

# 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 INNATE REASONING IS NOT ENOUGH: IN-CONTEXT LEARNING ENHANCES REASONING LARGE LANGUAGE MODELS WITH LESS OVERTHINKING

Anonymous authors

Paper under double-blind review

## ABSTRACT

Recent advances in Large Language Models (LLMs) have introduced Reasoning Large Language Models (RLLMs), which employ extended thinking processes with reflection and self-correction capabilities, demonstrating the effectiveness of test-time scaling. RLLMs exhibit innate Chain-of-Thought (CoT) reasoning capability obtained from training, leading to a natural question: “Is CoT prompting, a popular In-Context Learning (ICL) method for chat LLMs, necessary to enhance the reasoning capability of RLLMs?” In this work, we present the first comprehensive analysis of the impacts of Zero-shot CoT and Few-shot CoT on RLLMs across mathematical reasoning tasks. We examine models ranging from 1.5B to 32B parameters, finding that contrary to concerns, CoT prompting significantly enhances RLLMs’ performance in most scenarios. Our results reveal distinct patterns: large-capacity models show minimal improvement on simple tasks but substantial gains on complex problems, while smaller models exhibit the opposite behavior. Further analysis demonstrates that CoT prompting effectively controls the distribution of the numbers of thinking tokens and reasoning steps, reducing excessive reflections by approximately 90% in some cases. Moreover, attention logits analysis reveals the RLLMs’ overfitting to reflection-related words, which is mitigated by external CoT guidance. Notably, our experiments indicate that for RLLMs, one-shot CoT consistently yields superior performance compared to Few-shot CoT approaches. Our findings provide important insights for optimizing RLLMs’ performance through appropriate prompting strategies.

## 1 INTRODUCTION

Recent advances in Large Language Models (LLMs) have introduced Reasoning Large Language Models (RLLMs) such as OpenAI o1 (2024), DeepSeek-R1 (2025), and Qwen QwQ (2025). These models generate innate Chain-of-Thought (CoT) before answering at inference time, demonstrating the effectiveness of test-time scaling (Muennighoff et al., 2025). Innate CoT is characterized by reflection and self-correction, which can significantly enhance a model’s ability to solve complex reasoning tasks (Kumar et al., 2024).

Over the past three years, as one of the most effective reasoning methods, CoT prompting (Wei et al., 2022) has been widely applied to LLMs and has helped models externally establish chain-like thinking for reasoning problems. This approach has significantly improved performance across various reasoning tasks (Zhang et al., 2022; Wang et al., 2022; Lyu et al., 2023; Ge et al., 2024). However, researchers from DeepSeek Guo et al. (2025) have expressed concern that Few-shot CoT might actually impair the performance of reasoning LLMs, implying that Few-shot CoT may not benefit RLLMs as it does traditional LLMs. This raises a concern: “*Is CoT prompting, a popular In-Context Learning (ICL) method for chat LLMs, necessary to enhance the reasoning capability of RLLMs?*”

In this paper, we present the first comprehensive analysis examining the impact of Zero-shot CoT (Kojima et al., 2022) & Few-shot CoT (Wei et al., 2022) on RLLMs for mathematical problems. Our experiments examine models with parameter sizes ranging from 1.5B to 32B, revealing that Zero-shot CoT & Few-shot CoT significantly enhance the performance of reasoning LLMs on mathematical

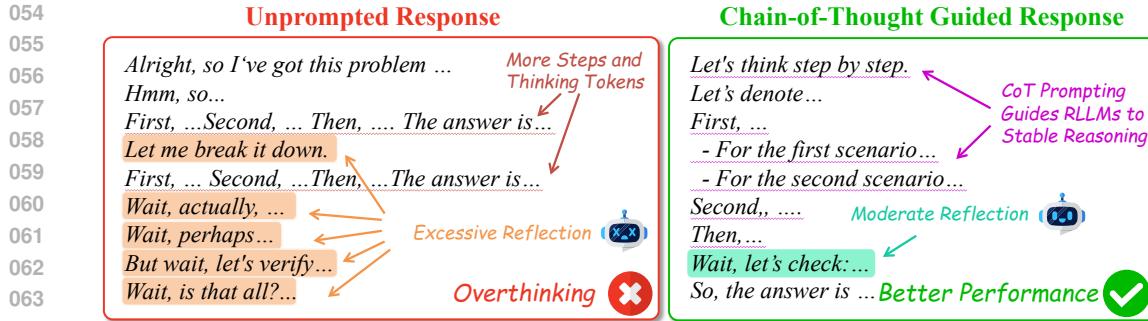


Figure 1: CoT prompting continues to play an important role in reasoning LLMs: (1) improving reasoning performance, (2) controlling the number of thinking tokens, (3) regulating the number of reasoning steps, and (4) mitigating overthinking.

tasks in most cases. Notably, for large-capacity models: the improvement on simple datasets is minimal; however, the improvement on complex datasets is substantial; whereas for small-capacity models, the situation is reversed. Additionally, we find that Zero-shot CoT & Few-shot CoT play a significant role in regulating the numbers of thinking tokens and reasoning steps. To explore why direct prompting generates so many thinking tokens, we have statistically analyzed the reflection frequency of RLLM outputs. Surprisingly, on complex datasets, the average frequency of reflection per question is as high as over 800 for the 1.5B capacity model and over 400 for the 32B capacity model. It is worth noting that CoT prompting can reduce the average number of reflections by about 90% in some cases. To explore the reason for RLLM’s excessive reliance on reflection, we conducted deep analysis of the attention logits (before softmax) and attention scores (after softmax) of R1-8B and found its particular focus on words such as “Wait,” “Double-Check,” and “Ensure.” In contrast, its base model, LLAMA3.1-8B-INSTRUCT, did not exhibit such special attention. Thus, we believe that the reason for this overthinking is the RLLM’s overfitting to reflection and self-correction mechanisms. Through external prompt guidance, such as Zero-shot CoT & Few-shot CoT, the overfitting phenomenon can be alleviated. After observing the significant impact of Zero-shot CoT & Few-shot CoT on the performance of RLLMs, we conducted experiments to test the effect of the number of shots on performance. Unlike the common trend in LLMs where more shots generally lead to better performance, we found that one-shot CoT achieved the best performance across all datasets and models.

Our main contributions are as follows:

- We are the first to conduct a comprehensive analysis examining the impact of Zero-shot CoT & Few-shot CoT on RLLMs for mathematical problems.
- In terms of breadth, our experiments have revealed that Zero-shot CoT & Few-shot CoT plays a crucial role in controlling the distribution of the numbers of thinking tokens and reasoning steps of RLLM, as well as in suppressing overthinking.
- In terms of depth, we conducted a visualization analysis of the attention logits of RLLMs and discovered the overfitting of RLLM to reflection words. Additionally, our findings indicate that one-shot CoT achieved the best performance for RLLMs.

## 2 RELATED WORKS

### 2.1 RESEARCH AND ANALYSIS ON COT PROMPTING

Chain-of-Thought prompting, first introduced by Wei et al. (2022), showed that providing exemplars of intermediate reasoning steps can significantly boost LLMs’ performance on complex tasks. Soon after, Kojima et al. (2022) discovered that even without any demonstrations, simply appending a prompt like “Let’s think step by step” enables strong Zero-shot CoT reasoning. Subsequent efforts focused on automating and refining CoT prompts. Zhang et al. (2022) proposed Auto-CoT, which automatically generates diverse reasoning chains for Few-shot prompts. In parallel, Wang et al. (2022) introduced a self-consistency decoding strategy: by sampling multiple distinct reasoning paths and selecting the most consistent final answer, they achieved striking performance gains in CoT

108 prompting. Beyond new prompting strategies, researchers also analyzed how CoT content affects  
 109 outcomes. Jin et al. (2024) found that longer reasoning sequences, even containing minor mistakes,  
 110 substantially enhance LLM reasoning accuracy, whereas overly concise chains degrade it. Building  
 111 upon this work, Wu et al. (2025) demonstrate that a nuanced relationship exists between CoT length  
 112 and performance, identifying an optimal length that balances decomposition benefits against error  
 113 accumulation based on model capability and task complexity.

114

## 115 2.2 REASONING LARGE LANGUAGE MODELS WITH INNATE COT 116

117 Despite the excellent performance of CoT prompting, inherent limitations in adaptability persist;  
 118 consequently, reasoning LLMs, such as OpenAI’s o1 (2024) have been introduced to generate reasoning  
 119 internally, offering enhanced efficiency and broader generality. Shortly after its introduction,  
 120 the research community responded with projects: Open-O1 (2024). Subsequently, Alibaba launched  
 121 both Marco-O1 (2024) and QWQ-preview (2024). The former integrates search algorithms and  
 122 reflective prompting within a small-scale model to achieve step-by-step problem solving despite  
 123 limited resources, whereas the latter illustrates that medium-scale open-source models, when com-  
 124 bined with reinforcement learning, can approach the reasoning capabilities of larger proprietary  
 125 models. DeepSeek-R1 (2025) represents the apex of this evolutionary trajectory by adopting an  
 126 extreme “large-scale + pure reinforcement learning” route to achieve reasoning performance on par  
 127 with OpenAI’s O1 models. However, researchers from DeepSeek Guo et al. (2025) have expressed  
 128 concern that Few-shot CoT might impair the performance of RLLMs, implying that Few-shot CoT  
 129 may not benefit RLLMs as it does traditional LLMs. The concern sparked our curiosity. Although  
 130 previous research has explored the impact of CoT prompting on LLMs, our work is the first detailed  
 131 study on the influence of CoT prompting on reasoning LLMs.

132

## 133 3 DOES COT PROMPTING STILL MATTER FOR REASONING LLMs? 134

135 In this section, we first highlight our experimental findings, then introduce our experimental setup,  
 136 followed by details of each experiment and data analysis.

137

We begin by highlighting some of the most exciting results from our analysis here:

138

- 139 • In most cases, CoT prompting plays important roles in improving the performance of reasoning  
 140 LLMs. The magnitude of improvement is influenced by the model’s parameter size and the  
 141 difficulty of the dataset.
- 142 • RLLMs suffer from serious overthinking, the average frequency of reflection per question is as  
 143 high as over 800 for the 1.5B capacity model and over 400 for the 32B capacity model on complex  
 144 datasets. CoT prompting can effectively alleviate this issue.
- 145 • Setting the number of shots to 1 provides the maximum performance of RLLMs.

146

### 147 3.1 PRELIMINARY

148

We employ three external CoT prompting methods in our experiments and briefly introduce these  
 149 methods here: (1) **Direct**: only provides the problem description directly. (2) **Zero-shot CoT** (Kojima  
 150 et al., 2022): involves appending a thought inducing phrase “*Let’s think step by step.*” (3) **Few-shot**  
 151 **CoT** (Wei et al., 2022): provides the LLM with a few exemplars, including task descriptions and  
 152 expected outputs, to guide its reasoning. See Appendix A.1 for details.

153

154

### 155 3.2 SETUP

156

157

**Models** Our experimental subjects are open-source reasoning LLMs: DeepSeek’s DEEPSEEK-R1-  
 158 DISTILL-QWEN-1.5, DEEPSEEK-R1-DISTILL-QWEN-7B, DEEPSEEK-R1-DISTILL-LLAMA-8B,  
 159 DEEPSEEK-R1-DISTILL-QWEN-14B, DEEPSEEK-R1-DISTILL-QWEN-32B, abbreviated as R1-  
 160 1.5B, R1-7B, R1-8B, R1-14B, R1-32B respectively (Guo et al., 2025). Additionally, we included  
 161 open-source models from the community: OPENO1-LLAMA-8B-v0.1 (Open Source O1, 2024),  
 MARCO-O1 (Zhao et al., 2024), abbreviated as OPENO1-8B, MARCO-7B. For detailed decoding  
 162 configurations, see Appendix A.2.

162 163 164 165 166 167 168 169 170 171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215	MODEL	PROMPT	GSM8K	ASDIV	SAT_MATH	MATH	AIME24	AMC23
DEEPEEK-R1 SERIES								
R1-1.5B	Direct	5.7 <sub>(-)</sub>	11.7 <sub>(-)</sub>	46.9 <sub>(-)</sub>	14.4 <sub>(-)</sub>	3.3 <sub>(-)</sub>	10.0 <sub>(-)</sub>	
	Few-shot CoT	31.3 <sub>(\uparrow 449.1)</sub>	<b>50.9</b> <sub>(\uparrow 335.0)</sub>	<b>93.8</b> <sub>(\uparrow 100.0)</sub>	<b>55.4</b> <sub>(\uparrow 284.7)</sub>	<b>6.7</b> <sub>(\uparrow 103.0)</sub>	<b>40.0</b> <sub>(\uparrow 300.0)</sub>	
	Zero-shot CoT	<b>32.8</b> <sub>(\uparrow 475.4)</sub>	43.8 <sub>(\uparrow 274.4)</sub>	71.9 <sub>(\uparrow 53.3)</sub>	37.7 <sub>(\uparrow 161.8)</sub>	3.3 <sub>(0.0)</sub>	30.0 <sub>(\uparrow 200.0)</sub>	
R1-7B	Direct	35.7 <sub>(-)</sub>	60.2 <sub>(-)</sub>	81.2 <sub>(-)</sub>	29.4 <sub>(-)</sub>	6.7 <sub>(-)</sub>	17.5 <sub>(-)</sub>	
	Few-shot CoT	<b>81.1</b> <sub>(\uparrow 127.2)</sub>	<b>88.8</b> <sub>(\uparrow 47.5)</sub>	<b>96.9</b> <sub>(\uparrow 19.3)</sub>	<b>67.2</b> <sub>(\uparrow 128.6)</sub>	<b>20.0</b> <sub>(\uparrow 198.5)</sub>	<b>57.5</b> <sub>(\uparrow 228.6)</sub>	
	Zero-shot CoT	69.6 <sub>(\uparrow 95.0)</sub>	70.8 <sub>(\uparrow 17.6)</sub>	78.1 <sub>(\downarrow 3.8)</sub>	65.8 <sub>(\uparrow 123.8)</sub>	6.7 <sub>(0.0)</sub>	42.5 <sub>(\uparrow 142.9)</sub>	
R1-8B	Direct	78.2 <sub>(-)</sub>	84.9 <sub>(-)</sub>	65.6 <sub>(-)</sub>	64.3 <sub>(-)</sub>	<b>36.7</b> <sub>(-)</sub>	<b>62.5</b> <sub>(-)</sub>	
	Few-shot CoT	69.1 <sub>(\downarrow 11.6)</sub>	81.7 <sub>(\downarrow 3.8)</sub>	<b>87.5</b> <sub>(\uparrow 33.4)</sub>	66.7 <sub>(\uparrow 3.7)</sub>	23.3 <sub>(\downarrow 36.5)</sub>	50.0 <sub>(\downarrow 20.0)</sub>	
	Zero-shot CoT	<b>79.4</b> <sub>(\uparrow 1.5)</sub>	<b>85.1</b> <sub>(\uparrow 0.2)</sub>	84.4 <sub>(\uparrow 28.7)</sub>	<b>68.4</b> <sub>(\uparrow 6.4)</sub>	3.3 <sub>(\downarrow 91.0)</sub>	52.5 <sub>(\downarrow 16.0)</sub>	
R1-14B	Direct	82.9 <sub>(-)</sub>	76.3 <sub>(-)</sub>	71.9 <sub>(-)</sub>	35.3 <sub>(-)</sub>	6.7 <sub>(-)</sub>	15.0 <sub>(-)</sub>	
	Few-shot CoT	<b>89.8</b> <sub>(\uparrow 8.3)</sub>	<b>93.7</b> <sub>(\uparrow 22.8)</sub>	87.5 <sub>(\uparrow 21.7)</sub>	<b>72.1</b> <sub>(\uparrow 104.3)</sub>	<b>33.0</b> <sub>(\uparrow 392.5)</sub>	<b>70.0</b> <sub>(\uparrow 366.7)</sub>	
	Zero-shot CoT	82.2 <sub>(\downarrow 0.8)</sub>	82.3 <sub>(\uparrow 7.9)</sub>	<b>90.6</b> <sub>(\uparrow 25.9)</sub>	61.5 <sub>(\uparrow 74.2)</sub>	13.3 <sub>(\uparrow 98.5)</sub>	37.5 <sub>(\uparrow 150.0)</sub>	
R1-32B	Direct	82.7 <sub>(-)</sub>	85.7 <sub>(-)</sub>	81.2 <sub>(-)</sub>	42.3 <sub>(-)</sub>	10.0 <sub>(-)</sub>	20.0 <sub>(-)</sub>	
	Few-shot CoT	83.2 <sub>(\uparrow 0.6)</sub>	<b>92.5</b> <sub>(\uparrow 7.9)</sub>	<b>100.0</b> <sub>(\uparrow 23.2)</sub>	<b>79.0</b> <sub>(\uparrow 86.8)</sub>	<b>43.3</b> <sub>(\uparrow 333.0)</sub>	<b>57.5</b> <sub>(\uparrow 187.5)</sub>	
	Zero-shot CoT	<b>92.0</b> <sub>(\uparrow 11.3)</sub>	90.1 <sub>(\uparrow 5.1)</sub>	81.2 <sub>(0.0)</sub>	75.6 <sub>(\uparrow 78.7)</sub>	13.3 <sub>(\uparrow 33.0)</sub>	55.0 <sub>(\uparrow 175.0)</sub>	
COMMUNITY MODELS								
MARCO-7B	Direct	52.4 <sub>(-)</sub>	55.7 <sub>(-)</sub>	56.2 <sub>(-)</sub>	47.8 <sub>(-)</sub>	0.0 <sub>(-)</sub>	32.5 <sub>(-)</sub>	
	Few-shot CoT	35.8 <sub>(\downarrow 31.7)</sub>	55.7 <sub>(0.0)</sub>	<b>78.1</b> <sub>(\uparrow 39.0)</sub>	<b>57.1</b> <sub>(\uparrow 19.5)</sub>	3.3 <sub>(\uparrow \infty)</sub>	<b>40.0</b> <sub>(\uparrow 23.1)</sub>	
	Zero-shot CoT	<b>55.6</b> <sub>(\uparrow 6.1)</sub>	<b>59.2</b> <sub>(\uparrow 6.3)</sub>	56.2 <sub>(0.0)</sub>	38.2 <sub>(\downarrow 20.1)</sub>	<b>6.7</b> <sub>(\uparrow \infty)</sub>	20.0 <sub>(\downarrow 38.5)</sub>	
OPENO1-8B	Direct	<b>74.9</b> <sub>(-)</sub>	77.0 <sub>(-)</sub>	75.0 <sub>(-)</sub>	36.1 <sub>(-)</sub>	<b>3.3</b> <sub>(-)</sub>	<b>32.5</b> <sub>(-)</sub>	
	Few-shot CoT	71.2 <sub>(\downarrow 4.9)</sub>	<b>79.2</b> <sub>(\uparrow 2.9)</sub>	<b>84.4</b> <sub>(\uparrow 12.5)</sub>	<b>45.4</b> <sub>(\uparrow 25.8)</sub>	<b>3.3</b> <sub>(0.0)</sub>	17.5 <sub>(\downarrow 46.2)</sub>	
	Zero-shot CoT	<b>74.9</b> <sub>(0.0)</sub>	<b>79.2</b> <sub>(\uparrow 2.9)</sub>	78.1 <sub>(\uparrow 4.1)</sub>	36.1 <sub>(0.0)</sub>	<b>3.3</b> <sub>(0.0)</sub>	22.5 <sub>(\downarrow 30.8)</sub>	

Table 1: Accuracy (%) of various RLLMs across multiple datasets under different prompting settings: Direct (baseline), Few-shot CoT, and Zero-shot CoT. For non-baseline methods, performance changes compared to Direct prompting are shown below as percentages (%). For Few-shot CoT, the default number of shots is 5. Bold numbers represent the highest accuracy achieved for each model-dataset combination across the three prompting methods.

**Datasets** We conduct our experiments on six mainstream English mathematical benchmarks, which cover difficulty levels ranging from elementary school to competition level: GSM8K (Cobbe et al., 2021), ASDiv (Miao et al., 2021), SAT\_MATH (Zhong et al., 2023), MATH (Hendrycks et al., 2021), AIME2024 (AI-MO, 2024a), and AMC2023 (AI-MO, 2024b).

**Metrics** We employed four metrics to analyze the experimental results: (1) Accuracy: The ratio of correct samples to the total number of samples; (2) Number of thinking tokens: The outputs of RLLMs comprise thinking and result parts. ‘Thinking tokens’ refers to the token count within the thinking component; (3) Number of reasoning steps: The number of steps contained in the thinking parts. For example, “Firstly,...; Secondly, ...; Finally, ...” contains three steps; and (4) Number of reflections: The number of reflections per instance. For instance, “Wait, ...” constitutes one reflection. See Appendix A.3 for details.

### 3.3 THE IMPACT OF COT PROMPTING ON ACCURACY OF REASONING LLMs

As shown in Table 1, in 72% of cases, Zero-shot CoT and Few-shot CoT prompting continue to have a general impact on improving the accuracy of reasoning LLMs. Notably, CoT prompting demonstrates

216 significant performance enhancements in specific model and dataset combinations. For instance,  
 217 Zero-shot CoT improved R1-1.5B’s performance on the GSM8K dataset by 475.4%, while Few-shot  
 218 CoT enhanced R1-14B’s performance on the AIME24 dataset by 392.5%. For large-capacity models,  
 219 the improvement on simple datasets is minimal; however, the improvement on complex datasets  
 220 is substantial. Conversely, for small-capacity models, this situation is reversed. Taking R1-32B as  
 221 an example of a large-capacity model, it shows improvements of up to 23.1% and as low as 0.6%  
 222 on simple datasets such as GSM8K, ASDiv, and SAT\_MATH. However, on complex datasets like  
 223 MATH, AIME24, and AMC23, R1-32B demonstrated improvements ranging from 33% to 333%. For  
 224 small-capacity models, exemplified by R1-1.5B, improvements on the same simple datasets ranged  
 225 from 53.3% to 475.44%, while on complex datasets, improvements ranged from no enhancement to  
 226 up to 300%.

### 227 3.4 THE DISTRIBUTION OF THINKING TOKENS

228 As illustrated in Figure 2, the distributions of thinking tokens across DeepSeek series models are  
 229 presented. From the analysis of prompting differences, under Direct prompting, the token distribution  
 230 is highly dispersed, with numerous instances where the number of thinking tokens is less than 30.  
 231 Few-shot CoT effectively regulates token distribution, with a substantial concentration of correct  
 232 samples at approximately 100 tokens. This phenomenon can be attributed to LLMs’ tendency to  
 233 emulate the examples provided in Few-shot CoT prompts. The token distribution for Zero-shot CoT  
 234 can be interpreted as an intermediate state between Direct and Few-shot CoT: samples with extremely  
 235 few thinking tokens persist, while simultaneously exhibiting clusters of correct samples concentrated  
 236 within specific ranges of thinking token counts. This indicates that CoT prompting not only influences  
 237 accuracy but also affects the distribution of the number of thinking tokens. Additionally, from  
 238 the perspective of model capacity, under both Direct and Zero-shot CoT conditions, the primary  
 239 distribution of thinking token quantities decreases as model capacity increases.

240 Paradoxically, we observe that beyond a certain threshold in token distribution, accuracy actually  
 241 decreases as the number of output tokens increases. This phenomenon appears to diverge from  
 242 previous research findings (Muennighoff et al., 2025; Jin et al., 2024), which led us to conduct more  
 243 in-depth experiments in the following sections.

### 244 3.5 THE RELATIONSHIP BETWEEN NUMBER OF REASONING STEPS AND ACCURACY

245 As shown in Figure 3, across two complex datasets under two prompting settings, there exists a  
 246 generally proportional relationship between the average numbers of thinking steps and accuracy.  
 247 Notably, the slope of the trend line for Few-shot CoT exceeds that of Direct. Additionally, the  
 248 distribution of step counts in Few-shot CoT tends to be smaller.

249 However, this does not resolve our confusion from the previous section, which prompted us to  
 250 consider: *why does accuracy increase with additional steps while decreasing with additional tokens?*  
 251 Upon analyzing the outputs of reasoning LLMs (see Appendix B for detail), we discovered the  
 252 following phenomenon: reasoning LLMs engage in substantial reflection within individual reasoning  
 253 steps to ensure answer correctness. Due to this reflection behavior, some responses contain few  
 254 reasoning steps yet comprise numerous thinking tokens. This observation reminds us that the numbers  
 255 of reasoning steps and thinking tokens are not proportionally related.

### 256 3.6 EXCESSIVE REFLECTION: THE UNNECESSARY EXHAUSTION OF THINKING TOKENS

257 As shown in Table 2, responses from reasoning LLMs contain numerous reflection statements,  
 258 indicating excessive self-correction and reflection. For example, R1-1.5B generates an average of  
 259 838.2 reflections per instance on the AIME24 dataset; while even the large-capacity model R1-32B  
 260 averages 414.2 reflections per instance under the same conditions. Although AIME24 is a more  
 261 challenging benchmark, generating hundreds of reflections per instance is clearly unreasonable.  
 262 Furthermore, we observed a positive relationship between the average number of reflections per  
 263 instance and dataset difficulty, meaning that as problem complexity increases, reasoning LLMs tend  
 264 to produce even more frequent reflections. This aligns with our intuition: the more challenging the  
 265 problem, the higher the model’s perplexity, leading to increased self-correction and reflection.

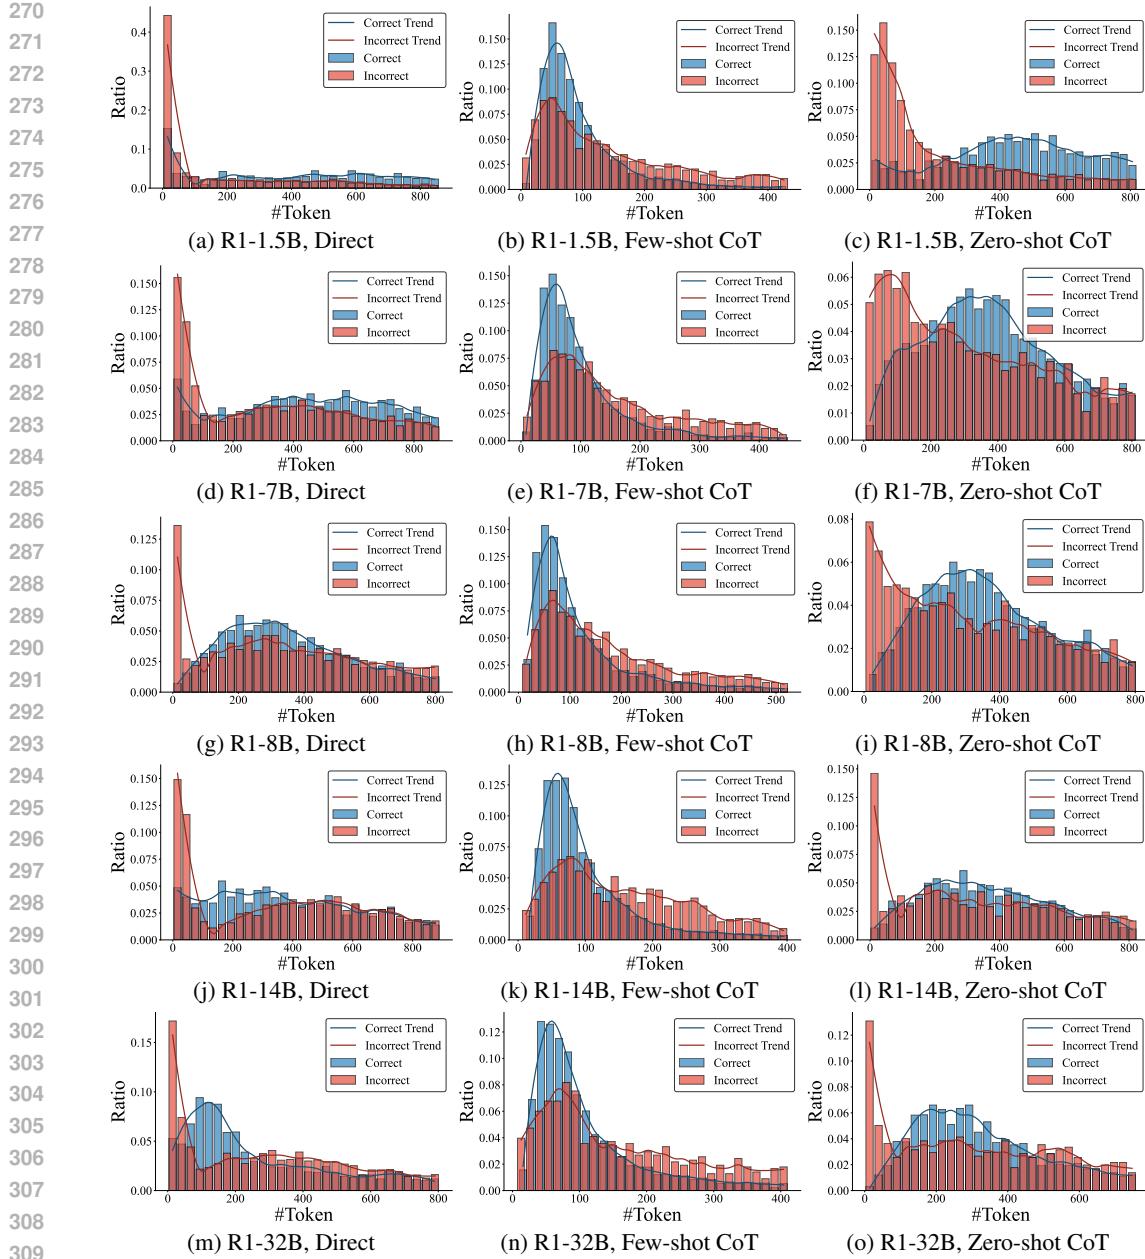


Figure 2: Distributions of thinking tokens across various RLLMs under three prompting methods evaluated on the MATH benchmark. The horizontal axis indicates the number of thinking tokens in the thinking parts (#Token), and the vertical axis represents the corresponding ratio. Histograms labeled “Correct” and “Incorrect” depict the distribution of token counts for correctly and incorrectly solved problems, while the trend lines (“Correct Trend” and “Incorrect Trend”) represent smoothed regression fits of these distributions.

For instance, after implementing Few-shot CoT, R1-32B’s average reflections number per instance decreased from 414.2 to 2.56, while accuracy increased from 10% to 43.3%. It demonstrates that the majority of reflections per instance are redundant and produce unnecessary thinking tokens.

Furthermore, Zero-shot CoT demonstrates a stronger inhibitory effect on excessive self-correction and reflection compared to Few-shot CoT when applied to complex datasets. For example, on the AIME24 dataset, R1-14B averages 63.77 reflections per instance with Few-shot CoT, whereas with Zero-shot CoT, this average decreases to 8.80. Similarly, on the AMC23 dataset, R1-32B averages

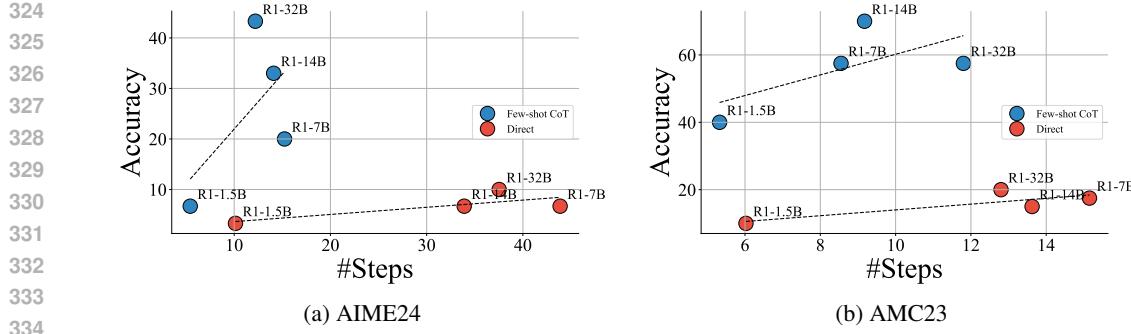


Figure 3: Relationship between accuracy and the average number of reasoning steps for different RLLMs evaluated on AIME24 and AMC23. The horizontal axis represents the average number of reasoning steps (#Steps), and the vertical axis represents accuracy. Dotted lines indicate regression fits illustrating the general correlation trends between average number of reasoning steps and accuracy.

MODEL	AIME24			AMC23			MATH		
	DIRECT	5-SHOT	0-CoT <sup>†</sup>	DIRECT	5-SHOT	0-CoT	DIRECT	5-SHOT	0-CoT <sup>†</sup>
R1-1.5B	838.2	90.3( <sub>489.2</sub> )	18.1( <sub>497.8</sub> )	497.5	170.7( <sub>465.7</sub> )	15.6( <sub>496.9</sub> )	23.5	5.0( <sub>478.8</sub> )	13.9( <sub>40.9</sub> )
R1-7B	435.1	167.8( <sub>461.4</sub> )	6.2( <sub>498.6</sub> )	406.4	123.6( <sub>469.6</sub> )	5.1( <sub>498.7</sub> )	14.2	2.0( <sub>486.0</sub> )	4.7( <sub>466.9</sub> )
R1-8B	59.9	49.2( <sub>418.0</sub> )	10.2( <sub>483.0</sub> )	41.4	60.6( <sub>446.4</sub> )	4.8( <sub>488.4</sub> )	2.0	1.5( <sub>425.6</sub> )	3.4( <sub>474.9</sub> )
R1-14B	529.0	63.8( <sub>487.9</sub> )	8.8( <sub>498.3</sub> )	307.9	52.4( <sub>483.0</sub> )	6.6( <sub>497.9</sub> )	10.5	1.6( <sub>484.6</sub> )	4.4( <sub>458.5</sub> )
R1-32B	414.2	2.6( <sub>499.4</sub> )	7.8( <sub>498.1</sub> )	268.0	101.9( <sub>462.0</sub> )	4.2( <sub>498.4</sub> )	6.9	1.1( <sub>483.5</sub> )	3.0( <sub>457.0</sub> )
QWEN2.5-MATH	0.0	0.0( <sub>0.0</sub> )	0.0( <sub>0.0</sub> )	0.0	0.0( <sub>0.0</sub> )	0.0( <sub>0.0</sub> )	0.1	0.0( <sub>433.3</sub> )	0.0( <sub>483.3</sub> )
LLAMA3.1-8B	0.0	0.0( <sub>0.0</sub> )	0.0( <sub>0.0</sub> )	0.0	0.0( <sub>0.0</sub> )	0.0( <sub>0.0</sub> )	0.2	0.0( <sub>4100.0</sub> )	0.5( <sub>4113.6</sub> )

Table 2: Average number of reflections per instance across different LLMs under three prompting settings on three mathematical datasets. Values shown in gray indicate results under the Direct baseline. For non-baseline methods, relative performance changes are shown below as percentages (%). Models QWEN2.5-MATH-7B and LLAMA3.1-8B-INSTRUCT, abbreviated as QWEN2.5-MATH and LLAMA3.1-8B respectively, are included as comparative baselines since they serve as the base models from which R1-7B and R1-8B are fine-tuned. <sup>†</sup>: 0-CoT refers to Zero-shot CoT.

101.85 reflections per instance with Few-shot CoT, while with Zero-shot CoT, this average is reduced to 4.17. These findings indicate that employing Zero shot CoT is a simpler and token efficient method for suppressing overthinking.

## 4 DEEPER EXPLORATION

### 4.1 OVERTHINKING IS COMMONPLACE

Our previous analysis identified excessive reflections as a significant phenomenon in reasoning LLMs tackling complex mathematical problems. This behavior manifests as numerous reflections that increase token counts without proportionally improving accuracy. We now examine whether this behavior persists in simpler mathematical problems.

As shown in Figure 4, we analyzed the relationship between accuracy and reasoning steps across different model capacities on simpler datasets (GSM8K and ASDiv). The results reveal a consistent pattern: accuracy initially increases with additional reasoning steps but begins to decline after reaching an optimal point (typically 2-3 steps). This inverted U-shaped relationship is particularly pronounced in smaller models. For example, in GSM8K with the R1-1.5B model (Figure 4a), accuracy peaks at 3 steps under Few-shot CoT prompting before dropping significantly at 4 steps. Similarly, in ASDiv, optimal performance occurs at 2 steps before declining. Since these datasets typically require no more than 3 steps to solve, additional steps represent redundant reasoning rather than productive problem-solving. This phenomenon appears universally across model sizes, though its severity varies. Smaller models (R1-1.5B and R1-7B) experience more dramatic performance degradation with

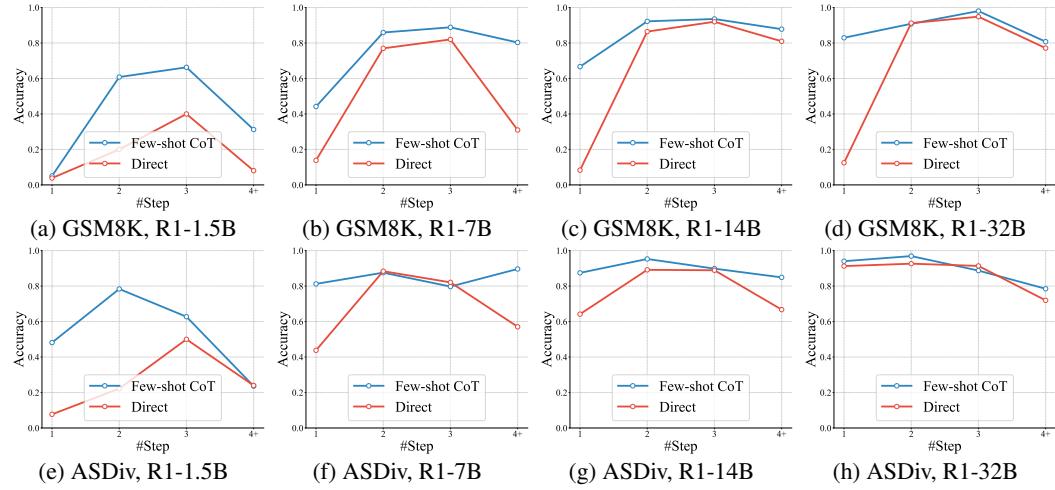


Figure 4: Relationship between the number of reasoning steps (#Step) and accuracy of RLLMs on the GSM8K and ASDiv datasets. The accuracy is averaged across individual reasoning steps provided by the RLLMs. Results show that accuracy initially increases with the number of steps but declines after reaching an optimal point (around 2-3 steps).

#SHOTS	R1-1.5B	R1-7B	R1-14B	R1-32B
0	3.3(–)	6.7(–)	6.7(–)	10.0(–)
1	<b>13.3</b> ( $\uparrow$ 303.0)	<b>23.3</b> ( $\uparrow$ 247.8)	<b>36.7</b> ( $\uparrow$ 447.8)	<b>56.7</b> ( $\uparrow$ 467.0)
2	3.3(0.0)	10.0( $\uparrow$ 49.3)	26.7( $\uparrow$ 298.5)	40.0( $\uparrow$ 300.0)
3	10.0( $\uparrow$ 203.0)	20.0( $\uparrow$ 198.5)	23.3( $\uparrow$ 247.8)	30.0( $\uparrow$ 200.0)
4	6.7( $\uparrow$ 103.0)	<b>23.3</b> ( $\uparrow$ 247.8)	33.3( $\uparrow$ 397.0)	43.3( $\uparrow$ 333.0)
5	6.7( $\uparrow$ 103.0)	20.0( $\uparrow$ 198.5)	33.0( $\uparrow$ 392.5)	43.3( $\uparrow$ 333.0)

Table 3: Accuracy (%) of RLLMs on the AIME24 dataset under different Few-shot CoT settings. “#Shots” indicates the number of Question-Answer pairs as examples provided to the model. The baseline (Direct, without any Question-Answer pair as example) is shaded in grey, with percentages below showing relative performance changes (%) compared to this baseline. Bold numbers represent the highest accuracy achieved for each model.

excessive steps, while larger models (R1-14B and R1-32B) maintain relatively stable performance even with additional steps. These findings confirm that excessive reflection remains prevalent even in simpler datasets, and that adding steps beyond necessity does not improve and often harms accuracy.

#### 4.2 IMPACT OF THE NUMBER OF SHOTS ON RLLM PERFORMANCE

In the preceding sections, our experiments have demonstrated that CoT prompting significantly enhances the performance of reasoning LLMs across most scenarios. However, in our default experimental configuration, we utilized a five-shot setting for Few-shot CoT prompting. This raises a question: What is the optimal number of exemplars for maximizing RLLM performance, and how does performance vary as the number of shots changes?

Table 3 presents the accuracy of various DeepSeek models on the challenging AIME24 dataset under different Few-shot CoT settings (0-5 shots). The results reveal a clear pattern: providing exactly one Question-Answer pair (one-shot) yields optimal or near-optimal performance for most model sizes. For R1-1.5B, R1-14B, and R1-32B, one-shot CoT prompting produces the highest accuracy, with improvements ranging from 303.0% for R1-1.5B to 467.0% for R1-32B compared to the Direct baseline. This finding suggests that minimal exemplification—just a single example—provides sufficient structural guidance for most RLLMs to navigate complex reasoning tasks. Additional examples beyond this point rarely improve performance and often lead to degradation, particularly in

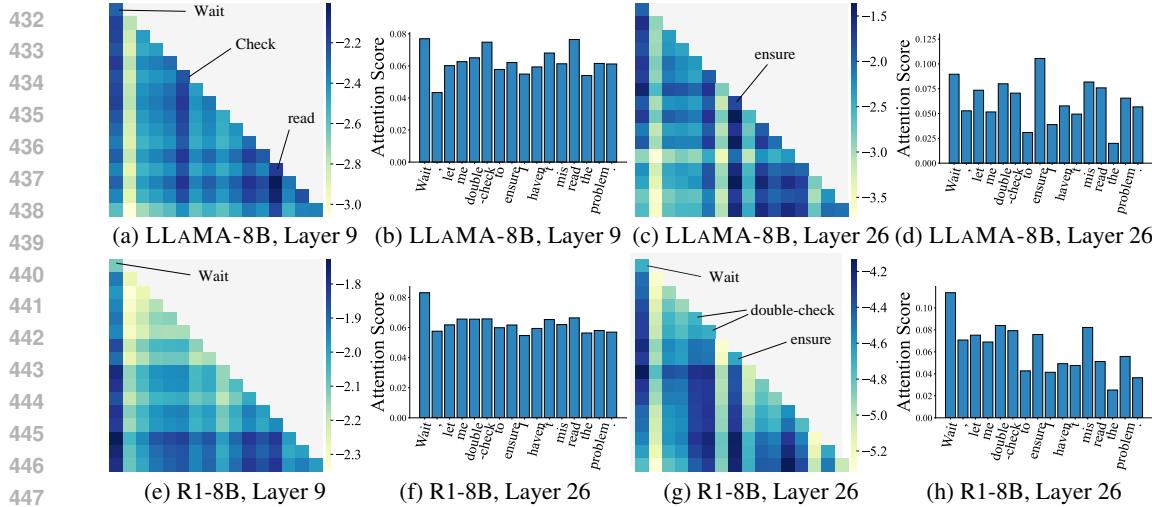


Figure 5: Visualization of attention distribution in LLAMA3.1-8B-INSTRUCT and R1-8B. The heatmaps (left side) show attention logits (before softmax), averaged over all heads per layer, and the corresponding bar graphs (right side) illustrate the softmax-normalized attention scores for the input sequence “Wait, let me double-check to ensure I haven’t misread the problem.” Subfigures (a)-(d) represent the LLAMA3.1-8B-INSTRUCT at layers 9 and 26, while subfigures (e)-(h) depict the R1-8B at the same layers. Here, the attention scores, denoted by  $\alpha$ , are computed as  $\alpha = \mathbb{E}_h[\sigma(\mathbf{A})]$ .

the 2-3 shot range. This pattern indicates that RLLMs may struggle with interference from multiple examples, with a single clear example providing the optimal balance between guidance and flexibility.

### 4.3 ATTENTION-BASED INSIGHTS INTO OVER-REFLECTION

To investigate the mechanistic origins of excessive reflection, we analyzed attention patterns in R1-8B and its base model LLAMA3.1-8B-INSTRUCT. Figure 5 visualizes their attention distribution mechanisms when processing a typical reflection phrase.

The visualizations reveal that R1-8B allocates significantly higher attention to reflection tokens compared to LLAMA3.1-8B-INSTRUCT. At both middle (layer 9) and deep (layer 26) layers, R1-8B exhibits intensified attention logits for tokens such as “Wait” and “double-check” (Figures 5e and 5g). This pattern is further confirmed in the normalized attention scores, where R1-8B allocates substantially higher attention to “Wait” (Figures 5f and 5h) than LLAMA3.1-8B-INSTRUCT (Figures 5b and 5d). These observations suggest that reasoning LLMs have demonstrated an oversensitivity to linguistic markers of verification and reassessment during their training process. This hyperattention to reflection cues likely contributes to the excessive reflection behavior observed in our experiments. The phenomenon appears to stem from an unintended consequence of RLLM training, wherein models overfit to reflection-related keywords.

## 5 CONCLUSION

Our study provides the first comprehensive analysis of CoT prompting for RLLMs, addressing concerns about its potential negative impact. Our experiments across models from 1.5B to 32B parameters on various mathematical tasks demonstrate that both Zero-shot CoT and Few-shot CoT significantly enhance RLLM performance in most scenarios. Large-capacity models showed minimal improvement on simple tasks but substantial gains on complex problems, while smaller models exhibited the opposite pattern. Notably, one-shot prompting consistently outperforms multi-shot approaches. Additionally, CoT effectively regulates thinking token distribution and reasoning steps, reducing excessive reflection. Attention analysis revealed the mechanism behind this phenomenon: RLLMs overfit to reflection-related linguistic tokens, which CoT helps mitigate. Our findings provide crucial insights for optimizing RLLM performance through appropriate prompting strategies, confirming that external CoT remains vital for enhancing mathematical reasoning in RLLMs.

486 ETHICS STATEMENT  
487488 In conducting our research, we prioritize ethical standards to ensure integrity and contribute positively  
489 to the scientific community. We exclusively utilize open-source datasets, ensuring our work builds  
490 upon accessible and transparent resources. Our methods employ widely recognized models with  
491 established reliability within the academic community. We have designed our methodology to prevent  
492 generating harmful or misleading information, safeguarding our findings' integrity.  
493494 REPRODUCIBILITY STATEMENT  
495496 To ensure reproducibility, we utilize publicly available datasets with detailed processing procedures  
497 documented in the appendices. Our methodology is fully specified through pseudocode, mathe-  
498 matical formulations, and comprehensive descriptions. All configurations, hyperparameters, and  
499 evaluation protocols are explicitly documented. Theoretical contributions include complete proofs  
500 with assumptions clearly stated. Code and implementation details will be released upon acceptance.  
501502 REFERENCES  
503504 AI-MO. Ai-mo/aimo-validation-aime, 2024a. URL <https://huggingface.co/datasets/AI-MO/aimo-validation-aime>.  
505506 AI-MO. Ai-mo/aimo-validation-amc, 2024b. URL <https://huggingface.co/datasets/AI-MO/aimo-validation-amc>.  
507509 Baolong Bi, Shaohan Huang, Yiwei Wang, Tianchi Yang, Zihan Zhang, Haizhen Huang, Lingrui  
510 Mei, Junfeng Fang, Zehao Li, Furu Wei, et al. Context-dpo: Aligning language models for  
511 context-faithfulness. *arXiv preprint arXiv:2412.15280*, 2024a.  
512513 Baolong Bi, Shenghua Liu, Lingrui Mei, Yiwei Wang, Pengliang Ji, and Xueqi Cheng. Decod-  
514 ing by contrasting knowledge: Enhancing llms' confidence on edited facts. *arXiv preprint*  
515 *arXiv:2405.11613*, 2024b.  
516517 Baolong Bi, Shenghua Liu, Yiwei Wang, Lingrui Mei, Junfeng Fang, Hongcheng Gao, Shiyu Ni, and  
518 Xueqi Cheng. Is factuality enhancement a free lunch for llms? better factuality can lead to worse  
519 context-faithfulness. *arXiv preprint arXiv:2404.00216*, 2024c.  
520521 Baolong Bi, Shenghua Liu, Yiwei Wang, Yilong Xu, Junfeng Fang, Lingrui Mei, and Xueqi Cheng.  
522 Parameters vs. context: Fine-grained control of knowledge reliance in language models. *arXiv*  
523 *preprint arXiv:2503.15888*, 2025.  
524525 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
526 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
527 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.  
528529 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu,  
530 Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for  $2+3=?$  on the overthinking of  
531 o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024.  
532533 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
534 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
535 Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*,  
536 2021.  
537538 Zenghao Duan, Wenbin Duan, Zhiyi Yin, Yinghan Shen, Shaoling Jing, Jie Zhang, Huawei Shen, and  
539 Xueqi Cheng. Related knowledge perturbation matters: Rethinking multiple pieces of knowledge  
540 editing in same-subject. *arXiv preprint arXiv:2502.06868*, 2025.  
541542 Yuyao Ge, Zhongguo Yang, Lizhe Chen, Yiming Wang, and Chengyang Li. Attack based on data: a  
543 novel perspective to attack sensitive points directly. *Cybersecurity*, 6(1):43, 2023.  
544

540 Yuyao Ge, Shenghua Liu, Baolong Bi, Yiwei Wang, Lingrui Mei, Wenjie Feng, Lizhe Chen, and  
 541 Xueqi Cheng. Can graph descriptive order affect solving graph problems with llms? *arXiv preprint*  
 542 *arXiv:2402.07140*, 2024.

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

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

549 Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and  
 550 Mengnan Du. The impact of reasoning step length on large language models. In Lun-Wei Ku, Andre  
 551 Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 1830–1842, Bangkok, Thailand, August 2024. Association for Computational  
 552 Linguistics. doi: 10.18653/v1/2024.findings-acl.108. URL <https://aclanthology.org/2024.findings-acl.108/>.

553 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
 554 language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:  
 555 22199–22213, 2022.

556 Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli,  
 557 Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via  
 558 reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.

559 Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki,  
 560 and Chris Callison-Burch. Faithful chain-of-thought reasoning. In *The 13th International Joint  
 561 Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter  
 562 of the Association for Computational Linguistics (IJCNLP-AACL 2023)*, 2023.

563 Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, and Xueqi Cheng. Slang: New concept  
 564 comprehension of large language models. *arXiv preprint arXiv:2401.12585*, 2024a.

565 Lingrui Mei, Shenghua Liu, Yiwei Wang, Baolong Bi, Ruibin Yuan, and Xueqi Cheng. Hid-  
 566 deguard: Fine-grained safe generation with specialized representation router. *arXiv preprint*  
 567 *arXiv:2410.02684*, 2024b.

568 Shen-Yun Miao, Chao-Chun Liang, and Keh-Yih Su. A diverse corpus for evaluating and developing  
 569 english math word problem solvers. *arXiv preprint arXiv:2106.15772*, 2021.

570 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke  
 571 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time  
 572 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

573 Open Source O1. Open O1: A Model Matching Proprietary Power with Open-Source Innovation.  
 574 <https://github.com/Open-Source-O1/Open-O1>, 2024.

575 OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms>, 2024. Accessed: September 24, 2025.

576 Qwen. Qwq-32b: The power of scaling rl, March 2025. URL <https://qwenlm.github.io/blog/qwq-32b/>.

577 Mingjie Sun, Xinlei Chen, J. Zico Kolter, and Zhuang Liu. Massive activations in large language  
 578 models. *arXiv preprint arXiv:2402.17762*, 2024.

579 Qwen Team. Qwq: Reflect deeply on the boundaries of the unknown, November 2024. URL  
 580 <https://qwenlm.github.io/blog/qwq-32b-preview/>.

581 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-  
 582 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.  
 583 *arXiv preprint arXiv:2203.11171*, 2022.

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

597

598 Yuyang Wu, Yifei Wang, Tianqi Du, Stefanie Jegelka, and Yisen Wang. When more is less: Under-  
599 standing chain-of-thought length in llms. *arXiv preprint arXiv:2502.07266*, 2025.

600 Wanli Yang, Fei Sun, Xinyu Ma, Xun Liu, Dawei Yin, and Xueqi Cheng. The butterfly effect of model  
601 editing: Few edits can trigger large language models collapse. *arXiv preprint arXiv:2402.09656*,  
602 2024.

603

604 Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. Automatic chain of thought prompting in  
605 large language models. *arXiv preprint arXiv:2210.03493*, 2022.

606

607 Yu Zhao, Huifeng Yin, Bo Zeng, Hao Wang, Tianqi Shi, Chenyang Lyu, Longyue Wang, Weihua  
608 Luo, and Kaifu Zhang. Marco-01: Towards open reasoning models for open-ended solutions, 2024.  
609 URL <https://arxiv.org/abs/2411.14405>.

610 Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu  
611 Chen, and Nan Duan. Agieval: A human-centric benchmark for evaluating foundation models.  
612 *arXiv preprint arXiv:2304.06364*, 2023.

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

Prompt Style	Prompt Template
Direct	Question: ;Question;\nAnswer:
Zero-shot CoT	Question: ;Question;\nAnswer: Let's think step by step.
Few-shot CoT	Question: ;Example question;\nAnswer: ;Example Answer;\n... (more Few-shot CoT examples)\nQuestion: ;Question;\nAnswer:

Table 4: Prompt styles and their corresponding templates for DeepSeek series models distilled from the Qwen family.

## A DETAIL OF EXPERIMENT

### A.1 PROMPT

Table 4 presents detailed prompt templates for DeepSeek series models distilled from the Qwen family (e.g., R1-1.5B, R1-7B). When adapting to a LLaMA-based model (e.g., R1-8B, LLAMA3.1-8B-INSTRUCT), the template undergoes a replacement of “Question:” with “User:” and “Answer:” with “Assistant:”. For MARCO-7B, the template undergoes a replacement of “Question:” with “User:” and “Answer:” with “Content:”.

Below we provide examples of the CoT prompting templates used in our experiments. These examples demonstrate how we formatted different types of mathematical problems and their corresponding CoT solutions across various datasets.

For simpler datasets like GSM8K and ASDiv, we selected question-answer pairs that demonstrate basic arithmetic reasoning, as shown in the following example:

#### Example from GSM8K

Question: There are 15 trees in the grove. Grove workers will plant trees in the grove today. After they are done, there will be 21 trees. How many trees did the grove workers plant today?

Answer: There are 15 trees originally. Then there were 21 trees after some more were planted. So there must have been  $21 - 15 = 6$ . The answer is 6.

For more complex datasets like AMC23, AIME24 and MATH, we utilized examples that demonstrate advanced mathematical reasoning with explicit step-by-step solutions:

#### Example from MATH

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

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

$$\frac{\frac{1}{3}(1 - (\frac{2}{3})^5)}{1 - \frac{2}{3}} = \frac{211}{243}.$$

For multiple choice datasets like SAT\_MATH, we included multiple-choice format examples:

#### Example from SAT\_MATH

Question: If  $\frac{x-1}{3} = k$  and  $k = 3$ , what is the value of  $x$ ? Answer Choices: (A) 2 (B) 4 (C) 9 (D) 10

Answer: If  $k = 3$ , then  $x - 1 = 3 \times 3$ , therefore,  $x - 1 = 9$  and  $x = 10$ . The answer is D.

702 A.2 IMPLEMENTATION DETAILS  
703704 We deploy the open-source LLMs for our experiments on a  $4 \times$  NVIDIA A800 server. The decoding  
705 temperature was set to zero (Greedy decoding). We set the number of maximum new tokens according  
706 to the level of datasets. Specifically, for simple datasets (GSM8K, ASDiv, SAT\_MATH), the number  
707 of maximum token per call is set to 2048; for complex datasets (AIME24, AMC23), it is set to 32768.  
708709 A.3 METRICS  
710

## 711 Accuracy

712 
$$\text{Accuracy} = \frac{\#\text{correct answers}}{\#\text{total questions}} \quad (1)$$
  
713

714 where # represents the number of instances.  
715716  
717 **Number of Thinking Tokens** For OPENO1-8B, the thinking part is wrapped in 'Thoughts' tags.  
718 For others, the content before the last final answer keywords is defined as thinking parts. The answer  
719 keywords contain: "the answer is", "The answer is", "Final Answer", "final answer is", "\*\*Final  
720 Answer", "\*\*Conclusion:\*\*", "\*\*Answer:\*\*".  
721722 **Number of Reasoning steps** We employ LLAMA3.1-8B-INSTRUCT to analyze the thinking part  
723 of RLLMs' responses, specifically to quantify the number of reasoning steps. The prompt templates  
724 are utilized as follows:  
725726 **Prompt Template for Answering The Number of Reasoning Steps**727 Analyze the following mathematical solution and count how many distinct thinking steps are used.  
728 A step is defined as a logical unit where a specific calculation or deduction is made.  
729 Equations that are directly derived from previous ones count as the same step if they're part of the same  
730 logical operation.  
731 Here's a solution example:732 To find the total meters James runs in a week, we need to break down the problem step by step. First,  
733 determine how many sprints James does each week. He runs 3 sprints 3 times a week, so the total number of  
734 sprints is 3 multiplied by 3, which equals 9 sprints.  
735 Next, calculate the total distance by multiplying the number of sprints by the distance of each sprint. Each  
736 sprint is 60 meters, so 9 sprints multiplied by 60 meters per sprint equals 540 meters.  
737 Therefore, James runs a total of 540 meters each week.738 **Solution:**

739 To determine the total number of meters James runs in a week, follow these steps:

740 1. **Calculate the total number of sprints per week:**

741 James runs 3 sprints each day and does this 3 times a week.

742 
$$\text{Total sprints per week} = 3 \text{ sprints/day} \times 3 \text{ days} = 9 \text{ sprints}$$

743 2. **Calculate the total distance run:**

744 Each sprint is 60 meters. Multiply the total number of sprints by the distance of each sprint.

745 
$$\text{Total distance} = 9 \text{ sprints} \times 60 \text{ meters/sprint} = 540 \text{ meters}$$

746 **Final Answer:**

747 540 meters

748 For this example solution, the answer would be:

749 `>{"num_steps": 2}`

750 Because there are 2 distinct thinking steps:

751 1. Calculating the total number of sprints per week 2. Calculating the total distance run

752 Now analyze the following solution:

753 Solution to analyze:

754 *RLLMs' response content*

755 Provide only a JSON object with the following format:

`>{"num_steps": number of distinct thinking steps}`

756 **Number of reflection** We quantify the frequency of reflective keywords within the thinking parts  
 757 of RLLMs' responses.  
 758

759 The reflective keywords contain: "wait", "Let me think", "I need to reconsider", "Let me recalculate",  
 760 "Let me double-check", "Let me verify", "Let me correct", "Let me revise", "Let me rethink", "Let me  
 761 try again", "Hold on", "I see a problem", "Let me backtrack", "Let me restart", "I need to rethink",  
 762 "Let me approach this differently", "Let me redo", "No, that's not right", "Let's break down the  
 763 reasoning".

764 **B CASE STUDY**  
 765

766 To further illustrate our findings, we present a detailed case study comparing the reasoning process  
 767 of the R1-32B model across different prompting methods. We selected a problem, as shown below,  
 768 from the AIME24 dataset that was solved correctly under both Direct and Few-shot CoT settings.  
 769

770 **Pure Question from AIME24**

771 Alice chooses a set  $A$  of positive integers. Then Bob lists all finite nonempty sets  $B$  of positive integers with  
 772 the property that the maximum element of  $B$  belongs to  $A$ . Bob's list has 2024 sets. Find the sum of the  
 773 elements of  $A$ .

774 Under direct prompting, R1-32B generates 1088 thinking tokens across 15 reasoning steps, exhibiting  
 775 extensive self-reflection patterns. The model frequently interrupts its own reasoning with phrases  
 776 like "Wait" and "Let me check" (highlighted in red), demonstrating the excessive reflection tendency  
 777 identified in our experimental analysis. Despite reaching the correct answer (55) midway through its  
 778 reasoning, the model continues to question its conclusion, alternating between accepting and rejecting  
 779 its own results before finally committing to the correct answer.

780 **Answer of R1-32B of Problem Described Directly**

781 Okay, so... Hmm, so... Let's break this down... Well, for... So, the number... Since... Therefore...  
 782 Sum... We need... *Wait*, but 2024 is equal... So, if we can express...  
 783 But 2024 is equal to... Hmm, but 24 is  $16 + 8$ , which is... So,  $2024 = \dots$   
 784 *But wait*, that would mean... But in terms of binary... *Let me check*:...  
 785  $2^10$  is 1024,  $2^11$  is 2048. So... 24 is  $16 + 8$ , so 2024 is...  
 786 But in binary, 2024 is... *Let's verify*:  
 787  $11111101000$  is:  $1 * 2^{10} = \dots 1 * 2^9 = \dots 1 * 2^8 = \dots 1 * 2^7 = \dots 1 * 2^6 = \dots 1 * 2^5 = \dots 0 * 2^4 = \dots 1 * 2^3 =$   
 788  $\dots 0 * 2^2 = \dots 0 * 2^1 = \dots 0 * 2^0 = \dots$   
 789 Adding these up... Yes, that's correct.  
 790 So, 2024 in binary is... which corresponds to... So, the exponents are...  
 791 But in our earlier equation... So, each term... Therefore, the exponents... So,  $a - 1 = \dots$  Therefore ...  
 792 Therefore, the sum... *Let's compute that*:  
 793  $4 + 6 = \dots 10 + 7 = \dots 17 + 8 = \dots 25 + 9 = \dots 34 + 10 = \dots 44 + 11 = \dots$  So, the sum is 55.  
 794 *Wait*, but the initial answer was 2024, which is wrong because... So, the initial thought was incorrect.  
 795 Therefore, *the correct answer should be 2024*.  
 796 *But let me double-check*.  
 797 We have... Expressed in binary, 2024 is... which has 1s at positions... So, the exponents... so  $a = \dots$   
 798 Compute step by step:  $4 + 6 = \dots 10 + 7 = \dots 17 + 8 = \dots 25 + 9 = \dots 34 + 10 = \dots 44 + 11 = 55$ .  
 799 Yes, that's correct. So, the sum is 55.  
 800 Therefore, *the initial answer of 2024 was incorrect. The correct answer is 55*.  
 801 *But wait, let me think again*. Is there another way...  
 802 The problem says:... So, for each  $a$  in  $A$ ... Therefore, the total number... So, we need... So, 2024 is...  
 803 Hence, *the answer is 55*.  
 804 *But wait*, the initial answer was 2024, which is... So, the user...  
 805 Therefore, *the correct answer is 55*.  
 806 *But let me check once more*.  
 807 Suppose  $A = \dots$  Then, for each  $a$  in  $A$ ... So, for  $a=4$ : ...  $a=6$ : ...  $a=7$ : ...  $a=8$ : ...  $a=9$ : ...  $a=10$ : ...  $a=11$ : ...  
 808 Now, let's sum these up:  $8 + 32 = \dots 40 + 64 = \dots 104 + 128 = \dots 232 + 256 = \dots 488 + 512 = \dots$   
 809 Yes, that's correct. So, the total number of sets is indeed 2024, and the sum of  $A$  is 55.  
 810 Therefore, *the answer is 55*. 55

811 In contrast, when prompted with Few-shot CoT, the same model produces only 420 thinking tokens  
 812 across 12 reasoning steps—a 61.4% reduction in token consumption. The reasoning process displays  
 813 markedly higher information density per step, with significantly fewer reflective interruptions. While

810 the model still engages in some verification (as seen in the “Wait, let’s compute” statement), the  
 811 overall reasoning follows a more direct path toward the solution with minimal redundancy.  
 812

813 **Answer of R1-32B of Problem Described with Few-shot CoT Prompting**

814 Let’s think... Let’s denote... Each element... Therefore, the total... Given that...  $\text{Sum}_{a \in A} 2^{a-1} = 2024$ .  
 815 Let’s factor 2024:  $2024 \div 2 = \dots 1012 \div 2 = \dots 506 \div 2 = \dots 253 \div 11 = \dots$

816 So,  $2024 = 2^3 \times 11 \times 23$ .

817 We need... But 2024 in binary...

818 *Wait, let’s compute 2024 in binary:*

819  $2024 \div 2 = \dots 1012 \div 2 = \dots 506 \div 2 = \dots 253 \div 2 = \dots 126 \div 2 = \dots 63 \div 2 = \dots 31 \div 2 = \dots 15 \div 2 = \dots 7 \div 2 = \dots 3 \div 2 = \dots 1 \div 2 = \dots$

820 So writing... 11111101000.

821 So  $2024 = 1024 + 512 + 256 + 128 + 64 + 32 + 8$ .

822 *Wait, let’s check:*

823  $1024 + 512 = \dots 1536 + 256 = \dots 1792 + 128 = \dots 1920 + 64 = \dots 1984 + 32 = \dots 2016 + 8 = 2024$ .

824 Yes, correct.

825 So  $2024 = 2^{10} + 2^9 + 2^8 + 2^7 + 2^6 + 2^5 + 2^3$ .

826 But in our problem... So each term... Therefore, the exponents... Thus,  $a-1 = \dots$  Therefore,  $a = \dots$

827 So the set A is 4,6,7,8,9,10,11.

828 Therefore, the sum... Compute this:

829  $4 + 6 = \dots 10 + 7 = \dots 17 + 8 = \dots 25 + 9 = \dots 34 + 10 = \dots 44 + 11 = 55$ .

830 So the sum is 55.

831 *The answer is 55*

832 This case study provides concrete evidence of how Few-shot CoT prompting effectively regulates  
 833 reasoning behavior in RLLMs by reducing excessive reflections while maintaining or improving  
 834 solution accuracy. The higher information density in the Few-shot example demonstrates that while  
 835 the raw number of reasoning steps is only slightly reduced, the efficiency of those steps is substantially  
 836 improved through the elimination of redundant reflections. This supports our broader finding that  
 837 appropriate prompting strategies can mitigate the reflection overfitting observed in RLLMs, leading  
 838 to more streamlined reasoning without sacrificing performance.

839 **C ADDITIONAL RESULTS**

840 In addition to the MATH dataset analyzed in the main text, we further examined the distribution of  
 841 thinking tokens on the relatively simpler ASDIV and GSM8K datasets. Figures 6 and 7 present  
 842 the corresponding histograms of the number of thinking tokens under the three prompting methods:  
 843 Direct, Zero-shot CoT, and Few-shot CoT.

844 Overall, we observe trends that are consistent with those identified on the more complex MATH  
 845 dataset. First, Few-shot CoT generally yields a more concentrated distribution of thinking tokens,  
 846 mirroring the effect of example-based guidance seen in more challenging tasks. Meanwhile, Direct  
 847 prompting tends to produce outputs that vary more widely in the number of thinking tokens, with a  
 848 notable fraction of responses exhibiting very short or very long thinking parts. Interestingly, Zero-shot  
 849 CoT prompts again lie between these two extremes, indicating that a brief, generic instruction to  
 850 reason step by step partially constrains the model’s thinking process but does not standardize it as  
 851 strongly as providing explicit exemplars.

852 Nevertheless, compared to MATH, the distributions on both ASDIV and GSM8K show that the  
 853 majority of questions require fewer thinking tokens overall. This result aligns with the fact that  
 854 these two datasets are simpler than MATH, which naturally leads to shorter solution paths and fewer  
 855 opportunities for extensive reflections or self-corrections. We also note that, for larger-capacity  
 856 models, the differences in thinking token distributions among correct and incorrect solutions are  
 857 somewhat less pronounced than those observed in the MATH experiments, suggesting that complex  
 858 tasks accentuate the benefits and nuances of prompting more sharply.

859 Despite these dataset-specific distinctions, the overarching pattern remains consistent: CoT prompting  
 860 not only enhances the accuracy of reasoning LLMs but also regulates their reasoning length. In  
 861 particular, the inclusion of even a short chain-of-thought instruction reduces the propensity for  
 862 excessive self-reflection and focuses the models on more concise, goal-oriented reasoning steps.

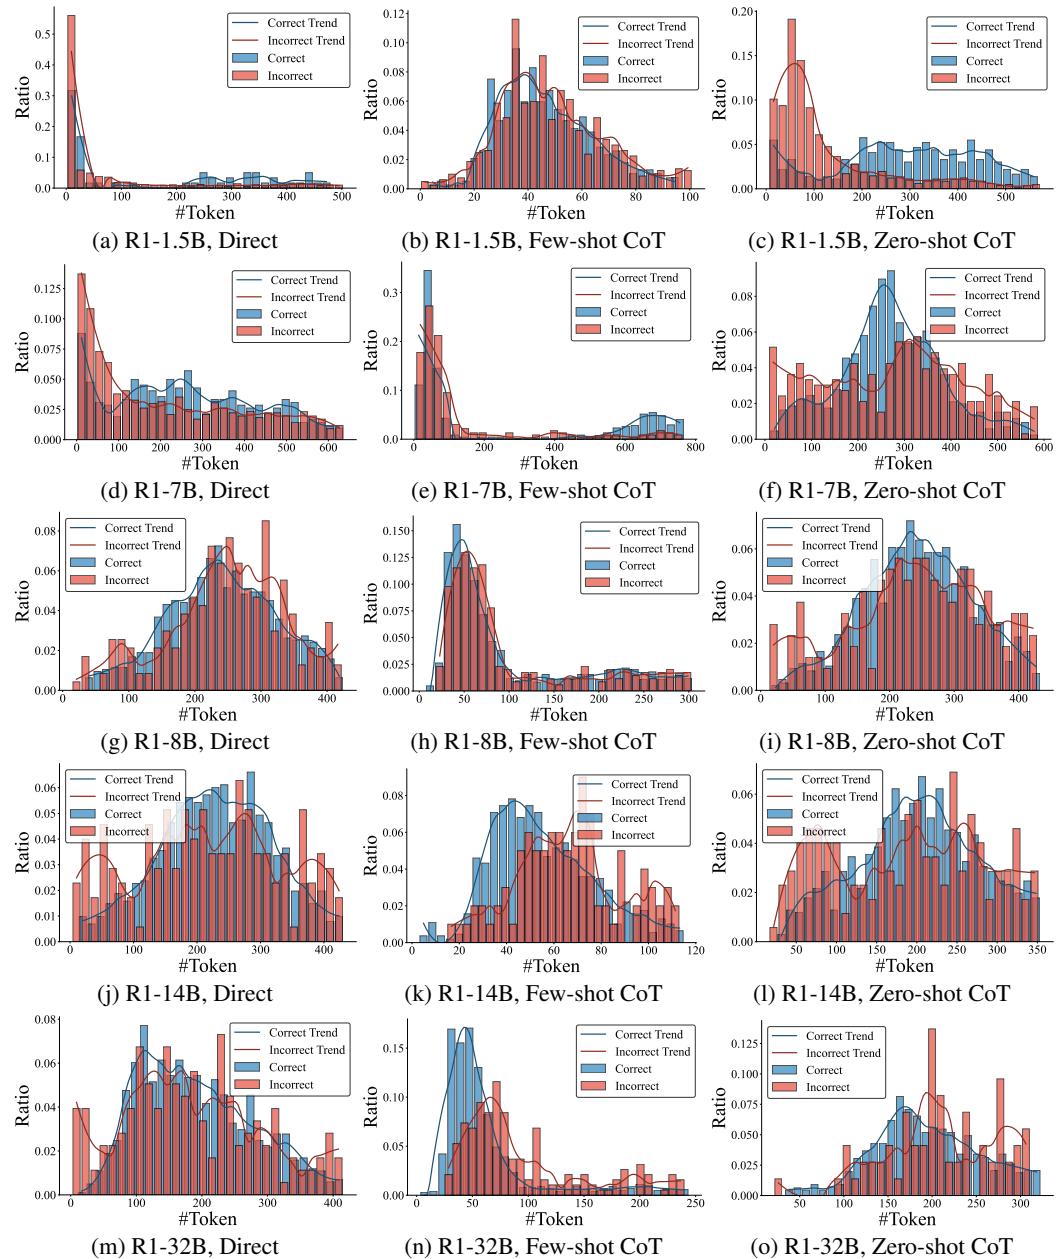


Figure 6: Distributions of thinking tokens across various RLLMs under three prompting methods evaluated on the GSM8K benchmark. The horizontal axis indicates the number of thinking tokens in the thinking parts (#Token), and the vertical axis represents the corresponding ratio. Histograms labeled “Correct” and “Incorrect” depict the distribution of token counts for correctly and incorrectly solved problems, respectively, while the trend lines (“Correct Trend” and “Incorrect Trend”) represent smoothed regression fits of these distributions.

## D LARGE LANGUAGE MODEL USAGE

We employed Claude Sonnet 4 as a grammar expert to assist with proofreading this manuscript. Specifically, Claude Sonnet 4 was used solely to identify and correct linguistic issues including verb tense inconsistencies, grammatical errors, punctuation mistakes, and subordinate clause structures. The LLM’s role was strictly limited to language polishing without any contribution to the research content, methodology, or scientific conclusions.

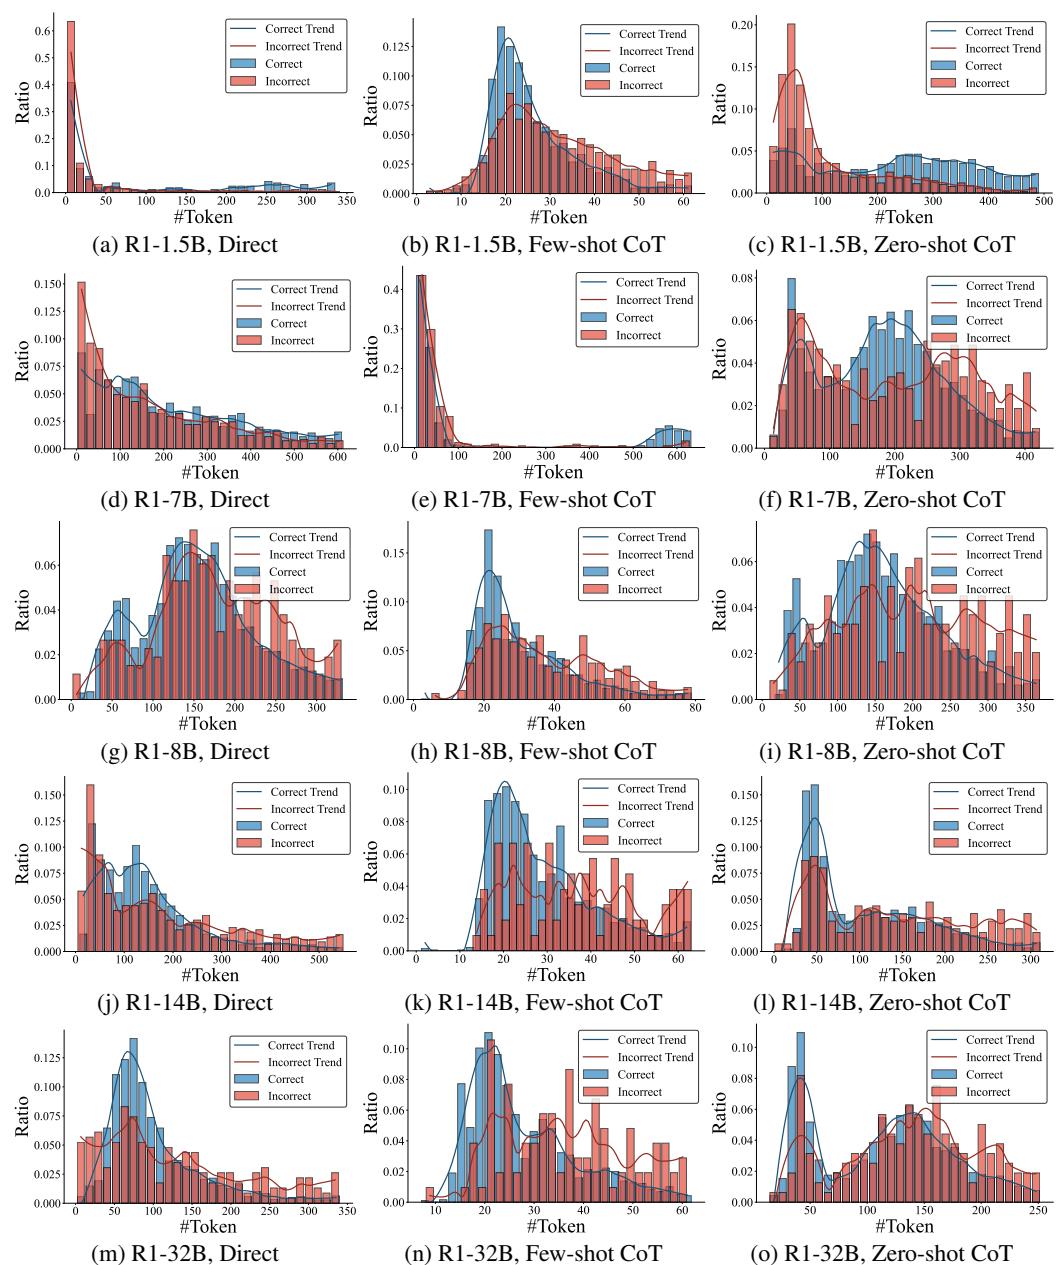


Figure 7: Distributions of thinking tokens across various RLLMs under three prompting methods evaluated on the ASDIV benchmark. The horizontal axis indicates the number of thinking tokens in the thinking parts (#Token), and the vertical axis represents the corresponding ratio. Histograms labeled “Correct” and “Incorrect” depict the distribution of token counts for correctly and incorrectly solved problems, respectively, while the trend lines (“Correct Trend” and “Incorrect Trend”) represent smoothed regression fits of these distributions.