rStar** (Qi et al., 2024): Use Monte Carlo Tree Search during inference to select actions.
- **AFlow** (Zhang et al., 2024b): An automated framework that uses MCTS to efficiently explore LLM agentic workflow.

1169 **Others.**  
1170

- **DSPy** (Khattab et al., 2023): A programming model that can express and optimize sophisticated LM pipelines.
- **Graph of Thoughts (GoT)** (Besta et al., 2024): A graph-based model which allows arbitrary combination of thoughts according to the manually designed graph architecture. We carefully design graphs for each specific dataset.
- **Buffer of Thoughts (BoT)** (Yang et al., 2024d): A Retrieval-Augmented Generation (RAG) model that enhance reasoning via reasoning patterns in the buffer. These patterns are continuously updated during inference time.

1180 As shown in Table 10, we compare these methods on representative datasets, GSM8K, GPQA, and  
1181 StrategyQA, which covers three different domains. Results show that our model outperforms all  
1182 baseline methods among all datasets.  
1183

By flexibly organizing logic blocks, our design can outperform these baselines. To compare the efficiency and effectiveness, we show the token consumption per question for representative baselines:

Unlike search-based methods that require multiple explorations, our design directly generates task-specific logical trajectories, allowing for better performance with a similar or lower cost.

1188  
 1189  
 1190  
 1191  
 1192  
 1193  
 1194  
 1195  
 1196

Table 10: Extended baseline comparisons. The **bold** numbers indicate the best performance in each group of experiments, and the underlined numbers indicate the best baseline method. All results are averaged across 5 repeated runs with standard deviation.

LLM	Method	MATH-500	GSM8K	GPQA	StrategyQA	Average
Qwen2.5-Instruct-7B	DirectQA	74.96±1.37	90.95±1.06	31.16±0.23	68.53±0.53	66.40±0.46
	Self-Refine	72.32±1.20	88.39±0.38	32.23±0.61	74.06±0.79	66.75±0.33
	Least-to-Most	73.28±0.84	89.10±0.82	31.79±0.54	73.60±0.64	66.94±0.25
	SelfCheck	74.36±1.60	90.24±1.09	33.84±0.61	72.87±0.77	67.83±0.35
	DeAR	69.92±0.88	87.84±0.54	35.71±0.37	71.03±0.85	66.13±0.39
	One-step-greedy	71.24±1.43	89.49±0.99	32.63±0.50	72.43±0.56	66.45±0.54
	Basic MCTS	72.40±0.79	89.51±0.72	35.13±0.59	75.72±0.80	68.19±0.19
	LiteSearch	72.52±1.29	89.86±0.86	33.17±0.30	76.04±0.65	67.90±0.52
	Q*	73.76±0.81	<u>92.28±0.61</u>	34.38±0.76	76.80±0.99	69.30±0.50
	rStar	75.88±1.02	92.05±0.89	38.17±0.32	77.09±0.83	70.80±0.68
	AFlow	75.72±1.16	90.95±0.98	37.37±0.63	77.09±0.56	70.28±0.13
	DSPy	72.48±0.86	88.75±0.77	34.11±0.68	73.89±0.94	67.31±0.49
	Graph of Thoughts (GoT)	73.24±1.06	88.76±1.07	33.26±0.66	75.69±0.91	67.74±0.39
	Buffer of Thoughts (BoT)	75.16±0.74	91.98±0.64	34.91±0.58	74.38±0.44	69.11±0.32
	RLoT (ours)	<b>76.88±1.35</b>	<b>92.95±0.63</b>	<b>44.78±0.71</b>	<b>79.56±0.64</b>	<b>73.54±0.18</b>

1204  
 1205  
 1206  
 1207  
 1208  
 1209  
 1210  
 1211  
 1212  
 1213  
 1214  
 1215  
 1216  
 1217  
 1218  
 1219  
 1220  
 1221  
 1222  
 1223  
 1224  
 1225  
 1226  
 1227

Table 11: Token consumption, solving time (minutes) per question, and performance scores for representative baselines.

LLM	Method	GSM8K			GPQA			StrategyQA			Average		
		Score	Token	Time	Score	Token	Time	Score	Token	Time	Score	Token	Time
Qwen2.5-Instruct-7B	DirectQA	91.58	240	0.05	31.25	705	0.17	68.85	301	0.05	63.89	415	0.09
	Self-Refine	88.55	678	0.12	32.59	1784	0.41	73.94	939	0.16	65.03	1134	0.23
	LiteSearch	89.76	912	0.20	33.04	2139	0.55	75.84	1228	0.24	66.21	1426	0.33
	rStar	91.96	2636	0.77	38.62	4127	1.44	77.15	2987	0.82	69.24	3483	1.01
	RLoT (ours)	92.87	1759	0.61	44.64	3001	0.95	79.04	2341	0.53	72.18	2367	0.70

1235  
 1236  
 1237  
 1238  
 1239  
 1240  
 1241

1242 E.2 MATH DOMAIN BASELINES  
12431244 To better demonstrate the effectiveness of our design, we compare RLoT with more baselines that  
1245 have been specifically designed for math reasoning tasks. We involve fine-tuning designs which are  
1246 computationally expensive.1247  
1248 

- **AceMath** (Liu et al., 2024): A supervised fine-tuning model designed specifically for  
1249 mathematical reasoning. It first develops a math-specialized reward model using public  
1250 datasets and then performs fine-tuning and reasoning guided by this reward model.
- **PFPO** (Jiao et al., 2024): A supervised fine-tuning approach guided by pseudo reward  
1251 feedback. The feedback is generated either through a self-consistency mechanism or with  
1252 the assistance of more powerful LLMs.

  
12531254 We also include inference-time designs that hire complicated algorithms like Monte-Carlo Tree  
1255 Search (MCTS).  
12561257  
1258 

- **LLM2** (Yang et al., 2024c): A lightweight model aimed at enhancing LLM reasoning  
1259 during inference. It achieves this by training a verifier to distinguish and prioritize better  
LLM-generated responses.
- **LLaMA-Berry** (Zhang et al., 2024a): An inference-time Monte Carlo Tree Search (MCTS)  
1260 method that explores reasoning paths using a trained reward model.
- **HiAR-ICL** (Wu et al., 2024): A method that matches problems with multiple reasoning  
1261 templates at inference time. These templates are previously generated using MCTS on a  
subset of the dataset.

  
1262  
12631264 Table 12: Performance comparison of RLoT’s with more baselines. The **bold** numbers indicate the  
1265 best performance in each category.  
1266

Category	Method	MATH	GSM8K	Average
Fine-tuning	AceMath (Liu et al., 2024)	<b>64.42</b>	<b>90.45</b>	<b>77.44</b>
	PFPO (Jiao et al., 2024)	57.80	89.60	73.70
Inference-time	LLM2 (Yang et al., 2024c)	48.60	88.00	68.30
	LLaMA-Berry (Zhang et al., 2024a)	54.80	89.80	72.30
	HiAR-ICL (Wu et al., 2024)	55.00	<b>90.70</b>	72.85
	RLoT (ours)	<b>56.56</b>	90.07	<b>73.32</b>

1267 We test with Llama3.1-8B-Insturct on MATH and GSM8K benchmarks and show the results in  
1268 Table 12. For the baselines, we use the performance reported in the original papers, and for RLoT,  
1269 we test with the same navigator model as in Section 4.3 in the main text. From the results, we can  
1270 observe that our method outperforms all inference-time baselines. Meanwhile, with substantially  
1271 lower computational consumption, it reaches comparable performance to some fine-tuning methods  
1272 that require modification on the parameters of LLMs.  
1273  
1274

1296 **F ANALYSES ON THE SELF-EVALUATION STATE**  
12971298 **F.1 RELIABILITY**  
12991300  
1301 Previous work shows that LLMs can evaluate and correct their own outputs (Han et al., 2024).  
1302 Meanwhile, multiple works (Madaan et al., 2023; Xue et al., 2024) have reached success with self-  
1303 correcting. Based on the literature, we design the structured self-evaluation that prompts the LLM  
1304 across granular aspects (Table 1), including the correctness of modeling and calculation. This design  
1305 aims to be more specific and reliable than the single judgment.  
13061307 To validate the robustness, we randomly sample 100 intermediate reasoning texts and manually check  
1308 the self-evaluation. We find that the LLM’s assessment was accurate in 82 out of 100 cases. Also,  
1309 we deliberately introduced modeling and calculation errors in a reasoning path, and test whether our  
1310 self-evaluation mechanism is able to correctly identify. As shown in Table 13, the problem is from  
1311 GSM8K (Cobbe et al., 2021), which includes the relationship of three variables. In addition to the  
1312 correct reasoning step, we modify it to obtain two wrong reasoning steps, which respectively makes a  
1313 mistake on modeling and calculation. The result shows that the LLM is able to identify and classify  
1314 different kinds of mistakes, demonstrating the effectiveness of self-evaluation state extraction.  
13151316 **Table 13: An example of self-assessing of intermediate states.**  
1317

1317 <b>Question</b>	1318 <b>Reasoning step</b>	1319 <b>Self-evaluated correctness</b>	
		1320 <b>Modeling</b>	1321 <b>Calculation</b>
1322 Mark has a garden with 1323 flowers. He planted plants 1324 of three different colors in it. 1325 Ten of them are yellow, 1326 and there are 80% more of 1327 those in purple. There are only 1328 25% as many green flowers as 1329 there are yellow and purple 1330 flowers. How many flowers 1331 does Mark have in his garden?	1332 In this step, we aim at calculate the number 1333 of purple flowers, which is 80% more 1334 than the yellow ones. 1335 We calculate by $10 \times (1 + 0.8) = 18$	1336 True	1337 True
	1338 In this step, we aim at calculate the number 1339 of purple flowers, which is 80% more 1340 than the yellow ones. 1341 We calculate by $10 \times 0.8 = 8$	1342 False	1343 True
	1344 In this step, we aim at calculate the number 1345 of purple flowers, which is 80% more 1346 than the yellow ones. 1347 We calculate by $10 \times (1 + 0.8) = 17$	1348 True	1349 False

1350 Based on the validation, we consider our structured self-evaluation provides a sufficiently reliable  
1351 state representation for training the navigator.  
13521353 **F.2 NOISE IMPACT**  
13541355 To investigate the navigator’s reliance on accurate state perceptions, we introduced synthetic noise  
1356 into the self-evaluation state vectors during inference. As shown in Table 14, replacing the state with  
1357 random values results in a significant performance drop, falling even below DirectQA. This confirms  
1358 that the navigator actively utilizes the state information for decision-making rather than following  
1359 a fixed, state-agnostic policy. However, the framework exhibits commendable robustness. Under  
1360 extreme conditions with 50% noise (where half of the state bits are flipped), RLoT still maintains an  
1361 average accuracy of 64.99%, being comparable to the DirectQA baseline. As the noise decreases to  
1362 30%, performance recovers rapidly to 69.11%.1363 Meanwhile, to minimize potential noise inherent in self-evaluation, we deliberately designed the eval-  
1364 uation criteria as a binary classification task rather than a multi-grade continuous scoring. Continuous  
1365 or fine-grained scoring (e.g., 1-10 scales) requires precise calibration, which is challenging for LLMs.  
1366 In contrast, a binary standard simplifies the decision boundary, significantly reducing the cognitive  
1367 load on the LLM. This design choice makes the evaluation process more robust, ensuring that the  
1368 state vector remains a stable and reliable signal for the navigator even with less capable LLMs.  
1369

1350

1351

Table 14: Impact of noise in self-evaluation states.

LLM	State	GSM8K	GPQA	StrategyQA	Average
Qwen2.5-Instruct-7B	DirectQA	91.58	31.25	68.85	63.89
	Random	89.76	28.13	65.50	61.13
	50% noise	91.74	34.82	68.41	64.99
	30% noise	92.27	43.75	71.32	69.11
	Normal	<b>92.87</b>	<b>44.64</b>	<b>79.04</b>	<b>72.18</b>

1358

1359

## F.3 ALTERNATIVE DESIGNS

1360

Furthermore, we explore alternative designs of self-evaluation:

1361

1362

- **MLP.** Use the self-evaluation vector from a fixed LLM as input, only train an MLP as the navigator.
- **LLM (fixed).** Directly prompt the LLM to select the action based on the raw reasoning trajectory.
- **LLM (finetune).** Use the raw reasoning trajectory as input, train an LLM backbone with an MLP classification head as the navigator.

1363

1364

1365

1366

1367

1368

1369

1370

1371

1372

1373

1374

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

Table 15: Comparison of alternative self-evaluation designs.

LLM	Navigator	GSM8K	GPQA	StrategyQA	Average
Qwen2.5-Instruct-7B	DirectQA	91.58	31.25	68.85	63.89
	LLM (fixed)	89.16	37.95	74.38	67.16
	LLM (finetune)	88.32	34.15	70.01	64.16
	MLP (ours)	92.87	44.64	79.04	72.18

As the results in Table 15 show, LLM (fixed) performs badly since pre-trained LLMs are not specified for the task. LLM (finetune) even performs worse. Our method injects valuable human knowledge for explicit, structured self-evaluation. In contrast, LLM (finetune) requires an implicit evaluation from raw text, which is much more complex and has more parameters to be tuned. This requires substantially more consumption and specialized design to overcome the complexity and uncertainty. Therefore, our method is a comprehensive design at a low cost. While some self-evaluation cases may be inaccurate, the overall improvement demonstrates the general benefit.

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

1404 **G ANALYSES ON THE ROLE OF RL IN TRAINING THE NAVIGATOR**  
 1405

1406 To illustrate the necessity of using RL in training the navigator, we compare the RL navigator with  
 1407 several other decision methods for deciding the next reasoning block. The methods includes the  
 1408

- 1409 • **Fixed logic sequence.** Fixed "Decompose-Reason-Refine" sequence like human thinking.
- 1410 • **Supervised-trained navigator.** We first collected 20K state-action-score samples with  
 1411 random actions. Then, we trained a model to predict the outcome for each action given a  
 1412 state, and select the action with the highest predicted score during inference.
- 1413 • **LLM as navigator.** Directly prompt another LLM to select the action based on the reasoning  
 1414 context.

1415 We also add:

- 1416 • **Repeated strong blocks.** Repeat Debate or Refine until reaching the answer.

1417 **Table 16: Comparison with alternative navigator methods on various benchmarks.**

1418 <b>LLM</b>	1419 <b>Method</b>	1420 <b>GSM8K</b>	1421 <b>GPQA</b>	1422 <b>StrategyQA</b>	1423 <b>Average</b>
1424 Qwen2.5-Instruct-7B	1425 DirectQA	1426 91.58	1427 31.25	1428 68.85	1429 63.89
	1430 Repeated refine	1431 87.72	1432 32.37	1433 72.34	1434 64.14
	1435 Repeated debate	1436 89.99	1437 33.04	1438 72.05	1439 65.03
	1440 Fixed logic sequence	1441 88.38	1442 36.16	1443 71.13	1444 65.22
	1445 Supervise-trained navigator	1446 89.84	1447 36.83	1448 70.31	1449 65.66
	1450 LLM as navigator	1451 89.16	1452 37.95	1453 74.38	1454 67.16
<b>1455 RL-trained navigator (ours)</b>		<b>1456 92.87</b>	<b>1457 44.64</b>	<b>1458 79.04</b>	<b>1459 72.18</b>

1430 The results in Table 16 show that these methods are worse than RLoT:  
 1431

- 1432 • **Fixed logic sequence.** Rigid, unable to design flexible logical structures for specific tasks.
- 1433 • **Supervise-trained navigator.** The randomly sampled state-action pairs lack the ability of  
 1434 RL to purposely exploit effective parts of the sample space.
- 1435 • **LLM as navigator.** Pre-trained LLMs are not specified for the logic selection task.

1436 These show the importance of using RL for our navigator to flexibly select logic blocks.  
 1437

## 1458 H ANALYSES ON THE PROCESS REWARD MODEL (PRM)

### 1460 H.1 THE ROLE OF PRM

1462 To investigation the role of PRM, we also test an ORM to train the navigator.

1464 Table 17: Comparison of RLoT using an Outcome Reward Model (ORM) vs. a Process Reward  
1465 Model (PRM).

1467 LLM	1468 Method	1469 GSM8K	1470 GPQA	1471 StrategyQA	1472 Average
1468 Qwen2.5-Instruct-7B	1469 DirectQA	1470 91.58	1471 31.25	1472 68.85	1473 63.89
	1469 RLoT with ORM	1470 91.96	1471 41.52	1472 73.94	1473 69.14
	<b>1470 RLoT with PRM (ours)</b>	<b>1471 92.87</b>	<b>1472 44.64</b>	<b>1473 79.04</b>	<b>1474 72.18</b>

1475 The results in Table 17 show that the PRM outperforms the ORM, indicating that the step reward  
1476 signal from PRM is crucial.

### 1477 H.2 GENERALIZATION CAPABILITY OF PRM

1478 Recent studies suggest that process reward models (PRMs) trained on mathematical data exhibit  
1479 strong transferability to related domains such as STEM (Zhang et al., 2025). To empirically verify this  
1480 and justify our reward design, we conducted a PRM calibration analysis. We calculated the correlation  
1481 between the intermediate PRM scores assigned at various reasoning steps and the correctness of the  
1482 final answer.

1483 Table 18: PRM calibration.

1484 LLM	1485 Correlation with Correctness	1486 MATH	1487 GPQA	1488 StrategyQA
1485 Qwen2.5-Instruct-7B	1486 Step1 PRM	1487 0.199	1488 0.184	1489 0.157
	1486 Step2 PRM	1487 0.222	1488 0.213	1489 0.195
	1486 Step3 PRM	1487 0.236	1488 0.215	1489 0.202
	1486 Step4 PRM	1487 0.240	1488 0.228	1489 0.218
	<b>1487 Average</b>	<b>1488 0.224</b>	<b>1489 0.210</b>	<b>1490 0.193</b>

1491 As presented in Table 18, the results reveal two insights. First, there is a consistent positive correlation  
1492 between intermediate PRM scores and the final outcome across all tested benchmarks. This confirms  
1493 that the PRM effectively gauges the quality of intermediate steps, validating its suitability as a  
1494 dense reward signal for training the navigator. Second, although the PRM was trained primarily on  
1495 mathematical data, its scores exhibit a notable correlation with solution correctness on out-of-domain  
1496 tasks, including GPQA (STEM) and StrategyQA (commonsense). This demonstrates the PRM’s  
1497 intrinsic ability to generalize its verification logic across different domains.

1498 This calibration analysis explains the underlying mechanism of our method’s transferability. Because  
1499 the math-trained PRM provides reliable feedback even on unseen tasks like GPQA and StrategyQA,  
1500 the navigator trained with these rewards can successfully learn effective reasoning policies across  
1501 diverse domains.

### 1502 H.3 IMPACT OF PRM QUALITY

1503 Further, to explore the impact of PRM quality, we test:

- 1504 • ORM: As a minimal quality PRM.
- 1505 • Degraded PRM: Add a std=0.1 Gaussian noise to our raw PRM.
- 1506 • Better PRM: Qwen2.5-Math-PRM-7B, a stronger PRM.

1507 These results in Table 19 show that a higher-quality PRM is beneficial. However, the navigator is  
1508 robustly beneficial with a lower-quality PRM, as long as the reward signal is directionally meaningful.

1512  
 1513  
 1514  
 1515  
 1516  
 1517  
 1518  
 1519  
 1520  
 1521  
 1522  
 1523  
 1524  
 1525  
 1526  
 1527  
 1528  
 1529  
 1530  
 1531  
 1532  
 1533  
 1534  
 1535

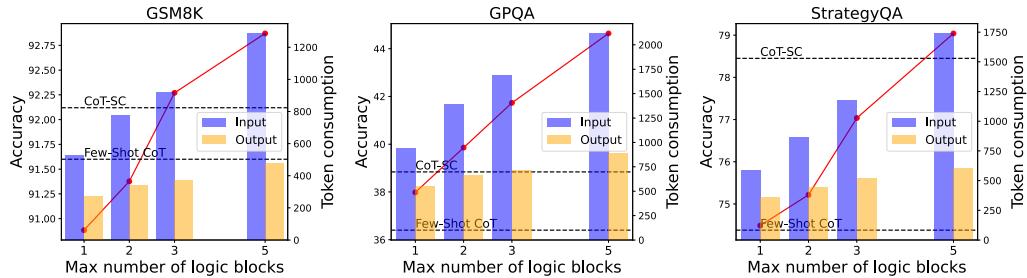
1536 Table 19: Impact of PRM quality on navigator performance.

1537 LLM	1538 PRM	1539 GSM8K	1540 GPQA	1541 StrategyQA	1542 Average
1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565	None-DirectQA	91.58	31.25	68.85	63.89
	ORM	91.96	41.52	73.94	69.14
	MathShepherd-Mistral-7B + disturb	92.49	42.41	77.15	70.68
	MathShepherd-Mistral-7B	92.87	44.64	79.04	72.18
	Qwen2.5-Math-PRM-7B	93.10	45.31	80.06	72.82

## 1566 I ANALYSES ON THE DESIGN OF LOGIC BLOCKS

### 1568 I.1 EFFECT OF LOGIC BLOCK NUMBER

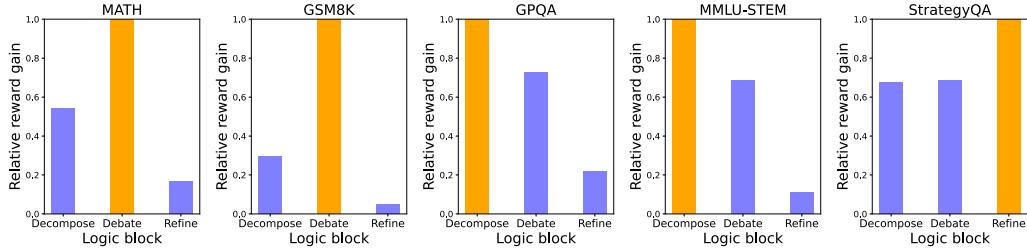
1570 We first check the test-time scaling for the logic blocks. Results in Figure 8 show that a longer  
1571 sequence of logic blocks brings a performance gain.



1582 Figure 8: Test-time scaling with different number of logic blocks.

### 1586 I.2 ROLE OF DIFFERENT LOGIC BLOCKS ON DIFFERENT DATASETS

1588 We also analyze the impact of different logic blocks on different datasets. Results in Figure 9  
1589 show the average PRM reward gain of each logic block. The most useful block varies in different  
1590 tasks, specifically Debate for math tasks, Decompose for STEM tasks, and refine for common-sense  
1591 reasoning.



1602 Figure 9: Contribution of each logic block in the reasoning sequence.

### 1605 I.3 ALTERNATIVE BLOCKS

1607 Further, we test some possible new blocks:

- **Follow-up-question:** Propose and answer a sub-question based on the current reasoning.
- **Rephrasing:** Reorganize the current reasoning.
- **Compression:** Condense the current reasoning.

1614 Table 20: Performance with additional logic blocks.

1615 LLM	1616 Logic blocks	1617 MATH	1618 GPQA	1619 StrategyQA	Average
1617 Qwen2.5-Instruct-7B	Original in RLoT	76.70	44.64	79.04	66.79
	Original + Follow-up-question	76.78	43.97	79.24	66.66
	Original + Rephrasing	77.30	45.54	80.20	67.68
	Original + Compression	76.06	43.53	77.87	65.82

1620 The results in Table 20 showed that "Rephrasing" is beneficial, while "Follow-up-question" (perhaps  
1621 self-ask-and-answer has limited effect) and "Compression" (likely due to information loss) are not.  
1622 This shows the extensibility of RLoT: it provides a framework for the community to design more  
1623 reasoning blocks to enhance the performance.

1624

1625

1626

1627

1628

1629

1630

1631

1632

1633

1634

1635

1636

1637

1638

1639

1640

1641

1642

1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

1654

1655

1656

1657

1658

1659

1660

1661

1662

1663

1664

1665

1666

1667

1668

1669

1670

1671

1672

1673

1674 **J IMPLEMENTATION DETAILS**  
16751676 **J.1 SETTINGS**  
16771678 In this section, we provide the main implementation settings for reproducibility in Ta-  
1679 ble 21. Please refer to our source code at [https://anonymous.4open.science/r/](https://anonymous.4open.science/r/RL-LLM-Reasoning-1A30)  
1680 RL-LLM-Reasoning-1A30 for the exact usage of each hyper-parameters and more details.  
16811682 Table 21: Implementation details.  
1683

Module	Element	Detail
System	OS	Ubuntu 22.04.2
	CUDA	11.7
	Python	3.11.4
Double-Dueling DQN	$\gamma$	0.9
	Number of episodes	3000
	Batch Size	64
	Interval of target network updating	50
	Optimizer	Adam
	Learning rate	0.01
	Learning rate decay	0.5 per 1000 episodes
	Replay buffer size	500
	Start epsilon	1
	Min epsilon	0.0
RLoT framework	Epsilon decay	0.9995 per step
	Maximum number of actions	5
	Number of trials for self-consistency	3
	Temperature for LLMs	1.0

1701 **J.2 ILLUSTRATION OF THE FULL PIPELINE**  
17021703 Here we illustrate how every parts in our design work together as an whole pipeline.  
17041705 **1 DATA CONSTRUCTION**  
17061707 We pick out "hard problems", namely those that the LLM wrongly answers when directly prompted,  
1708 for training the navigator. Here are more details: we used Qwen2.5-14B-Instruct as a reference model  
1709 to select "hard problems" from the training sets. The direct prompts are:  
17101710 Prompts:  
1711 {Problem} Please generate the answer for the problem.  
1712 (MATH) Wrap the answer with boxed{{answer}}.  
1713 (GPQA) End the answer with 'The answer is (CHOICE)'.  
1714 (StrategyQA) End the answer with 'YES/NO'.  
17151716 The statistics on benchmarks involved in our training are shown in Table 22.  
17171718 Table 22: Statistics on benchmarks involved in our training.  
1719

Dataset	Total	Hard	Proportion (%)
MATH	7500	1736	23.15
GPQA	448	255	56.92
StrategyQA	1600	341	21.31

1724 By selecting "hard problems", our navigator dedicates to strategies that enhance LLMs on problems  
1725 that they cannot solve directly. For our RL approach, we do not construct (input, expected\_output)  
1726 pairs like supervised learning. The training data consists of individual questions. During training, the  
1727 model is fed with questions, and the outputs are evaluated to generate rewards, which are then fed  
back to train the model. We will detail the key components and entire process of training below.  
1728

1728 2 STATE SELF-EVALUATION (STATE SPACE)  
17291730 We obtain the 7-aspect state vector through a prompted self-evaluation mechanism. We feed the  
1731 problem and the existing reasoning into the LLM with the prompt in Appendix K.1. This design  
1732 compresses the textual reasoning trajectory into a low-dimensional state vector, making the policy  
1733 learning process more efficient and tractable.1734 Here is an input-output example of state self-evaluation, as seen in our case study in Appendix F:  
1735

1736 Input

1737 (Problem) Mark has a garden with flowers. He planted plants of three different  
1738 colors in it. Ten of them are yellow, and there are 80% more of those in  
1739 purple. There are only 25% as many green flowers as there are yellow and  
1740 purple flowers. How many flowers does Mark have in his garden?1741 (Existing reasoning) In this step, we aim at calculating the number of purple  
1742 flowers, which is 80% more than the yellow ones. We calculate by  $10 * (1 + 0.8) = 17$ 

1743 Output

1744 (State vector) {Modelling: True, Calculation: False, ...}

## 1745 3 LOGICAL BLOCKS (ACTION SPACE)

1746

1747 **First**, these logic blocks are inspired by well-established cognitive strategies in human problem  
1748 solving. "Decompose," "Debate," and "Refine" are the core actions, supplemented by "Reason one  
1749 step" and "Terminate" to ensure complete reasoning flows.1750 For instance, "Decompose" reflects the widely studied divide-and-conquer approach, while "Debate"  
1751 aligns with strategies involving comparative evaluation (Wang & Chiew, 2010). Also, they are  
1752 empirically effective. Prior studies (Madaan et al., 2023; Xue et al., 2024) have demonstrated that  
1753 "Decompose" and "Refine" can significantly enhance the performance of LLMs on complex tasks.1754 **Second**, while preserving the high-level cognitive intuition, we have operationalized them into  
1755 specific primitives that the LLM can execute. The key differences are:

1756

1757

- **Formalization with prompts:** Each block corresponds to a prompt template that instructs  
1758 the LLM to perform a specific function (Appendix K.2).
- **Standard Input/Output:** All blocks operate on a unified interface. The input is the original  
1759 problem and the existing reasoning, and the output is a new segment of reasoning. This  
1760 allows flexible composition of logic blocks.

1761 **Third**, here is an input-output example for the 'Debate' block (as Appendix B.1):

1762

1763 Input

1764 (Problem) How many square units are in the region satisfying the inequalities  
1765  $\$y \geq |x|$  and  $\$y \leq -|x| + 3$ ?1766 (Existing reasoning) Let's start by analyzing the given inequalities. The next  
1767 step is to determine the points of intersection... Now we know the points of  
1768 intersection are...

1769

1770 Output

1771 (New reasoning) The most promising plan is: Given the diagonals of the rhombus,  
1772 we can now calculate the area.

1773

1774 The contributions of the logic blocks are discussed in Appendix I.

1775

## 1776 4 PROCESS REWARD

1777 The PRM takes in the problem and existing reasoning and outputs a numerical score evaluating its  
1778 quality.

1779

1780 Here is an input-output example for the PRM:

1781

1782 Input:  
 1783 (Problem) How many square units are in the region satisfying the inequalities  
 1784  $y \geq |x|$  and  $y \leq -|x|+3$ ?  
 1785 (Existing reasoning) Let's start by analyzing the given inequalities. The next  
 1786 step is to determine the points of intersection... Now we know the points of  
 1787 intersection are...

1788 Output:  
 1789 (Reward) 0.765

## 1791 5 DEFINITION OF MDP

1793 Combining the above components, our MDP is defined as follows:  
 1794

- 1795 • **State:** A low-dimensional vector generated from the LLM's self-evaluation of the current  
 1796 reasoning step.
- 1797 • **Action:** One of five logic blocks that guide the next step reasoning.
- 1798 • **Reward:** A score from a PRM that evaluates the quality of the reasoning step after an action  
 1799 is taken.

## 1801 6 TRAINING PROCESS

1803 Follow the MDP framework, we train the navigator with Deep Q-Learning algorithm. The training  
 1804 process is:  
 1805

- 1806 - For training episodes:
  - 1807 - Randomly sample a Problem. Existing reasoning={}.  
  - 1808 - WHILE True:
    - 1809 - Problem + Existing reasoning --[Self-evaluation]--> State vector
    - 1810 - State vector --[MLP navigator]--> Logic block (action)
    - 1811 - Problem + Existing reasoning --[Logic block]--> New reasoning
    - 1812 - Existing reasoning <-- Existing reasoning + New reasoning
    - 1813 - Problem + Existing reasoning --[PRM]--> Reward
    - 1814 - Use the reward to train the MLP navigator
    - 1815 - IF reach the answer: BREAK

## 1816 7 INFERENCE PROCESS

1818 With the trained navigator, the inference process is:  
 1819

- 1820 - Problem. Existing reasoning={}. Answers={}.  
- 1821 - For self-consistency candidates number:
  - 1822 - WHILE True:
    - 1823 - Problem + Existing reasoning --[Self-evaluation]--> State vector
    - 1824 - State vector --[MLP navigator]--> Logic block (action)
    - 1825 - Problem + Existing reasoning --[Logic block]--> New reasoning
    - 1826 - Existing reasoning <-- Existing reasoning + New reasoning
    - 1827 - IF reach the answer: BREAK
  - 1828 - Answers <-- Answers + New answer
  - 1829 - Answers --[Self-consistency]--> Final answer

## 1830 8 SELF-CONSISTENCY EVALUATION

1831 To enhance the robustness of our final answers, we employ a self-consistency mechanism. Since the  
 1832 final answers can be in different formats, we use `sympy.parsing.latex` library to parse the answers and  
 1833 determine if they are mathematically equivalent. We select the final answer using majority voting. If  
 1834 no majority exists, one answer is selected at random.  
 1835

Here is an input-output example for self-consistency:

```
1836 Input:  
1837 (Answers) {0.5} {0.7} {1/2}  
1838  
1839 Output:  
1840 (Final answer) 0.5  
1841  
1842  
1843  
1844  
1845  
1846  
1847  
1848  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889
```

---

1890 **K PROMPTS**
1891

1892 In this section, we provide the prompts for obtaining states and performing actions to ensure repro-
1893 ducibility. In the following prompt blocks, the text enclosed in braces "{}" denotes problem-specific
1894 content, such as intermediate reasoning steps, while the remaining texts serve as fixed templates.
1895

1896 **K.1 PROMPT FOR OBTAINING STATES**
1897

1898 We use the following prompt to extract a state vector from an intermediate reasoning step. This
1899 prompt guides the LLM in systematically evaluating the current step across multiple aspects. During
1900 our experiments, we observed that LLMs effectively identify these aspects and provide detailed scores
1901 for each aspect.

1902 **Listing 1: Prompts for obtaining states**


---

1903 1 {Problem and reasoning steps}
1904 2 Please evaluate the current step from the following aspects.
1905 3 A) Correctness
1906 4 A1: Correctness of modeling:
1907 5 Whether the current step is correctly derived from the origin
1908 problem.
1909 6 A2: Clarity for further reasoning:
1910 7 Whether the current step is clearly presented, without ambiguity
1911 , to support further reasoning.
1912 8 A3: Correctness of calculation:
1913 9 Whether the numerical computation in the current step is
1914 performed correctly.
1915 10 B) Complexity
1916 11 B1: Complexity to reach the final answer:
1917 12 Whether it still requires complex reasoning or calculation to
1918 reach the final answer from the current step.
1919 13 B2: Alternative methods in further reasoning:
1920 14 Whether there exist multiple alternative methods to solve the
1921 problem in the current step.
1922 15 C) Completeness
1923 16 C1: Closeness to the final solution:
1924 17 Whether the current step is close enough to directly reach the
1925 final answer.
1926 18 C2: Completeness within the step:
1927 19 Whether all necessary elements within this specific step are
1928 known from the problem or previous steps.
1929 20 For each aspect, please score 1 for False, 2 for Unsure, and 3 for
1930 True, and score 0 if the current step does not involve this
1931 aspect. Please attach the reason for each score.
1932 21 Use the format 'A1 score=[SCORE] reason=[REASON]'.
1933 22 Only score the current reasoning step here, and DONOT conduct
1934 further reasoning.

---

1935 **K.2 PROMPTS FOR ACTIONS**
1936

1937 **Reason one step:** The prompt below is designed to conduct the action of “Reason one step”. We
1938 emphasize that it should focus on one step at a time, which better controls the output and prevents
1939 mistakes in long reasoning paths.

1940 **Listing 2: Prompts for action “Reason one step”**


---

1941 1 Here is a problem and several reasoning steps.
1942 2 {Problem and previous steps}
1943 3 Please reason exactly ONE more step based on the current step here,
1944 and DONOT reason too many steps at once.

---

1945 **Decompose:** A “Decompose” action consists of the following three prompts. First, we use prompt 3
1946 to break down the current problem into multiple subtasks. Next, prompt 4 is applied sequentially to

1944 each subtask to complete its execution. Since the current subtask may depend on previous ones, the  
 1945 results of earlier subtasks are included (see line 4 of the prompt). Finally, the execution results are  
 1946 summarized using prompt 5, which captures the key steps and outcomes. Only the summarized result  
 1947 is utilized for subsequent reasoning actions.

1948 **1949 Listing 3: Prompts for obtaining subtasks in action “Decompose”**

---

1950 1 Here is a problem and several reasoning steps.  
 1951 2 {Problem and previous steps}  
 1952 3 Please decompose the current task into subtasks, where we can solve  
     the original problem by combining these results of subtasks.  
 1953 4 Only provide subtasks decomposition here, and DONOT conduct specific  
     reasoning or calculation.  
 1954 5 Use the format ‘### Subtask1: subtask1’.

---

1956 **1957 Listing 4: Prompts for executing subtasks in action “Decompose”**

---

1958 1 Here is a problem and several reasoning steps.  
 1959 2 {Problem and reasoning steps before decomposition}  
 1960 3 For the next step, the task is decomposed into subtasks, here are  
     the reasonings in the first few subtasks.  
 1961 4 {Executing results of previous subtasks}  
 1962 5 Please conduct the following Subtask{subtask\_id} to continue the  
     reasoning.  
 1963 6 DONOT conduct a more detailed decomposition for the subtask.

---

1965 **1966 Listing 5: Prompts for summarize subtasks in action “Decompose”**

---

1967 1 Here are a few detailed reasoning subtasks of a problem.  
 1968 2 {Executing results of subtasks}  
 1969 3 Please give a clear and concise summary of these subtasks, keeping  
     the key reasoning and results in each subtask.  
 1970 4 Only provide the summary here, and DONOT conduct more reasoning or  
     calculation.

---

1973 **Debate:** A “Debate” action involves multiple rounds of question-answering. First, the LLM generates  
 1974 different plans for the task using prompt 6. Next, the plans are compared, and the most promising  
 1975 one is selected using prompt 7, mimicking how human experts debate and discuss to reach a solution.  
 1976 Based on the chosen plan, reasoning is advanced by one step through prompt 8. Similar to the  
 1977 “Decompose” action, only the results of the final one-step reasoning (output of prompt 8) are retained  
 1978 for subsequent reasoning processes.

1979 **1980 Listing 6: Prompts for obtaining various plans in action “Debate”**

---

1981 1 Here is a problem and several reasoning steps.  
 1982 2 {Problem and previous reasoning steps}  
 1983 3 Please propose three different alternative plans for solving the  
     problem in the current step.  
 1984 4 Only provide plans here, and DONOT conduct specific reasoning or  
     calculation.  
 1985 5 Use the format ‘### Plan1: plan1’.

---

1987 **1988 Listing 7: Prompts for analysing and comparing plans in action “Debate”**

---

1989 1 Here is a problem and several reasoning steps.  
 1990 2 {Problem and previous reasoning steps}  
 1991 3 Currently, we have several alternative plans for solving the problem  
     in the current step.  
 1992 4 {Generated Plans}  
 1993 5 Please review and compare these plans carefully, and tell which one  
     is most promising for further reasoning. Only compare the plans  
     here, and DONOT conduct further reasoning or calculation.  
 1994 6 Use the format ‘The most promising plan is Plan[INDEX]: [REASON]’,  
     where [INDEX] is an integer index of the plan and [REASON] is a  
     detailed analysis.

---

1998

1999

2000

2001

2002

2003

2004

2005

2006

2007

2008

2009

2010

2011

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

2047

2048

2049

2050

2051

---

**Listing 8: Prompts for executing the plan in action “Debate”**

---

```

1 Here is a problem and several reasoning steps.
2 {Problem and previous reasoning steps}
3 For the next step, we have decided on the most promising plan:
4 {Plan}
5 Please reason **exactly one** more step according to the plan here,
   and DONOT reason too many steps at once.

```

---

**Refine:** We use prompt 9 to perform the “Refine” action, which instructs the LLM to review and improve the reasoning steps for clarity and correctness.

**Listing 9: Prompts for action “Refine”**

---

```

1 Here is a problem and several reasoning steps
2 {Problem and previous reasoning steps}
3 Please check and refine the current thought here, and DONOT conduct
   further reasoning or calculation.

```

---

**Terminate:** The prompt for the “Terminate” action concludes the reasoning process by generating the final answer. The only variation lies in the output format, which is adapted to the specific requirements of each dataset.

**Listing 10: Prompts for action “Terminate”**

---

```

1 Here is a problem and several reasoning steps
2 {Problem and previous reasoning steps}
3
4 ## GSM8K
5 Please generate the answer for the problem. Please end the answer
   with 'The answer is numerical_answer'.
6
7 ## MATH
8 Please generate the answer for the problem. Wrap the answer with \\
   boxed{{answer}}.
9
10 ## MMLU-STEM and GPQA
11 End the answer with 'The answer is (CHOICE)'.
12
13 ## StrategyQA
14 Please generate the answer for the problem. At the end of your
   answer, conclude the answer with 'The answer is yes' or 'The
   answer is no'.

```

---

**K.3 SENSITIVE TO ACTION-PROMPTS PHRASING.**

To evaluate whether the performance of RLoT is sensitive to the specific wording of the action prompts, we conducted a robustness test by rephrasing the natural language instructions for the Decompose, Debate, and Refine actions.

**Decompose:** We rephrase the prompts for the action of Decompose as:

**Listing 11: Prompts for obtaining subtasks in action “Decompose”**

---

```

1 This is the problem and the reasoning progress so far.
2 {Problem and previous steps}
3 Break down the current problem into smaller, manageable sub-problems
.
4 Solving these sub-problems sequentially should lead to the solution
   of the original question.
5 Only provide subtasks decomposition here, and DONOT conduct specific
   reasoning or calculation.
6 Format: '### Subtask1: subtask1'.

```

---

2052

## 2053 Listing 12: Prompts for executing subtasks in action “Decompose”

2054

---

- 1 This is the problem context and the steps taken.
- 2 {Problem and reasoning steps before decomposition}
- 3 We have divided the main task into several parts. Below is the reasoning progress for the initial parts.
- 4 {Executing results of previous subtasks}
- 5 Now, please execute the specific instruction for Subtask{subtask\_id} to advance the solution.
- 6 DONOT conduct a more detailed decomposition for the subtask.

---

2056

2059

2060

2061

2062

## 2063 Listing 13: Prompts for summarize subtasks in action “Decompose”

2064

2065

2066

2067

2068

2069

2070 **Debate:** We rephrase the prompts for the action of Debate as:

2071

2072

## 2073 Listing 14: Prompts for obtaining various plans in action “Debate”

2074

2075

2076

2077

2078

2079

2080

2081

---

- 1 This is a problem and existing reasoning steps.
- 2 {Problem and previous reasoning steps}
- 3 Formulate three distinct strategies to address the problem from the current state.
- 4 Only provide plans here, and DONOT conduct specific reasoning or calculation.
- 5 Use the format ’### Plan1: plan1’.

---

2082

2083

2084

2085

2086

2087

2088

2089

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

## 2080 Listing 15: Prompts for analysing and comparing plans in action “Debate”

---

- 1 This is the problem context and the steps taken.
- 2 {Problem and previous reasoning steps}
- 3 We have generated a set of potential strategies to address the problem at this stage. {Generated Plans}
- 4 Critically evaluate these options and identify the most viable strategy to advance the reasoning.
- 5 Only compare the plans here, and DONOT conduct further reasoning or calculation.
- 6 Use the format ’The most promising plan is Plan[INDEX]: [REASON]’, where [INDEX] is an integer index of the plan and [REASON] is a detailed analysis.

---

2090

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

## 2092 Listing 16: Prompts for executing the plan in action “Debate”

---

- 1 Here is a problem and several reasoning steps.
- 2 {Problem and previous reasoning steps}
- 3 We have selected the optimal strategy to proceed: {Plan}
- 4 Carry out the immediate next logical step following this strategy, and DONOT reason too many steps at once.

---

2099 **Refine:** We rephrase the prompts for the action of Refine as:

2100

2101

2102

2103

2104

2105

## 2101 Listing 17: Prompts for action “Refine”

---

- 1 Here is a problem and several reasoning steps
- 2 {Problem and previous reasoning steps}
- 3 Please examine the current reasoning for potential improvements and revise it accordingly, and DONOT conduct further reasoning or calculation.

---

2106

2107

Table 23: Impact of prompt rephrasing.

LLM	Method	GSM8K	GPQA	StrategyQA	Average
Qwen2.5-Instruct-7B	DirectQA	91.58	31.25	68.85	63.89
	RLoT (original)	<b>92.87</b>	44.64	79.04	72.18
	RLoT (rephrased)	92.49	<b>45.09</b>	<b>79.91</b>	<b>72.42</b>

2112

2113

The results, presented in Table 23, demonstrate that RLoT exhibits high robustness to prompt variations. The model with rephrased prompts achieves an average accuracy very similar to the original one. This consistency indicates that the effectiveness of RLoT stems from the adaptive logical structures constructed by the navigator, rather than overfitting to specific phrasing or reliance on extensive prompt engineering.

2119

2120

2121

2122

2123

2124

2125

2126

2127

2128

2129

2130

2131

2132

2133

2134

2135

2136

2137

2138

2139

2140

2141

2142

2143

2144

2145

2146

2147

2148

2149

2150

2151

2152

2153

2154

2155

2156

2157

2158

2159

2160 **L USE OF LLMs**  
2161

2162 The authors used LLMs to aid or polish paper writing, but all content has been carefully reviewed by  
2163 the author. The authors used LLMs for literature retrieval and discovery, but all related works have  
2164 been carefully reviewed and organized by the author. The research ideation in this work was entirely  
2165 completed by the author and does not involve the use of LLMs.  
2166  
2167  
2168  
2169  
2170  
2171  
2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187  
2188  
2189  
2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202  
2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213