

000 FIRST TRY MATTERS: REVISITING THE ROLE OF RE- 001 002 FLECTION IN REASONING MODELS 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Large language models have recently demonstrated significant gains in reasoning
011 ability, often attributed to their capacity to generate longer chains of thought and
012 engage in reflective reasoning. However, the contribution of reflections to perfor-
013 mance improvement remains unclear. In this paper, we systematically analyze the
014 rollouts of eight reasoning models on five mathematical datasets. We focus on
015 reflective behaviours where the model has already produced an answer but con-
016 tinues reflecting before finalizing its output. Our analysis reveals that reflections
017 are predominantly confirmatory and rarely alter the model’s initial answer, a pat-
018 ttern consistent across models and datasets. To understand the role of reflections in
019 training, we construct supervised fine-tuning (SFT) datasets with varying amounts
020 of reflection steps. We observe that training models on rollouts with more reflec-
021 tion steps primarily enhances first-answer correctness rather than the ability to
022 correct initially wrong answers through reflections. This motivates us to propose
023 a question-aware early-stopping method that enhances inference-time token effi-
024 ciency by stopping the reasoning process once a few plausible candidate answers
025 are generated, thereby reducing unnecessary reflection steps. Motivated by this,
026 we further propose to dynamically truncate the reflections after a candidate answer
027 has appeared during generation, which reduces reasoning tokens by 24.5% across
028 five mathematical datasets, within a 2.9% drop in accuracy.¹

029 1 INTRODUCTION

030 Large language models (LLMs) have made remarkable progress in reasoning abilities, achieving
031 strong performance across domains such as mathematics, logic, and code synthesis (Cobbe et al.,
032 2021; Chen et al., 2021). This leap is largely attributable to the development of Chain-of-Thought
033 (CoT) reasoning pattern (Nye et al., 2021; Wei et al., 2022), which guides the model to break
034 down complex problems into a series of intermediate steps. Recent breakthroughs such as Open-
035 AI’s o1 (OpenAI et al., 2024) and DeepSeek-R1 (DeepSeek-AI et al., 2025) have brought LLMs
036 to the next paradigm, known as reasoning models (Ke et al., 2025; Zhang et al., 2025). Unlike
037 traditional CoT reasoning, which follows a single linear thought process, reasoning models trained
038 with Reinforcement Learning with Verifiable Rewards (RLVR) are believed to possess the ability to
039 internally reflect on their reasoning steps, detect potential errors, and adaptively adjust the reasoning
040 trajectories (Luo et al., 2025; Liu et al., 2025c; Yu et al., 2025; MiniMax et al., 2025; Li et al., 2025)

041 Within this paradigm, a consistent correlation is observed between the length of a model’s generated
042 response and its reasoning accuracy (Muennighoff et al., 2025). Models that generate more extensive
043 CoTs tend to exhibit higher accuracy, suggesting that longer reasoning processes are more beneficial.
044 One commonly observed pattern in these long reasoning rollouts is the presence of “reflections”,
045 where models elaborate or re-examine a solution after deriving a candidate answer. Intuitively, such
046 reflections are assumed to be productive, much like human problem-solving, where self-reflections
047 lead to correction or an “aha moment” (Chen et al., 2025b).

048 Despite their intuitive appeal, prior studies report mixed findings on the effects of reflective behav-
049 iors. Some emphasize the sophisticated internal mechanisms of reflection and their role in preventing
050 reasoning collapse (Yang et al., 2025c), while others argue that self-reflection patterns are often su-
051 perficial and do not improve outcomes (Liu et al., 2025c). Crucially, these studies provide limited

052
053 ¹The code and dataset of this work will be open-sourced upon acceptance.

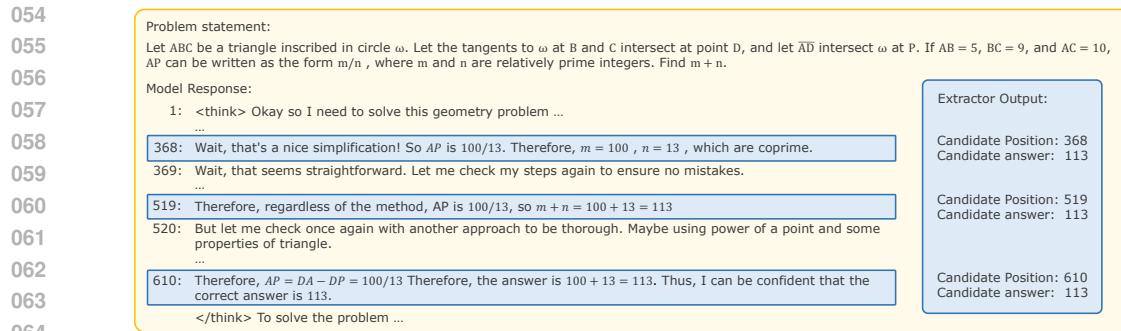


Figure 1: Illustration of a long CoT and the extraction result of candidate answers.

quantitative analysis on the reflective behavior of reasoning models, leaving unresolved whether reflections genuinely help models correct errors or merely confirm earlier conclusions.

To address this open question, we perform a systematic, large-scale quantitative study of reflection patterns in eight reasoning models across five mathematical benchmarks of varying difficulty. To extract these patterns, we design an LLM-based extractor that locates positions in the rollouts where candidate answers are produced. Since a rollout often contains multiple candidate answers, we define the portion of the rollout that follows the first candidate as reflections. This setup allows us to disentangle forward reasoning (steps leading to the first candidate) from reflective reasoning (subsequent steps) and to evaluate whether reflections genuinely contribute to error correction.

Our experiments quantitatively show that across various models and datasets, reflections are largely confirmatory and rarely corrective: once a candidate answer is proposed, subsequent reasoning steps seldom overturn it, instead mainly reiterating or justifying the initial answer. This finding challenges a wide belief that reflections are the primary mechanism for self-correction. It also raises two fundamental questions:

- *If reflections of reasoning models rarely change answers, why is their presence strongly correlated with accuracy?*
- *If reflections mostly confirm earlier conclusions, can we safely truncate them at inference time to reduce computation without significantly harming performance?*

To address these questions, we explore the role of reflections in both training and inference. On the training side, we conduct supervised fine-tuning (SFT) with datasets containing different amounts of reflective reasoning. Our results show that performance gains do not arise from teaching the model to self-correct after mistakes. Instead, **reflections in training data improve performance by increasing the likelihood that the model solves the problem correctly on the first try**. We hypothesize that rollouts with more reflections implicitly expose diverse reasoning paths toward the same problem, which enriches the training distribution and leads to better generalization on unseen problems. On the inference side, motivated by the observation that reflections are mostly confirmatory, we propose a simple yet effective strategy: early stopping of reflections when additional reasoning is unlikely to change the outcome. This method reduces token usage by 24.5% with less than a 2.9% drop in accuracy. Moreover, it allows a dynamic balance between token usage and performance by controlling the early stopping criteria. The contribution of this paper is three-fold:

- **Taxonomy of reflection behavior** (Section 2). We provide the first large-scale analysis of how reasoning models allocate tokens between forward reasoning and reflection, showing that reflections are mostly confirmatory rather than corrective, and the accuracy of the first try is the driving factor of improvement.
- **Training insights** (Section 3). We show that reflection-rich training data improve model accuracy by diversifying reasoning exposure and strengthening first-try correctness, not by enabling error correction.
- **Efficient inference technique** (Section 4). We propose an early-stopping method that reduces token consumption during inference and allows control over the balance between token usage and performance.

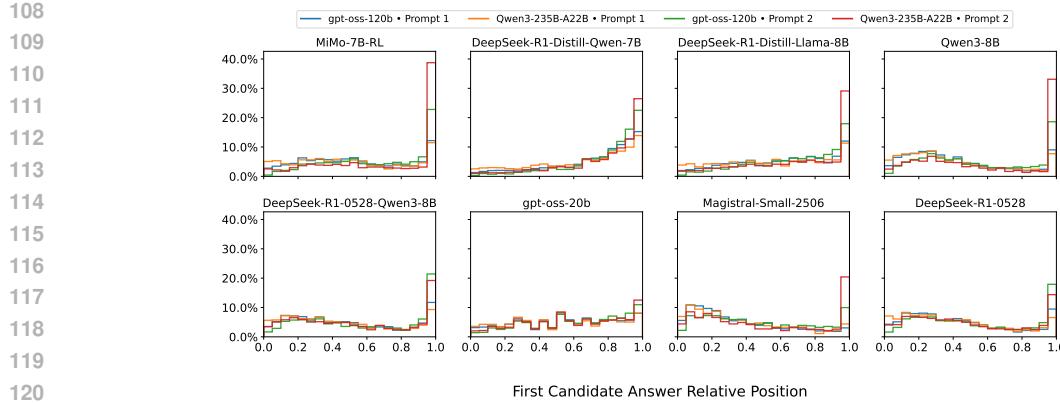


Figure 2: Distribution of first candidate answer positions across different LLMs and prompts. The x-axis denotes the relative position of the first candidate answer (line index divided by total lines), and the y-axis shows the proportion of rollouts in each bin.

2 ANALYZING REFLECTIONS IN REASONING MODELS

2.1 REFLECTION EXTRACTION

Conventional LLMs typically employ a single, linear generation process, concluding upon the initial derivation of a solution. Reasoning models, however, are capable of a more iterative and deliberative methodology (Ke et al., 2025). They can construct significantly more elaborated chain-of-thoughts (CoTs), not simply by extending the quantity of reasoning steps, but by generating and assessing multiple divergent reasoning trajectories and potential answers. This recursive process of refinement, frequently described as “**reflection**”, allows the model to compare alternatives, check intermediate claims, and potentially improve the final answer before committing.

While the capability of performing reflection often correlates with the reasoning capabilities of reasoning models, the internal mechanics of this process remain opaque. To unlock the full potential of these models and understand their decision-making, it is crucial to dissect their reflective patterns. We observe that within a long CoT, there can be multiple positions where the model has already derived a potential answer but opts to continue its reasoning before committing to a final output. Analyzing these critical points is fundamental to understanding the model’s reflection process.

Method In this work, we define “**reflection**” as the contents occurring between two successive candidate answers in the reasoning process. To extract reflections, we introduce an LLM-based candidate answer extractor that parses long CoT outputs and identifies the positions of candidate answers, enabling a structured analysis of the model’s reflective behavior. Specifically, a CoT can be represented as a sequence $C = \{s_1, s_2, \dots, s_N\}$, where each s_i is a reasoning step delimited by a line break. We employ an LLM (see Appendix A for prompts and example input), to extract plausible candidate answers, which can be formally expressed as follows:

$$\text{Extract}(C) = \{(i, a_i) \mid i \in [1, N] \wedge \text{IsCandidateAnswer}(s_i)\}, \quad (1)$$

where `IsCandidateAnswer`(s_i) determines whether step s_i contains a candidate answer, and a_i denotes the extracted candidate answer. Note that the extraction process only requires understanding what quantity the question is asking, and whether a reasoning step derived it, without any requirement on the ability to actually solve the question. As a result, this process yields a structured set of candidate answers and their corresponding positions for subsequent analysis, shown in Figure 1.

Setup We apply our LLM-based extractor on the rollouts of five mathematical benchmarks: AIME2024 (MAA, 2024), AIME2025 (AIME, 2025), AMC (AMC12, 2025), Olympiad Bench (He et al., 2024), Math500 (Hendrycks et al., 2021). Among these benchmarks, Math500 is considered easier, with state-of-the-art models reaching more than 95% accuracy, while AIME2024 and AIME2025 are considered harder, with model pass@1 performance ranging from 30% to 80%. For benchmarks with fewer problems, such as AIME, we increase the number of rollouts per problem to ensure a robust and consistent evaluation (See Table 3 of Appendix B for details). In

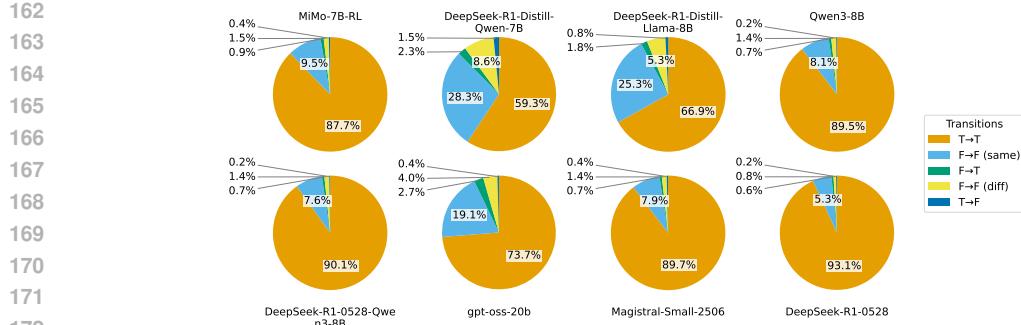


Figure 3: Reflections type statistics of long CoTs of different models. Long CoTs are collected on AIME2024 and AIME2025 (32 rollouts per question), AMC (4 rollouts per question), Olympiad Bench, and Math500 (1 rollout per question). Statistics are compiled for the union of all rollouts. More detailed breakdown of each dataset can be found in Figure 11 of Appendix D.

total, 3,427 rollouts are collected for each of the eight reasoning models evaluated. The studied models cover a wide spectrum of the reasoning model family, with sizes ranging from 7B to 685B, covering models trained with reinforcement learning (RL) (MiMo-7B-RL (Xiaomi et al., 2025), gpt-oss-20b (OpenAI, 2025), Magistral-Small-2506 (Mistral-AI et al., 2025), DeepSeek-R1-0528 (DeepSeek-AI et al., 2025)) and distillation (DeepSeek-R1-Distill-Qwen-7B (DeepSeek-AI et al., 2025), DeepSeek-R1-Distill-Llama-8B (DeepSeek-AI et al., 2025), Qwen3-8B (Yang et al., 2025a), DeepSeek-R1-0528-Qwen3-8B (DeepSeek-AI et al., 2025)).

Robustness Analysis To evaluate the robustness and accuracy of our method, we first conduct a human evaluation. We randomly sample 100 rollouts and ask human participants to evaluate whether the candidate positions and candidate answers extracted by our extractor are reasonable. Across 100 rollouts with 426 extracted candidates in total, human participants labeled 94% of the candidate extraction as correct, which demonstrates the high reliability of the proposed extractor (See Appendix C for more details).

Further, we evaluate the sensitivity of our method by testing four extractor variants, constructed using two different LLMs (Qwen3-235B-A22B (Yang et al., 2025a) and gpt-oss-120b (OpenAI, 2025)) and two distinct extraction prompts (see Appendix A for both complete prompts). To analyze consistency across these configurations, we focus on the relative position of the first candidate answer, which serves as a stable and comparable reference point despite differences in the number and length of reflections across models. The histogram in Figure 2 shows that the distribution of the first candidate answer’s relative position is consistent across different extractors for each model. This demonstrates the robustness and insensitivity of our extractor to the choice of LLMs or prompts, as it consistently captures reflection patterns across different models’ outputs. We use gpt-oss-120b and prompt 1 in Appendix A for the extractor in the rest of the paper.

2.2 REFLECTION ANALYSIS

Reflection Types Using the LLM-based extractor, we identify and extract n candidate answers for each CoT, here we denote them as $\{a_1, a_2, \dots, a_n\}$ indexed by their appearing order. We then evaluate the correctness of each answer a_i using a rule-based verifier², which returns **True** (T) if a_i is correct and **False** (F) otherwise. For two consecutive candidate answers from the same CoT, a_{i-1} and a_i , the **reflection type** is determined by whether the answer’s correctness changes from the previous to the current attempt: (1) T → T: both a_{i-1} and a_i are correct; (2) F → F (same): both a_{i-1} and a_i are incorrect, and $a_{i-1} = a_i$; (3) F → T: a_{i-1} is incorrect, and a_i is correct; (4) F → F (diff): both a_{i-1} and a_i are incorrect, and $a_{i-1} \neq a_i$; (5) T → F: a_{i-1} is correct, and a_i is incorrect. Specifically, we define T → T and F → F (same) reflections as *confirmatory*, since the answer is not changed. And F → T reflections as *corrective*, since it changes an incorrect answer to correct.

²<https://github.com/huggingface/Math-Verify>

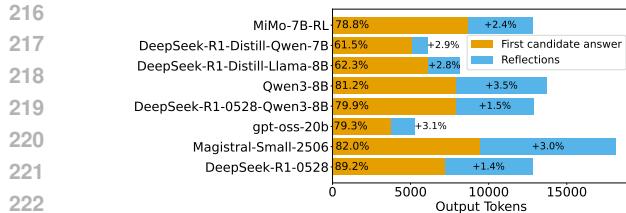


Figure 4: Breakdown of long CoTs: orange bars show the token count up to the first candidate answer, and blue bars show the token count in subsequent reflections. Numbers on bars indicate the accuracy of the first candidate answer, and the accuracy improvement brought by reflections.

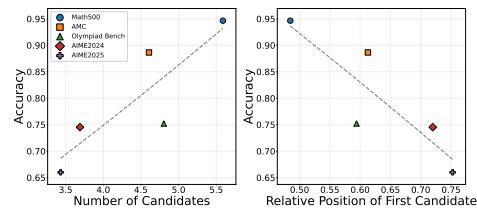


Figure 5: Left: Average number of candidate answers per rollout across different datasets. Right: Relative position of the first candidate. Values are averaged over 8 models.

Analysis on Reflection Types The analysis of reflection types within long CoT, as depicted in Figure 3, reveals that over 90% of the reflections are confirmatory, i.e., $T \rightarrow T$ and $F \rightarrow F$ (same), rather than corrective. This indicates that **most reflections reaffirm an existing answer instead of changing it**. This trend is universal across all tested models and datasets (see Figure 11 in Appendix D for a detailed breakdown). Crucially, the proportion of corrective reflections ($F \rightarrow T$) that actually improve performance is exceptionally small (mostly less than 2%).

Impact of Reflection on Performance To better understand the impact of reflections on performance and their token usage, we compare the accuracy of the first candidate answer with that of the final answer in the rollouts as shown in Figure 4. The accuracy of the first candidate answer is given on the orange bar segments, and the contribution of subsequent reflections is shown on the blue segments. Reported accuracies are averaged over AIME2024, AIME2025, AMC, Olympiad Bench, and Math500. We report this averaged accuracy unless mentioned otherwise throughout the paper. We observe that while reflections after the first candidate answer consume a large portion of the total tokens (ranging from 16.8% to 47.8%), the resulting performance gain is limited (ranging from 1.4% to 3.5%). This suggests that the final accuracy strongly correlates with the correctness of the first answer, highlighting its dominant influence. In other words, **the first try matters**. We provide a detailed breakdown in Tables 5 and 6 of Appendix G.

Effect of Data Difficulty on Reflection Patterns To analyze the reasoning models' reflection patterns across various mathematical datasets of different difficulties, we plot the average number of candidates and the average relative position of the initial candidate in Figure 5 (See Table 4 in Appendix E for detailed statistics). Our analysis reveals that on more challenging datasets, such as AIME2024 and AIME2025, the model allocates more tokens to forward reasoning, delaying the appearance of the first candidate. Conversely, on easier datasets such as Math500, the first candidate appears much earlier in the reasoning trajectory. This presents a counterintuitive pattern: **models perform more reflections on easier problems and fewer on difficult ones**, indicating that the reflection mechanism of reasoning models is not well aligned with task difficulty.

3 THE ROLE OF REFLECTION IN REASONING MODEL TRAINING

Analysis of reflection patterns in reasoning model rollouts reveals a counterintuitive phenomenon: the majority of reflections neither alter the candidate answer nor contribute meaningfully to performance improvement. This observation raises a critical question: why do reasoning models still achieve a substantial performance boost after being trained on long CoTs containing many confirmatory reflections? In this section, we study this by conducting supervised fine-tuning (SFT) on curated datasets with different reflection characteristics and comparing their performance.

3.1 TRAINING WITH VARYING AMOUNT OF REFLECTIONS

To investigate the role of reflection in reasoning, we begin by examining how the amount of reflections included in training data affects model performance.

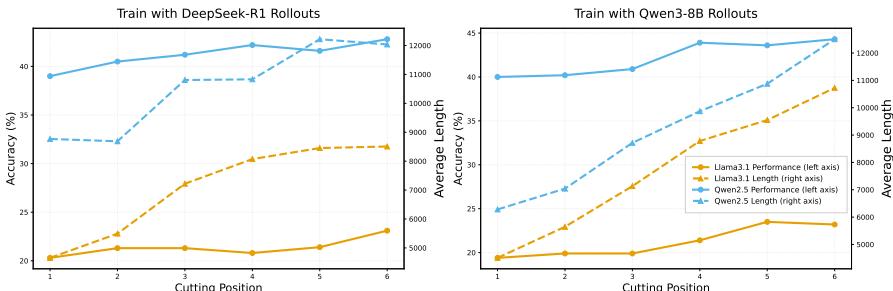


Figure 6: Comparison of performance and rollout length after SFT when training on rollouts cut at different positions. Qwen2.5-7B-Instruct and Llama3.1-8B-Instruct are trained using processed rollouts from DeepSeek-R1 and Qwen3-8B, respectively. Accuracies are averaged over five datasets.

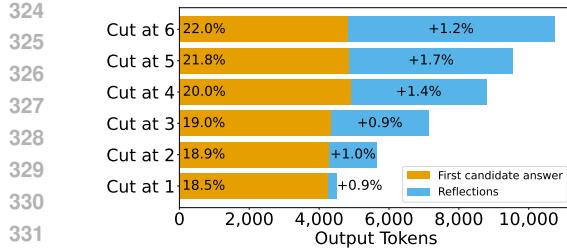
Training Data Construction To study the effect of the number of reflections in long CoT data, we carefully manipulate the rollouts by truncating them after the occurrence of different candidate answers, thereby constructing a long CoT SFT dataset with a controllable number of reflections. Specifically, we use the MiroMind-M1-SFT dataset (Li et al., 2025), which contains mathematical problems curated from diverse sources, along with their corresponding DeepSeek-R1 rollouts. Additionally, we generate one rollout per problem using Qwen3-8B, providing an alternative rollout source for comparison. We then apply the candidate answer extractor to filter out rollouts that satisfy: (1) produce a correct final answer, and (2) the correct answer appears more than six times as a candidate. This filtering step ensures that each selected rollout includes sufficient reflections, allowing flexible truncation from the first to the sixth candidate answers to construct datasets with a varying number of reflections.

To create a rollout with exactly i reflection steps, we truncate a filtered rollout at the i -th candidate answer, append the stop-thinking symbol `</think>`, then feed the truncated rollout into the reasoning model that used to generate it (e.g., DeepSeek-R1 or Qwen3-8B) to continue generation and produce the final answer. This yields a reasoning rollout with exactly i reflection steps. See Appendix Figure 12 for an example. By continuing the truncated thinking, we ensure that the rollouts we used for training are still coherent, without abrupt stops. We filter out rollouts whose answers in the continued generation are different from the candidate answer from which we truncated. This step removes less than 0.5% of rollouts, showing that once the model settles on a candidate answer, stopping the thinking process and prompting for a final response reliably yields that same answer.

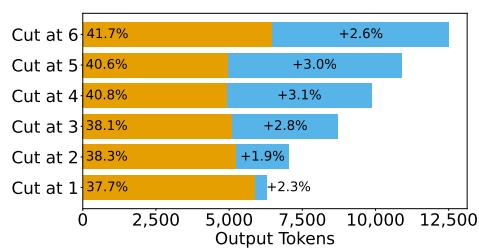
Applying this procedure on the MiroMind-M1-SFT dataset, we curate six SFT datasets. The i -th dataset, termed the “cut-at- i ” dataset, contains long CoTs truncated at the i -th reflection, resulting in exactly i reflections per example. These datasets share the same set of problems, and all rollouts correctly solve the problem. The difference between them lies in the controlled number of reflections. Another variable we need to control is that, CoTs with more reflections typically contain more tokens. Therefore, for fair comparison, we downsample these datasets to ensure all six have the same number of training tokens. After this process, each generated dataset contains 28 million tokens, with the cut-at-1 dataset having 6,754 questions and the cut-at-6 dataset having 3,405 questions.

Impact of Reflection Amount on SFT Performance Given the curated dataset, we perform SFT on Llama3.1-8B-Instruct (Grattafiori et al., 2024), and Qwen2.5-7B-Instruct (Yang et al., 2024). We test on the combined set of AIME24, AIME25, Olympiad Bench, AMC, and Math500. The performance results and corresponding rollout lengths are presented in Figure 6. It shows that training on reflection-rich rollouts yields higher accuracy and longer generations across different datasets and model architectures. For SFT models, this suggests that, under a fixed token budget, constructing datasets with more reflection-rich rollouts is more effective than using the same budget to include more questions with shorter rollouts.

To better understand this improvement, we leverage the extractor and we split each rollout at the first candidate answer: the segment before it is denoted as “first candidate answer”, and the segment after it as “reflections”. We then compare their corresponding lengths and accuracies. Figure 7 (plotted in the same manner as Figure 4) shows results trained on Qwen3-8B rollouts, with similar DeepSeek-R1 results in Figure 13 of Appendix H.



(a) Llama3.1-8B-Instruct



(b) Qwen2.5-7B-Instruct

Figure 7: Token usage and accuracy after SFT using Qwen3-8B rollouts. Before SFT, Llama3.1-8B-Instruct achieves 7.9% accuracy, and Qwen2.5-7B-Instruct achieves 35.3%.

Figure 7 demonstrates a clear trend that models trained on rollouts with more reflections achieve higher final performance. Averaging over Llama3.1-8B-Instruct and Qwen2.5-7B-Instruct, cut-at-6 outperforms cut-at-1 by 4.05%. This performance gain is mainly due to higher accuracy of the first candidate answers, which increase by an average of 3.75% from cut-at-1 to cut-at-6, whereas the contribution from additional reflections is much smaller, averaging only 0.3%. Interestingly, while the accuracy of the first answer is increased, the token cost of generating the first answer remains consistent after trained with different cut-at- i datasets. The tokens spent on reflections account for most of the difference, with an average of 5,636 reflection tokens increase per rollout from cut-at-1 to cut-at-6.

In conclusion, training models with more reflections leads to better performance and longer responses, which is expected. Surprisingly, from the breakdown of the source of improvements, we find that this gain does not come from reflections fixing incorrect answers, but from higher accuracy in the first candidate answer. One possible explanation is that richer reflections expose the model to diverse problem-solving approaches, improving generalization and boosting initial answer quality rather than simply correcting mistakes.

Discussions Our analysis in Figure 7 shows that SFT distillation enhances overall performance primarily by improving first-try correctness, especially when the SFT data includes more reflections, while the improvement brought by reflections is marginal. Our previous analysis in Figure 3 also shows that both RL-trained and SFT-distilled reasoning models show a similar pattern that reflections are mostly confirmatory and do not bring improvement. This raises the question of whether the RL training stage of reasoning models is also improving accuracy by making better first tries. To investigate this, we compare behaviors before and after RL of open-source reasoning models: MiroMind-M1-RL-7B (Li et al., 2025) and its initialization, MiroMind-M1-SFT-7B; MiroMind-M1-RL-32B and its initialization, DeepSeek-R1-Distill-Qwen-32B. As illustrated in Figure 8, we see that for both models, the performance gains after RL mainly come from first-answer accuracy improvement (+4.6% for 32B model, and +7.7% for 7B model), while gains attributable to reflections are marginal (+0.3% for 32B model, and +0.1% for 7B model). The experimental results indicate that during reinforcement learning, reasoning models primarily enhance their ability to produce a correct answer on the first attempt, rather than improving the quality of their subsequent reflections.

3.2 TRAINING WITH CORRECTIVE REFLECTION PATTERNS

In the previous section, we investigated how reflections affect model performance in SFT, and showed that more reflections in training rollout mainly improve first answer accuracy, with limited improvement on its corrective behavior. In this section, we test whether reflection ability can be improved by adding more corrective reflections (i.e., $F \rightarrow T$) to the training dataset.

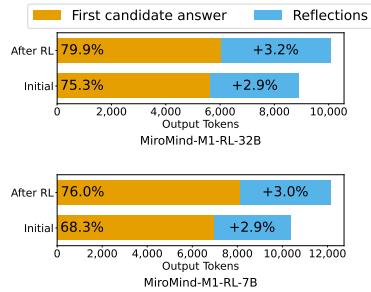


Figure 8: Changes of reasoning behavior after RL.

378 Table 1: Performance after SFT with different ratios of $F \rightarrow T$ reflections in the dataset.
379

380 381 Ratio (%)	382 Llama3.1-8B-Instruct			383 Qwen2.5-7B-Instruct		
	384 $p(F \rightarrow T)$	385 Accuracy (%)	386 Length	387 $p(F \rightarrow T)$	388 Accuracy (%)	389 Length
100	0.053	26.6	10830	0.036	44.1	9655
75	0.050	25.6	11085	0.043	44.4	8775
50	0.058	27.3	10746	0.045	43.1	9943
25	0.059	26.2	11295	0.046	44.8	10094
0	0.050	26.9	11419	0.041	44.1	9363

388 **Dataset Construction** We collect Qwen3-8B rollouts on math problems from the MiroMind-M1-
389 SFT dataset. For each question, we sample one rollout containing at least one $F \rightarrow T$ reflection
390 and one rollout consisting solely of $T \rightarrow T$ reflections. By filtering problems that have both type
391 of rollouts (corrective and confirmatory), we kept 6K problems. Using these problems and their
392 rollouts, we construct five datasets by varying the proportion of problems for which we select their
393 corrective rollout to include in the dataset: 0%, 25%, 50%, 75%, and 100%. For the remaining
394 problems in each dataset, we select their confirmatory rollouts to include in the dataset.

395 **Impact of Corrective Reflections on SFT Performance** We perform SFT on Llama3.1-8B-
396 Instruct and Qwen2.5-7B-Instruct using the constructed datasets, with results shown in Table 1.
397 Models trained on datasets with varying proportions of corrective reflections show similar response
398 lengths and accuracies. The performance difference between the best and worst models is just 1.7%,
399 and the maximum difference in response length is only about 1K tokens. Moreover, their ability
400 to flip an incorrect answer to a correct one, measured by $p(F \rightarrow T)$, the probability that the next
401 candidate is correct given that the current candidate is incorrect, shows no improvement. This indicates
402 that training on rollouts containing corrective reflections is not more beneficial than training
403 on rollouts with only confirmatory reflections. This echos with our earlier analysis, reasoning im-
404 provements are reflected mainly as higher first-answer accuracy, rather than increased $p(F \rightarrow T)$.
405

4 EFFICIENT REASONING BY EARLY STOPPING

406 Our studies in Section 2 show that the reflections of reasoning models are primarily confirmatory,
407 which suggests potential token efficiency gains by stopping once a few candidate answers are iden-
408 tified. In this section, we study the token-accuracy tradeoff under different early-stopping strategies.
409 Specifically, we propose a question-aware adaptive early-stopping approach to improve token effi-
410 ciency of the reasoning process.

411 **Candidate Answer Detector** A straightforward way to reduce confirmatory reflection tokens is
412 to monitor candidate answers and early-stop the reasoning process once a correct one is generated.
413 To achieve this, we train a Qwen3-1.7B-based candidate answer detector (CAD) to detect for each
414 sentence in the generation whether it contains the candidate answer. We construct CAD training
415 data from annotated rollouts in the MiroMind-M1-SFT dataset. The sentences in each rollout, de-
416 limited by \n, are annotated by gpt-oss-120b and labeled 1 if they contain a candidate answer, or 0
417 otherwise. The CAD takes the corresponding question and one sentence in the rollout as input, and
418 is trained to predict whether the sentence contains a candidate answer.
419

420 **Question-aware Reflection Controller** While reflections are mostly confirmatory in reasoning
421 rollouts, some mathematical problems may benefit more from reflections than others. To identify
422 such problems and give them more reflection budget, we train a question-aware reflection controller
423 (QRC) to predict for a problem whether we should stop at the first candidate, or allow more reflec-
424 tions before early-stopping. Specifically, we train a Qwen3-1.7B-based binary classifier that takes
425 in only the problem statement, and output a binary label. The training data is collected from the
426 annotated MiroMind-M1-SFT dataset, where a question is labeled 1 if its rollout contains $F \rightarrow T$
427 reflections, otherwise 0.

428 **Question-aware Adaptive Early-Stopping** With CAD and QRC, we can reduce unnecessary re-
429 flections in the reasoning process through question-aware adaptive early-stopping. During inference,
430 we first feed the question into the QRC to determine whether the reasoning process should terminate
431 at the first candidate answer or do more reflections. Then we use CAD to monitor the appearance of

432 Table 2: Question-aware adaptive early-stopping improves token efficiency. Without QRC, reasoning
 433 terminates immediately after the first candidate answer is generated. The classification thresh-
 434 olds of QRC and CAD are set as 0.05 and 0.5, respectively.

435

436 Dataset	437 Accuracy (%)			438 Length		
	439 Qwen3-8B	440 +CAD	441 +CAD, +QRC	442 Qwen3-8B	443 +CAD	444 +CAD, +QRC
AIME2024	82.1	77.9 (-4.2)	79.6 (-2.5)	18,962	13,517 (-28.7%)	14,869 (-21.6%)
AIME2025	70.8	65.0 (-5.8)	65.8 (-5.0)	22,998	17,664 (-23.2%)	18,014 (-21.7%)
AMC	93.0	90.0 (-3.0)	89.4 (-3.6)	13,279	8,432 (-36.5%)	8,756 (-34.1%)
Math500	97.4	94.4 (-3.0)	96.0 (-1.4)	5,755	2,912 (-49.4%)	3,593 (-37.6%)
Olympiad Bench	80.2	76.9 (-3.3)	78.4 (-1.8)	14,633	10,479 (-28.4%)	11,835 (-19.1%)
Average	84.7	80.9 (-3.8)	81.8 (-2.9)	15,125	10,601 (-29.9%)	11,414 (-24.5%)

445

candidate answers during generation and terminate thinking accordingly. In practice, we terminate
 at the first candidate if QRC labeled 0, otherwise the third candidate. We apply this approach to
 Qwen3-8B reasoning model and report the performance on five mathematical datasets in Table 2.

446

Table 2 illustrates that CAD saves on average 29.9% tokens across five mathematical datasets, with a modest 3.8% drop. With QRC, the performance drop is improved to 2.9%, while still enjoying a 24.5% token reduction. By controlling the classification thresholds of CAD and the QRC, our method provides a handle to balance between performance and token usage. Figure 9 illustrates the trade-off between token reduction and performance by adopting different threshold settings. On one extreme, a modest 1 percentage-point accuracy drop allows a 12.0% reduction in tokens; on the other extreme of the trade-off, an 8.12% accuracy drop corresponds to a 40.7% reduction in tokens.

450

5 RELATED WORKS

451

In the past year, scaling inference-time computation has emerged as a promising paradigm for improving the capabilities of large language models (LLMs) (Snell et al., 2024; Brown et al., 2025; Muennighoff et al., 2025). Building on this line of work, DeepSeek-R1 (DeepSeek-AI et al., 2025) demonstrates that inference-time scaling via reinforcement learning with verifiable rewards (RLVR) can unlock emergent reasoning abilities, yielding state-of-the-art results on challenging mathematical and coding benchmarks (Liu et al., 2025b; Ke et al., 2025; Xie et al., 2025; Jain et al., 2025).

452

Recently, a complementary body of research has examined why long chain-of-thought (CoT) reasoning is effective (Zhao et al., 2025; Chen et al., 2025a; Jiang et al., 2025). The dominant belief is that its success can be attributed to mechanisms such as recursive reflection, verification, and revision, which allow models to refine intermediate steps (Yang et al., 2025b; Wang et al., 2025). At the same time, excessively verbose traces are observed often to introduce redundancy, amplify hallucinations, and degrade the performance of reasoning models (Chen et al., 2025b; Zeng et al., 2025). In contrast, a separate line of literature argues that prolonged traces can strengthen reasoning ability and promote exploration of diverse solutions (Liu et al., 2025b;a). Despite these advances, the role of reflections within long CoTs remains underexplored. To address this gap, our work isolates and studies reflections within long CoTs across training and testing stages.

453

6 CONCLUSIONS

454

In this work, we systematically analyze the reflection pattern in long CoTs of reasoning models. We investigate their role in both the training and the inference phases. Through extensive experiments, we show that the reflections of reasoning models are mostly confirmatory, yet they are still helpful when included in training data. We also show that during inference time, confirmatory reflections consume a decent amount of tokens, while only introducing marginal improvements. To this end, we develop an efficient reasoning technique during inference to early stop excessive reflections while maintaining the performance. Together, these results provide a clearer understanding of the role of reflections and offer practical guidance for data design and inference efficiency.

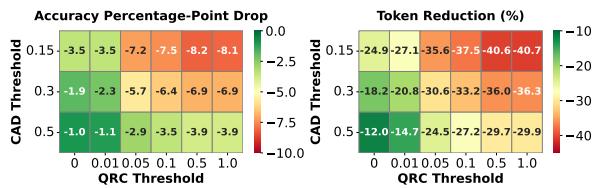


Figure 9: Accuracy drop and token reduction with varying classification thresholds of CAD and QRC.

486 ETHICS STATEMENT
487488 We did not collect any new human-subject data or process personally identifiable information. No
489 potential harms, such as bias amplification, toxic outputs, or dual-use risks, were observed in this
490 work.
491492 REPRODUCIBILITY STATEMENT
493494 We are committed to reproducibility. All details of the data used and generated are documented in
495 the paper, along with comprehensive descriptions of our experimental setup and hyperparameters.
496 Upon publication, we will release the source code and data to the public to support transparent
497 validation and further research.
498499 USE OF LLMs
500501 The use of LLM in this work is limited to rewriting sections to shorten them and save space for the
502 page limit.
503504 REFERENCES
505506 AIME. AIME Problems and Solutions, 2025. https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions, 2025. Accessed: September 23,
507 2025.509 AMC12. AMC 12 Problems and Solutions. https://artofproblemsolving.com/wiki/index.php/AMC_12_Problems_and_Solutions, 2025. Accessed: September
510 23, 2025.512 Bradley Brown, Jordan Juravsky, Ryan Saul Ehrlich, Ronald Clark, Quoc V Le, Christopher Re, and
513 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling,
514 2025. URL <https://openreview.net/forum?id=0xUEBQV54B>.
515516 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
517 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
518 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021. URL <https://arxiv.org/abs/2107.03374>.
519520 Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu,
521 Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-
522 of-thought for reasoning large language models, 2025a. URL <https://arxiv.org/abs/2503.09567>.
523524 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuwei
525 Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Do
526 NOT think that much for $2+3=?$ on the overthinking of long reasoning models. In *Forty-second*
527 *International Conference on Machine Learning*, 2025b. URL <https://openreview.net/forum?id=MSbU3L7V00>.
528529 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
530 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
531 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. URL <https://arxiv.org/abs/2110.14168>.
532533 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
534 Qihao Zhu, Shirong Ma, Peiyi Wang, et al. Deepseek-r1: Incentivizing reasoning capability in
535 llms via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>.
536537 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
538 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
539 models. *arXiv preprint arXiv:2407.21783*, 2024. URL <https://arxiv.org/abs/2407.21783>.
540

540 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,
 541 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for
 542 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*
 543 *arXiv:2402.14008*, 2024. URL <https://aclanthology.org/2024.acl-long.211/>.

544 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 545 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 546 *preprint arXiv:2103.03874*, 2021. URL <https://arxiv.org/abs/2103.03874>.

547 Arnav Kumar Jain, Gonzalo Gonzalez-Pumariega, Wayne Chen, Alexander M Rush, Wenting Zhao,
 548 and Sanjiban Choudhury. Multi-turn code generation through single-step rewards. In *Forty-*
 549 *second International Conference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=aJeLhLcsh0>.

550 Gangwei Jiang, Yahui Liu, Zhaoyi Li, Qi Wang, Fuzheng Zhang, Linqi Song, Ying Wei, and Defu
 551 Lian. What makes a good reasoning chain? uncovering structural patterns in long chain-of-
 552 thought reasoning, 2025. URL <https://arxiv.org/abs/2505.22148>.

553 Zixuan Ke, Fangkai Jiao, Yifei Ming, Xuan-Phi Nguyen, Austin Xu, Do Xuan Long, Minzhi Li,
 554 Chengwei Qin, Peifeng Wang, Silvio Savarese, Caiming Xiong, and Shafiq Joty. A survey of fron-
 555 tiers in LLM reasoning: Inference scaling, learning to reason, and agentic systems. *Trans. Mach.*
 556 *Learn. Res.*, 2025, 2025. URL <https://openreview.net/forum?id=SlsZZ25InC>.

557 Xingxuan Li, Yao Xiao, Dianwen Ng, Hai Ye, Yue Deng, Xiang Lin, Bin Wang, Zhanfeng Mo,
 558 Chong Zhang, Yueyi Zhang, Zonglin Yang, Ruilin Li, Lei Lei, Shihao Xu, Han Zhao, Weiling
 559 Chen, Feng Ji, and Lidong Bing. Miromind-m1: An open-source advancement in mathematical
 560 reasoning via context-aware multi-stage policy optimization, 2025. URL <https://arxiv.org/abs/2507.14683>.

561 Mingjie Liu, Shizhe Diao, Jian Hu, Ximing Lu, Xin Dong, Hao Zhang, Alexander Bukharin,
 562 Shaokun Zhang, Jiaqi Zeng, Makesh Narsimhan Sreedhar, Gerald Shen, David Mosallanezhad,
 563 Di Zhang, Jonas Yang, June Yang, Oleksii Kuchaiev, Guilin Liu, Zhiding Yu, Pavlo Molchanov,
 564 Yejin Choi, Jan Kautz, and Yi Dong. Scaling up rl: Unlocking diverse reasoning in llms via
 565 prolonged training, 2025a. URL <https://arxiv.org/abs/2507.12507>.

566 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
 567 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models,
 568 2025b. URL <https://arxiv.org/abs/2505.24864>.

569 Zichen Liu, Changyu Chen, Wenjun Li, Tianyu Pang, Chao Du, and Min Lin. There may not
 570 be aha moment in r1-zero-like training a pilot study. <https://oatlim.notion.site/oat-zero>, 2025c. Notion Blog.

571 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin
 572 Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-
 573 preview with a 1.5b model by scaling rl. <https://github.com/agentica-project/deepscaler>, 2025.

574 MAA. American Invitational Mathematics Examination (AIME) 2024. <https://maa.org/math-competitions/american-invitational-mathematics-examination-aime>,
 575 February 2024. Accessed: September 23, 2025.

576 MiniMax, :, Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan,
 577 Changqing Yu, et al. Minimax-m1: Scaling test-time compute efficiently with lightning attention,
 578 2025. URL <https://arxiv.org/abs/2506.13585>.

579 Mistral-AI, :, Abhinav Rastogi, Albert Q. Jiang, Andy Lo, Gabrielle Berrada, Guillaume Lample,
 580 Jason Rute, Joep Barmentlo, et al. Magistral, 2025. URL <https://arxiv.org/abs/2506.10910>.

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

594 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David
 595 Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Au-
 596 gustus Odena. Show your work: Scratchpads for intermediate computation with language models,
 597 2021. URL <https://arxiv.org/abs/2112.00114>.

598
 599 OpenAI. gpt-oss-120b & gpt-oss-20b model card, 2025. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2508.10925)
 600 2508.10925.

601
 602 OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden
 603 Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card,
 604 2024. URL <https://arxiv.org/abs/2412.16720>.

605 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally
 606 can be more effective than scaling model parameters, 2024. URL [https://arxiv.org/](https://arxiv.org/abs/2408.03314)
 607 abs/2408.03314.

608
 609 Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu,
 610 Juntao Li, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. Thoughts are
 611 all over the place: On the underthinking of o1-like llms, 2025. URL [https://arxiv.org/](https://arxiv.org/abs/2501.18585)
 612 abs/2501.18585.

613 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H.
 614 Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in
 615 large language models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave,
 616 K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35 (NeurIPS*
 617 2022), 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html.

618
 619 LLM-Core Xiaomi, :, Bingquan Xia, Bowen Shen, Cici, Dawei Zhu, Di Zhang, Gang Wang, Hailin
 620 Zhang, Huaqiu Liu, et al. Mimo: Unlocking the reasoning potential of language model – from
 621 pretraining to posttraining, 2025. URL <https://arxiv.org/abs/2505.07608>.

622
 623 Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu,
 624 Zhirong Wu, and Chong Luo. Logic-rl: Unleashing llm reasoning with rule-based reinforcement
 625 learning, 2025. URL <https://arxiv.org/abs/2502.14768>.

626
 627 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu,
 628 Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu,
 629 Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical
 630 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024. URL <https://arxiv.org/abs/2412.15115>.

631
 632 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 633 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*
 634 *arXiv:2505.09388*, 2025a. URL <https://arxiv.org/abs/2505.09388>.

635
 636 Shiming Yang, Yuxuan Tong, Xinyao Niu, Graham Neubig, and Xiang Yue. Demystifying long
 637 chain-of-thought reasoning. In *Forty-second International Conference on Machine Learning*,
 638 2025b. URL <https://openreview.net/forum?id=OLodUbcWjB>.

639
 640 Shu Yang, Junchao Wu, Xin Chen, Yunze Xiao, Xinyi Yang, Derek F. Wong, and Di Wang. Un-
 641 derstanding aha moments: from external observations to internal mechanisms, 2025c. URL
 642 <https://arxiv.org/abs/2504.02956>.

643
 644 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai,
 645 Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guang-
 646 ming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu,
 647 Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao
 Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingx-
 uan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025. URL
 648 <https://arxiv.org/abs/2503.14476>.

648 Zhiyuan Zeng, Qinyuan Cheng, Zhangyue Yin, Yunhua Zhou, and Xipeng Qiu. Revisiting the
649 test-time scaling of o1-like models: Do they truly possess test-time scaling capabilities? In
650 Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Pro-
651 ceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Vol-
652 ume 1: Long Papers)*, pp. 4651–4665, Vienna, Austria, July 2025. Association for Compu-
653 tational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.232. URL
654 <https://aclanthology.org/2025.acl-long.232/>.

655 Chong Zhang, Yue Deng, Xiang Lin, Bin Wang, Dianwen Ng, Hai Ye, Xingxuan Li, Yao Xiao,
656 Zhanfeng Mo, Qi Zhang, and Lidong Bing. 100 days after deepseek-r1: A survey on replication
657 studies and more directions for reasoning language models, 2025. URL <https://arxiv.org/abs/2505.00551>.

658 Chengshuai Zhao, Zhen Tan, Pingchuan Ma, Dawei Li, Bohan Jiang, Yancheng Wang, Yingzhen
659 Yang, and Huan Liu. Is chain-of-thought reasoning of llms a mirage? a data distribution lens,
660 2025. URL <https://arxiv.org/abs/2508.01191>.

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 A PROMPT FOR CANDIDATE EXTRACTION
703704 SYSTEM PROMPT 1
705706 SYSTEM PROMPT Candidate Answer Extractor
707708 Role: You read a math problem statement and a line-numbered model
709 solution (thinking only).
710711 Goal: For every line that presents a candidate answer to what the
712 problem asks, output a record of (line_number,
713 "candidate_answer_in_required_form"). Do not judge correctness
714 or re-derive the solution.
715716 1) Golden rule: lock the target first
717 Silently infer exactly what the original problem asks for:
718 - Target quantity: e.g., $m+n$, "remainder mod 1000", "sum of digits
719 of N ", "area", "number of solutions", etc.
720 - Required output form: e.g., integer, simplified fraction, decimal
721 to k d.p., radical with squarefree radicand, gcd/coprime
722 conditions, modulo residue, floor/ceil, units, etc.
723 - Trivial post-processing, if any (see §3).
724 Keep this target in working memory. All extraction converts to this
725 target form.
726727 2) What counts as a candidate answer
728 A line presents a candidate answer if it directly gives the target
729 or uniquely determines it after only the trivial final steps in
730 §3.
731732 Include lines that:
733 - State the target explicitly (e.g., "Thus $m+n=38$ ", "Answer: 456").
734 - Give an equivalent numeric/expression that becomes the target
735 after trivial conversion (e.g., a remainder before reduction,
736 ab before radical normalization when the target is $m+n$, a
737 fraction before simplification when the target is "lowest
738 terms", a raw integer before taking "last 3 digits", etc.).
739 - Re-present the candidate (e.g., boxed, restated, "Therefore "),
740 even if previously seen. Record every explicit presentation.
741742 Exclude lines that:
743 - Only give intermediate facts not uniquely tied to the target
744 (ranges/inequalities/bounds, generic identities, unspecialized
745 parameters) unless the problem asks for those.
746 - Require nontrivial algebra, case analysis, multi-step geometry,
747 or symbolic manipulation to reach the target (see §4).
748749 3) Trivial Final Steps (you must perform these when applicable)
750 When a line yields a value that is one obvious step away from the
751 target, perform the step and record the final target value:
752 - Simple arithmetic on explicit numerics/rationals
753 (add/subtract/multiply/divide; reduce fractions to lowest
754 terms).
755 - Modulo: compute $N \bmod m$; extract last k digits; compute parity.
- Digit ops: sum/product of digits; last digit.
- Floor/Ceil/Abs/Sign when directly evaluable on a numeric
expression.
- Rounding exactly as requested (e.g., to 3 d.p.).
- Radicals normalization: rewrite ab with squarefree b , absorb
perfect-square factors into a , ensure $\gcd(a,b)=1$.
- If the problem then asks for $m+n$, output that integer.
- Composite "reporting forms" common in contests:
- If target is $m+n$ from mn (squarefree n), or $p+q$ from a
reduced fraction p/q , compute and output the sum.
756

```

756
757     - If target is "remainder", "units digit", "sum of
758     coefficients", "sum of roots (given the polynomial)", etc., and
759     the line gives the immediately-evaluable precursor, do the
760     one-step conversion.
761     - Unit conversion if it's a fixed scalar multiply/divide stated by
762     the problem.
763     Never cross into multi-step derivations. If it's more than a short,
764     mechanical evaluation, do not include.
765
766     4) Nontrivial (do not do)
767     - No solving new equations, factoring beyond extracting perfect
768     squares for radicals normalization, trig/geometry multi-steps,
769     solving systems, case splits, or applying the quadratic formula
770     unless it's already fully computed in the line.
771     - No deducing implicit constraints unless the line states the value
772     that pins the target after a trivial step.
773
774     5) High-recall detection heuristics
775     When scanning a line, look for any of the following cues. If
776     present, attempt extraction.
777
778     Textual cues:
779     "so", "thus", "therefore", "hence", "we get", "equals", "is",
780     "becomes", "gives", "yields", "implies", "it follows",
781     "answer", "result", "final", "box/boxed".
782
783     Math cues:
784     - An equality or assignment (e.g.,  $=$ ,  $\equiv$ , if exact), explicit
785     numerals, simplified forms, isolated expressions at the end of
786     a derivation.
787     - Named quantity matching the target (e.g., "remainder = 456", "sum
788     of digits is 6").
789     - Expressions that trivially map to target form (e.g., 1218 when
790     target demands  $m+n$ ).
791
792     Repeat cues:
793     If a line reasserts or updates a candidate, record it again with
794     that line number.
795
796     6) Per-line extraction algorithm (do this for each line
797     independently)
798     1. Collect candidates on this line:
799         - Parse any explicit equalities/values/boxed content.
800         - Note any expression that can be trivially converted to the
801         target via  $\text{g3}$ .
802     2. Resolve to target form:
803         - Apply only  $\text{g3}$  operations; otherwise stop.
804         - If multiple possible candidates appear on the same line,
805         record each separately.
806     3. If successful, emit (line_number, "value_in_required_form").
807     If no candidate survives  $\text{g3}$ , skip the line.
808
809     7) Output format (STRICT)
810     - Output a Python list of tuples only: [(line, "value"), (line,
811         "value"), ...]
812     - Keep tuples in the order of increasing line number; if multiple
813         candidates on the same line, keep their left-to-right
814         occurrence order.
815     - "value" must be exactly what the problem asks for after trivial
816         conversion (e.g., put "38", not "36*sqrt(2)" when the target is
817          $m+n$ ).
818     - The very last line of your reply must be only that list so eval()
819         can parse it. No extra text.

```

```

810
811
812 8) Micro-examples (apply ȏ3 automatically)
813 - Remainder: Line has N = 123456; target is N mod 1000 $\to$ record
814     "456".
815 - Last 3 digits: Line has S = 7000456 $\to$ "456".
816 - Sum of digits: Line has N = 1002003 $\to$ "6".
817 - Reduced fraction: Line has 84/126 $\to$ "2/3".
818 - Radical m+n: Line has 1218 $\to$ normalize to 362 $\to$ m+n =
819     36+2 = "38".
820 - Floor: Line has 7.99 and target is the integer part $\to$ "7".
821 Edge case principle: When in doubt, include if the target is
822     uniquely determined by a single, trivial step.
823

```

SYSTEM PROMPT 2

You are given a text block that contains the original problem statement, followed by a line-numbered "model solution".

Your job is **NOT** to judge correctness or solve the problem. Instead, read the solution **line by line** and record every line that presents a ***candidate answer*** to the problem. You need to fully understand what the problem asks for to notice the candidate answer. Only the thinking part of the model solution is provided for analysis.

Definitions

Candidate answer any explicit value or statement that (a) directly answers what the problem asks **or** (b) uniquely determines it with only a trivial final step (e.g. once you know N, taking "N mod 1000" is immediate).
Candidate answer is not intermediate components (like individual addends when the question asks for their sum) unless the problem explicitly asks for each component.

If the line gives an expression that still needs a trivial final computation to directly answer the question (for instance, a fraction whose numerator and denominator you must sum), carry out that simple arithmetic and record the result as your "candidate answer."

There is likely multiple candidates answers in the model solution, and they are not necessarily the same as the model's final answer. You should not look for candidate answers by matching the model's final answer.

You can reason about the lines and decide whether they are candidate answers. For the final response, you should follow the format as below.

Final output format strict

1. For each qualifying line output a two-element tuple:
`(line_number, "candidate_answer")`
`line_number` is an integer.
`candidate_answer` is the exact answer text you extracted from that line (no boxing, no extra words) OR the answer that can be immediately implied from the line. Continuing the previous example, if the line indicates N=2016, the extracted candidate answer should be 16.
2. Collect the tuples in a Python list **in the order the lines appear**.
3. The **last line** of your reply must be **only** that list, so that `eval()` can parse it, for example, `[(12, "15"), (27, "3/4")]`

```

864
865   4. Do **not** output anything after that list.
866
867
868 EXAMPLE INPUT
869
870   Analyze the following problem and its model solution.
871
872   -----
873   Below is the problem statement **followed by** the line-numbered
874   model solution:
875   -----
876
877   Problem statement:
878   Find all prime numbers $p$ and positive integers $m$ such that
879    $2p^2 + p + 9 = m^2$ 
880
881   Model solution:
882   1: <think>
883   2: Okay, so I need to find all prime numbers p and positive
884   integers m such that the equation  $2p^2 + p + 9 = m^2$  holds. Hmm,
885   let's start by understanding the problem. I have to find primes
886   p and positive integers m where this quadratic in p becomes a
887   perfect square.
888   ...
889
890
891
892

```

B BENCHMARK STATISTICS

Table 3: Statistics of Datasets and Rollouts

Dataset	Problems	Rollouts per Problem	Total Rollouts
AIME 2024	30	32	960
AIME 2025	30	32	960
AMC	83	4	332
Olympiad Bench	675	1	675
MATH500	500	1	500
Total	1318	-	3427

C HUMAN EVALUATION

To validate our extraction method, we conduct a human annotation study with four participants. We task the annotators with labeling the correctness of each candidate answer in the rollout detected by our model. Each candidate is presented with the sentence that contains a candidate answer, alongside the corresponding problem statement of the rollout and the ground-truth solution.

The evaluation centers on two key questions, as illustrated in our user interface (Figure 10):

- **Q1:** This question tests whether the model correctly identifies a sentence containing a candidate answer. For example, a sentence such as “So the answer is 5?” qualifies as a valid candidate, whereas “Let’s try to solve this” does not.
- **Q2:** This question assesses whether the extracted candidate answer is an answer to the question, regardless of correctness. For example, if the question is asking some quantity that “can be represented as m/n , what is $m+n$ ”, then an extraction of “5” is valid in form, while “ $m=2, n=3$ ” is invalid, as it does not allow the math verifier to robustly evaluate its correctness.

918

919

920 Solve the following math problem step by step. You should first think about the reasoning process, and then provide the user with the answer. The reasoning process and answer are enclosed with `<think>` and `</answer>` tags respectively, i.e., `<think>` reasoning process here `</think>` `<answer>` answer here `</answer>`. The answer part should be a self-contained summary of the reasoning process i.e., a complete step-by-step solution. In the answer part, give the final answer to the problem in `\boxed{}`. The problem is: Ant Amelia starts on the number line at `0` and crawls in the following manner. For $n=1,2,3,$ Amelia chooses a time duration `$t_ns` and an increment `$x_ns` independently and uniformly at random from the interval `$(0,1)`. During the n th step of the process, Amelia moves `$x_ns` units in the positive direction, using up `$t_ns` minutes. If the total elapsed time has exceeded `3/5` minute during the n th step, she stops at the end of that step; otherwise, she continues with the next step, taking at most `33` steps in all. What is the denominator plus the numerator of the probability that Amelia's position when she stops will be greater than `15`?

921

922

923

924

925

926

927

928

929

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

Problem Statement:

Solve the following math problem step by step. You should first think about the reasoning process, and then provide the user with the answer. The reasoning process and answer are enclosed with `<think>` and `</answer>` tags respectively, i.e., `<think>` reasoning process here `</think>` `<answer>` answer here `</answer>`. The answer part should be a self-contained summary of the reasoning process i.e., a complete step-by-step solution. In the answer part, give the final answer to the problem in `\boxed{}`. The problem is: Ant Amelia starts on the number line at `0` and crawls in the following manner. For $n=1,2,3,$ Amelia chooses a time duration `$t_ns` and an increment `$x_ns` independently and uniformly at random from the interval `$(0,1)`. During the n th step of the process, Amelia moves `$x_ns` units in the positive direction, using up `$t_ns` minutes. If the total elapsed time has exceeded `3/5` minute during the n th step, she stops at the end of that step; otherwise, she continues with the next step, taking at most `33` steps in all. What is the denominator plus the numerator of the probability that Amelia's position when she stops will be greater than `15`?

Line 221

Line Content:

Wait, actually, let me check again. Because when she stops at step 2, the position is $x_1 + x_2$, which needs to be >1 . The probability that $x_1 + x_2 > 1$ is $1/2$, as we computed. But wait, actually, x_1 and x_2 are each uniform on $(0,1)$, so the probability that $x_1 + x_2 > 1$ is indeed $1/2$. Similarly, the probability that $x_1 + x_2 + x_3 > 1$ is $5/6$. Then, $P(A \cap B) = P(A) * P(B) = 1/2 * 1/2 = 1/4$, and $P(A^c \cap C) = P(A^c) * P(C) = 1/2 * 5/6 = 5/12$. So total probability is $1/4 + 5/12 = (3 + 5)/12 = 8/12 = 2/3$. Therefore, the probability is $2/3$, which is 2 over 3 . Therefore, denominator plus numerator is $3 + 2 = 5$. So the answer is 5 ? But wait, let me check if this is correct.

Extracted Candidate: 5

Ground Truth: 5.0

Q1: Does this line have a candidate?

Yes No

Q2: Is the extracted candidate valid, in the sense that it answers the question, regardless of correctness?

Yes No

Figure 10: The user interface for evaluating LLM extraction correctness for human participants.

Based on human verification of 100 randomly sampled rollouts (comprising 426 candidates), our model demonstrates high performance. It achieves 94.1% accuracy in identifying the correct position (Q1) and 94.0% accuracy in adhering to the target format (Q2).

D BREAKDOWN OF REFLECTIONS

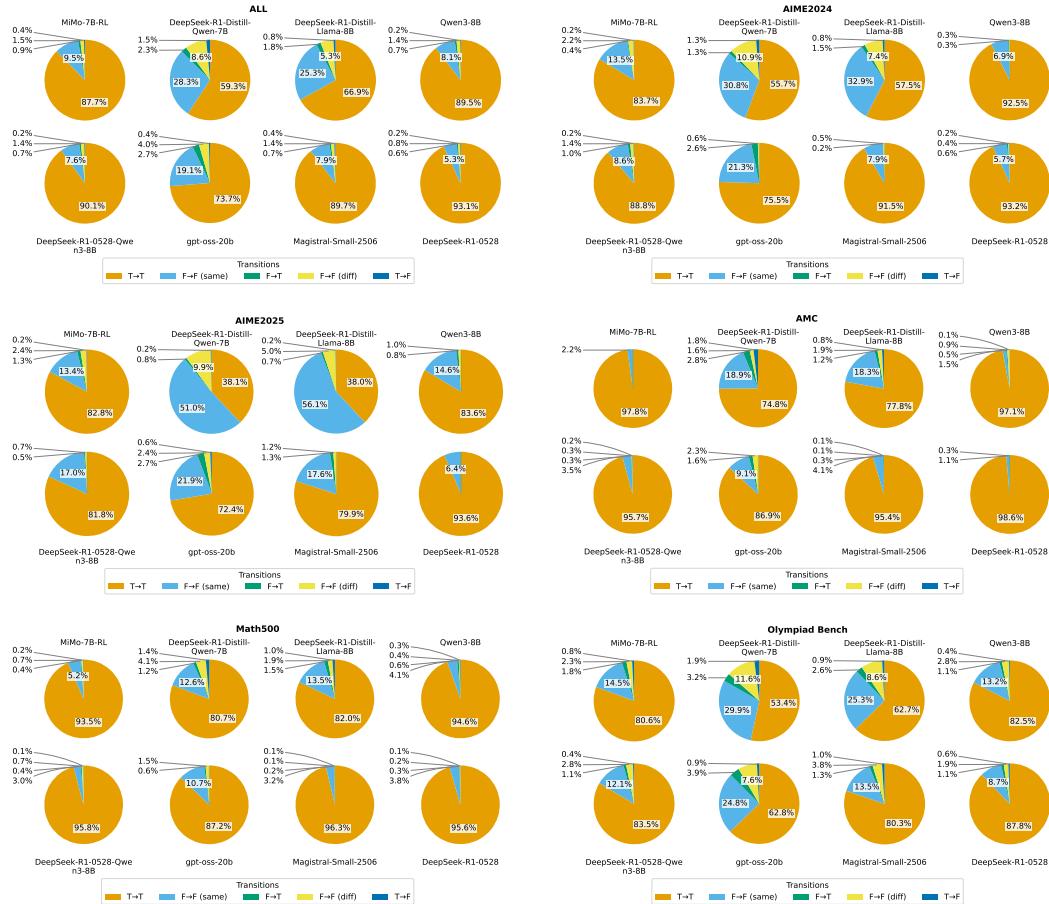


Figure 11: Reflection statistics of long CoTs of different models. Long CoTs are collected on AIME24 and AIME25 (32 rollouts per question), AMC (4 rollouts per question), Olympiad Bench, and Math500 (1 rollout per question).

E REFLECTION STATISTICS

Table 4: Number of candidate answers in rollouts across different reasoning models and datasets.

Model	AIME2024	AIME2025	Math500	Olympiad Bench	AMC
MiMo-7B-RL	3.64	3.52	5.79	4.69	4.64
DeepSeek-R1-Distill-Qwen-7B	3.50	3.12	2.47	3.76	2.96
DeepSeek-R1-Distill-Llama-8B	3.84	3.87	4.52	4.59	4.38
Qwen3-8B	4.27	3.72	6.94	6.10	5.77
DeepSeek-R1-0528-Qwen3-8B	3.55	3.17	6.32	4.74	4.57
gpt-oss-20B	3.22	3.12	3.01	3.48	3.31
Magistral-Small-2506	3.85	4.00	9.32	6.23	6.19
DeepSeek-R1-0528	3.65	2.96	6.34	4.85	5.03

1026 F EXAMPLE OF SFT DATA CURATION

1027

1028

1029

1030 Problem statement:

1031 Let ABC be a triangle inscribed in circle ω . Let the tangents to ω at B and C intersect at point D , and let \overline{AD} intersect ω at P . If $AB = 5$, $BC = 9$, and $AC = 10$, AP can be written as the form m/n , where m and n are relatively prime integers. Find $m + n$.

1032 Model Response:

1033 <think> Okay so I need to solve this geometry problem ...

1034 ...

1035 1st candidate Wait, that's a nice simplification! So AP is $100/13$. Therefore, $m = 100$, $n = 13$, which are coprime.

1036 Wait, that seems straightforward. Let me check my steps again

1037 to ensure no mistakes.

1038 ...

1039 2nd candidate Therefore, regardless of the method, AP is $100/13$, so $m + n = 100 + 13 = 113$

1040 But let me check once again with another approach to be

1041 thorough. Maybe using power of a point and some properties of

1042 triangle.

1043 ...

1044 3rd candidate Therefore, $AP = DA - DP = 100/13$. Therefore, the answer is $100 + 13 = 113$. Thus, I can be confident that the correct answer is 113.

1045 </think> To solve the problem ...

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

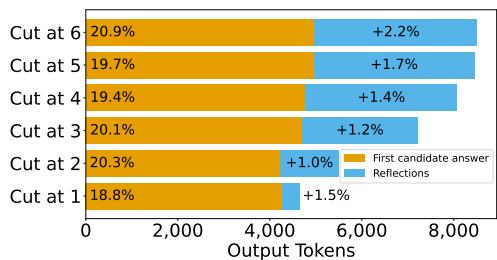
1067

1068

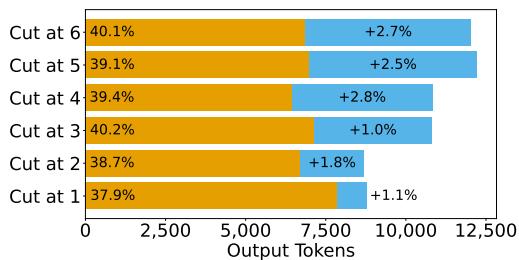
1069

1070

1

1080 **H TOKEN USAGE AND ACCURACY**
1081

(a) Llama3.1-8B-Instruct



(b) Qwen2.5-7B-Instruct

Figure 13: Token usage and accuracy after SFT using Deepseek-R1 rollouts, separated into two parts: the first candidate answers (before the first candidate) and the reflections (after the first candidate). In the figure, orange bars show tokens used before the first candidate answer, blue bars show tokens used for reflections, and performance is marked on each segment. Cut at i rows correspond to models trained with R1 rollouts truncated at different positions. Training data are controlled so that they have same number of training tokens. We can see that cut-at-6 is on average 3.8% better than cut-at-1 (averaging over two models), while the improvement in first answer accuracy contributes 2.65% improvement, and reflection contributes 1.15% improvement.

1100
1101 **I COMPARING DIVERSITY BETWEEN ROLLOUTS**
11021103 **Diversity Prompt**
1104

1105
1106 Please compare the two response and judge whether Response 2
1107 contains any new mathematical concept, or is simply redoing the
1108 previous calculations. Note, "new mathematical concept" can be
1109 any new information, imagine this is educational material to
1110 someone not familiar with math. Will Response 2 provide more
1111 information to them instead of simply repeating some previous
1112 calculation already in Response 1? In the last line of your
1113 response, please indicate "<YES>" or "<NO>".

1114 To assess whether datasets truncated at later positions include extra useful information, we compare
1115 two responses to the same problem: one from dataset cut-at- $\{i\}$ and the other from dataset cut-
1116 at- $\{i+1\}$, where the latter contains one more piece of reflection. Using an LLM (gpt-oss-120b),
1117 we evaluate whether this added text introduces any new mathematical concepts or problem-solving
1118 approaches. For all problems, we then compute the percentage for which cut-at- $\{i+1\}$ contributes
1119 additional concepts/approaches beyond cut-at- $\{i\}$. The statistics are shown in Table 7

1120 Table 7: Diversity analysis of longer rollouts. Percentage stands for the percentage of problems that
1121 the longer rollouts see additional problem-solving approaches.

	Percentage (%)
From cut-at-1 to cut-at-2	22.6
From cut-at-2 to cut-at-3	16.8
From cut-at-3 to cut-at-4	15.6
From cut-at-4 to cut-at-5	16.4
From cut-at-5 to cut-at-6	15.2

1122
1123 We provide two cases showing when the rollout with more reflections provides additional problem-
1124 solving approaches (Figure 14) and when the rollout with more reflections repeats previous calcula-
1125 tions (Figure 15).

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

Question. Even or odd sum of all natural numbers from 1 to 17?**Response in cut-at-1**

<think> Okay, so I need to figure out whether the sum of all natural numbers from 1 to 17 is even or odd. Hmm, let's start by recalling how to calculate the sum of numbers from 1 to a certain number. I think there's a formula for that. Let me remember... Oh yeah, it's the formula for the sum of an arithmetic series. The formula is $n(n + 1)/2$, where n is the last number in the sequence. So in this case, n is 17. Let me write that down:

$$\text{Sum} = 17 * (17 + 1) / 2$$

Calculating that, first I need to compute $17 + 1$, which is 18. Then multiply by 17: $17 * 18$. Hmm, what's 17 times 18? Let me do that step by step. $17 * 10$ is 170, $17 * 8$ is 136, so adding them together: $170 + 136 = 306$. Then divide by 2: $306 / 2 = 153$. So the total sum is 153. Now, I need to check if 153 is even or odd.

An even number is divisible by 2 without a remainder, and an odd number isn't. So, let me divide 153 by 2. 2 goes into 153 how many times? Well, $2 * 70$ is 140, subtract that from 153, you get 13. Then $2 * 6$ is 12, so that's 76 with a remainder of 1. So 153 divided by 2 is 76 with a remainder of 1, which means it's an odd number. Therefore, the sum is odd. ... Actually, that's correct, no further checking needed, I believe I have got the answer. Let's stop thinking and give a concrete answer based on the thoughts.

Response in cut-at-2

<think> Okay, so I need to figure out whether the sum of all natural numbers from 1 to 17 is even or odd. Hmm, let's start by recalling how to calculate the sum of numbers from 1 to a certain number. I think there's a formula for that. Let me remember... Oh yeah, it's the formula for the sum of an arithmetic series. The formula is $n(n + 1)/2$, where n is the last number in the sequence. So in this case, n is 17. Let me write that down:

$$\text{Sum} = 17 * (17 + 1) / 2$$

Calculating that, first I need to compute $17 + 1$, which is 18. Then multiply by 17: $17 * 18$. Hmm, what's 17 times 18? Let me do that step by step. $17 * 10$ is 170, $17 * 8$ is 136, so adding them together: $170 + 136 = 306$. Then divide by 2: $306 / 2 = 153$. So the total sum is 153. Now, I need to check if 153 is even or odd.

An even number is divisible by 2 without a remainder, and an odd number isn't. So, let me divide 153 by 2. 2 goes into 153 how many times? Well, $2 * 70$ is 140, subtract that from 153, you get 13. Then $2 * 6$ is 12, so that's 76 with a remainder of 1. So 153 divided by 2 is 76 with a remainder of 1, which means it's an odd number. Therefore, the sum is odd.

Wait, but maybe there's another way to figure this out without calculating the entire sum? Let me think. Since we're dealing with parity (even or odd), maybe we can find the number of odd numbers in the sequence from 1 to 17 and see if their count is odd or even. Because adding an even number of odd numbers would result in an even sum, and adding an odd number of odd numbers would result in an odd sum. But then again, even numbers don't affect the parity. So let's see.

From 1 to 17, how many odd numbers are there? Well, the numbers alternate between odd and even. Starting with 1 (odd), then 2 (even), 3 (odd), etc. Since 17 is odd, there are $(17 + 1)/2 = 9$ odd numbers? Wait, let me check. If you have numbers from 1 to n , the number of odd numbers is $\text{ceil}(n/2)$. For example, if n is even, like 16, then there are 8 odd numbers. If n is odd, like 17, then there are $(17 + 1)/2 = 9$ odd numbers. Let me count them manually to confirm. Numbers from 1 to 17:

1 (odd), 2 (even), 3 (odd), 4 (even), 5 (odd), 6 (even), 7 (odd), 8 (even), 9 (odd), 10 (even), 11 (odd), 12 (even), 13 (odd), 14 (even), 15 (odd), 16 (even), 17 (odd). Let me count how many odds: 1,3,5,7,9,11,13,15,17. That's 9 numbers. So 9 odd numbers.

Now, each odd number contributes 1 to the parity of the sum. So adding 9 odd numbers together. Since 9 is odd, the sum of 9 odd numbers is odd. Because adding two odd numbers gives even, adding another odd number makes it odd again. So odd + odd = even, even + odd = odd. So for 9 odds, since 9 is odd, the total sum would be odd. Therefore, the sum from 1 to 17 is odd.

... Actually, that's correct, no further checking needed, I believe I have got the answer. Let's stop thinking and give a concrete answer based on the thoughts.

Figure 14: One case where the additional reflection provides another problem-solving approach: To know the parity of the sum, one only needs to count the number of odd numbers.

1181

1182

1183

1184

1185

1186

1187

1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212

Question. What is the number? If I subtract a number from 12, I get 7. Calculate!

Response in cut-at-1

<think> Okay, let's see. The problem says, "If I subtract a number from 12, I get 7. Calculate!" So I need to find the number that when subtracted from 12 gives 7. Hmm, let me break this down.

First, let me translate the words into a mathematical equation. The problem states that subtracting a number from 12 results in 7. Let me denote the unknown number as x . So, the equation would be 12 minus x equals 7. That is, $12 - x = 7$.

Now, I need to solve for x . To do that, I can rearrange the equation. Let me think. If I subtract x from 12, the result is 7. So, to isolate x , I can subtract 7 from both sides? Wait, no. Let me recall how to solve linear equations.

Starting with $12 - x = 7$. If I want to solve for x , I can add x to both sides to get rid of the negative sign. So, adding x to both sides gives $12 = 7 + x$. Then, to get x by itself, I subtract 7 from both sides. That would be $12 - 7 = x$. Calculating that, 12 minus 7 is 5. So x equals 5. ... Actually, that's correct, no further checking needed, I believe I have got the answer. Let's stop thinking and give a concrete answer based on the thoughts.

Response in cut-at-2

<think> Okay, let's see. The problem says, "If I subtract a number from 12, I get 7. Calculate!" So I need to find the number that when subtracted from 12 gives 7. Hmm, let me break this down.

First, let me translate the words into a mathematical equation. The problem states that subtracting a number from 12 results in 7. Let me denote the unknown number as x . So, the equation would be $12 - x = 7$.

Now, I need to solve for x . To do that, I can rearrange the equation. Let me think. If I subtract x from 12, the result is 7. So, to isolate x , I can subtract 7 from both sides? Wait, no. Let me recall how to solve linear equations.

Starting with $12 - x = 7$. If I want to solve for x , I can add x to both sides to get rid of the negative sign. So, adding x to both sides gives $12 = 7 + x$. Then, to get x by itself, I subtract 7 from both sides. That would be $12 - 7 = x$. Calculating that, 12 minus 7 is 5. So x equals 5.

Wait, let me check that again. If I subtract 5 from 12, does that give 7? Let me verify. 12 minus 5 is indeed 7. Yeah, that works. So the number must be 5.

... Actually, that's correct, no further checking needed, I believe I have got the answer. Let's stop thinking and give a concrete answer based on the thoughts.

1213
 1214
 1215
 1216
 1217
 1218

Figure 15: One case where the additional reflection is simple repetition.

1219
 1220
 1221
 1222

J TESTS ON NON-MATHEMATICAL DOMAINS

1223
 1224
 1225
 1226
 1227
 1228

In addition to mathematical problems, in this section, we analyze the reasoning rollouts of Qwen3-8B, Qwen3-14B, and DeepSeek-R1-Distill-Llama-8B on three non-mathematical domains from MMLU: Formal Logic, High School Statistics, and Economics. Since MMLU are multiple-choice questions, we use the following prompt to extract the candidate answer positions.

1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

System Prompt for MMLU

You are given a text block that contains the original problem statement and a list of four choices to choose from to answer the problem, followed by a line-numbered model solution.

Your job is **NOT** to judge correctness or solve the problem. Instead, read the solution **line by line** and record every line where the model has implicitly or explicitly derived a candidate answer to the problem (it may not be correct, as long as it is an answer to the problem it counts). You need to fully understand what the problem asks for to notice the candidate answer. Only the thinking part of the model solution is provided for analysis.

Definitions

Candidate answer any explicit value or statement that (a) directly answers what the problem asks **or** (b) uniquely determines it with only a trivial final step (e.g. if the problem asks $N \bmod 1000$, then once you know N , taking $N \bmod 1000$ is immediate).

1242
 1243 *Candidate answer* is not intermediate components (like individual
 1244 addends when the question asks for their sum) unless the
 1245 problem explicitly asks for each component.
 1246 There is likely multiple candidates answers in the model solution,
 1247 and they are not necessarily the same as the model's final
 1248 answer. You should not look for candidate answers by matching
 1249 the model's final answer.
 1250
 1251 You can reason about the lines and decide whether they are
 1252 candidate answers. For the final response, you should follow
 1253 the format as below.
 1254
 1255 Final output format strict
 1256 1. For each qualifying line output a two-element tuple:
 1257 `line_number, choice`
 1258 `line_number` is an integer.
 1259 `choice` is one of A, B, C or D.
 1260 2. Collect the tuples in a Python list **in the order the lines
 1261 appear**.
 1262 3. The **last line** of your reply must be *only* that list,
 1263 so that `eval()` can parse it, for example, [(12, "A"), (27,
 1264 "C")]
 1265 4. Do **not** output anything after that list.
 1266
 1267

1268
 1269
 1270 We report results of Qwen3-8B, Qwen3-14B, and DeepSeek-R1-Distill-Llama-8B on three MMLU
 1271 tasks in Table 8, 9, and 10. We can see that on these non-mathematical tasks, the models exhibit a
 1272 similar pattern as in mathematical tasks, where the first answer accuracy closely matches the final
 1273 answer accuracy, and the relative position of the first answer is early.
 1274
 1275

Table 8: Performance of Qwen3-8B on three MMLU tasks.

Qwen3-8B	Formal Logic	High School Statistics	Economics
First answer accuracy	94.4	91.7	79.8
Final answer accuracy	96.8	94.4	80.7
First answer relative position	0.47	0.51	0.54

Table 9: Performance of Qwen3-14B on three MMLU tasks.

Qwen3-14B	Formal Logic	High School Statistics	Economics
First answer accuracy	96.8	93.1	78.9
Final answer accuracy	97.6	94.4	80.7
First answer relative position	0.51	0.56	0.59

Table 10: Performance of DeepSeek-R1-Distill-Llama-8B on three MMLU tasks.

DeepSeek-R1-Distill-Llama-8B	Formal Logic	High School Statistics	Economics
First answer accuracy	64.3	77.8	45.6
Final answer accuracy	66.7	79.2	48.2
First answer relative position	0.66	0.69	0.74

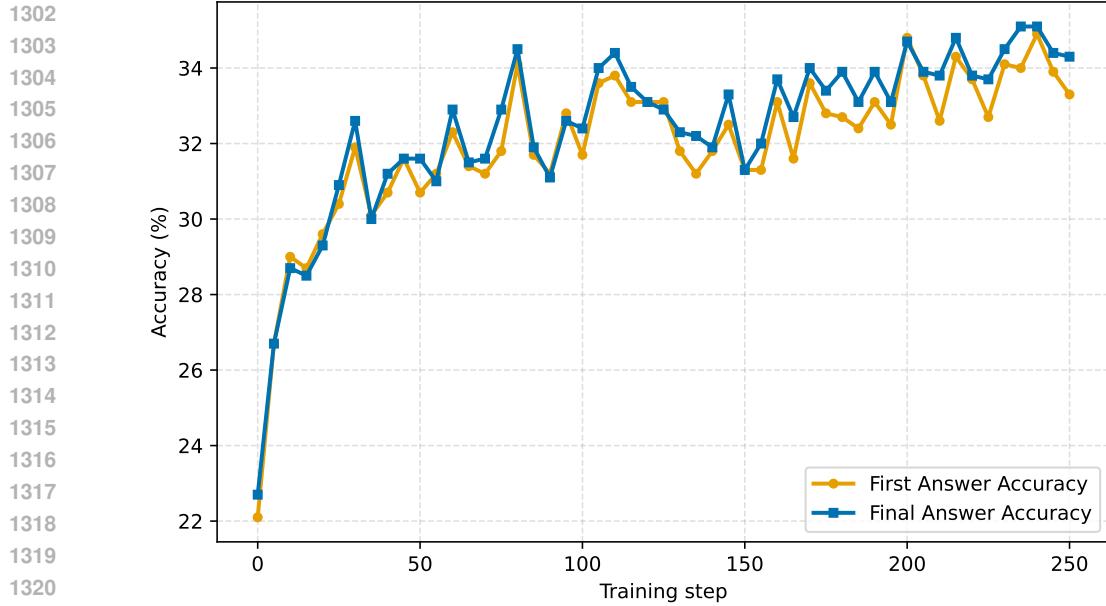
K DISCUSSION ON REINFORCEMENT LEARNING STAGE

1289
 1290 In Section 3, we studied the role of reflections in supervised-fine-tuning stage of the model training,
 1291 and briefly discussed the implications on RL stage. In this section, we further discuss how our
 1292 findings may apply to the reinforcement learning stage.
 1293
 1294

1295 We first conducted an RL training experiment, training a model initialized from Qwen2.5-Math-7B
 1296 with the DeepScaleR dataset and GRPO algorithm. We collect rollouts on the five mathematical

1296 datasets studied for every 5 steps of RL training and analyze their first answer accuracy and final
 1297 answer accuracy. As shown in Figure 16 the first-answer-accuracy closely matches the final answer
 1298 accuracy as training progresses. The echos with our findings that performance of a model is mainly
 1299 decided by its first answer accuracy, and models performance boost from RL also mainly originates
 1300 from the boost in first answer accuracy.

1301



1321

1322 Figure 16: Change of first answer accuracy and final accuracy as RL training progresses.
 1323

1324

L TRANSFERABILITY OF CAD AND QRC

1325

1326 In this section, we study the transferability of the proposed candidate answer detector (CAD) and
 1327 question-aware reflection controller (QRC) on other models and datasets. To do that, we use the
 1328 same CAD and QRC trained on the LLM-annotated Qwen3-8B rollouts on questions in M1-SFT
 1329 dataset, and apply to two models, Qwen3-14B and DeepSeek-R1-Distill-Llama-8B, on the five
 1330 mathematical datasets studied in the paper (Math500, AMC, Olympiad Bench, AIME2024, and
 1331 AIME2025), as well as three recent mathematical datasets at the time of writing: CMIMC, HMMT,
 1332 and BRUMO.

1333

1334 We report the results in Figure 17. We can see that the CAD and QRC can transfer well when
 1335 applied to Qwen3-14B on different datasets. Suggesting its robustness to model and problem for-
 1336 matting. Moreover, we also see that in some threshold configurations of CAD and QRC (e.g. 0.5
 1337 and 0), we achieve a 1.9% performance increase while reducing 13.5% of the tokens by early stop-
 1338 ping. Suggesting that excessive reflections may even have negative effects on the models reasoning
 1339 performance. Additionally, we show that CAD and QRC is not sensitive to problem difficulty by
 1340 showing their performance on AIME2025, which is considered the hardest among the five studied
 1341 datasets. Comparing Figure 17c with 17e and 17d with 17f, we can see that on AIME2025, CAD and
 1342 QRC do not introduce more performance drop than average, suggesting their robustness to problem
 1343 difficulties.

1344

1345

1346

1347

1348

1349

1350

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

1361

1362

1363

1364

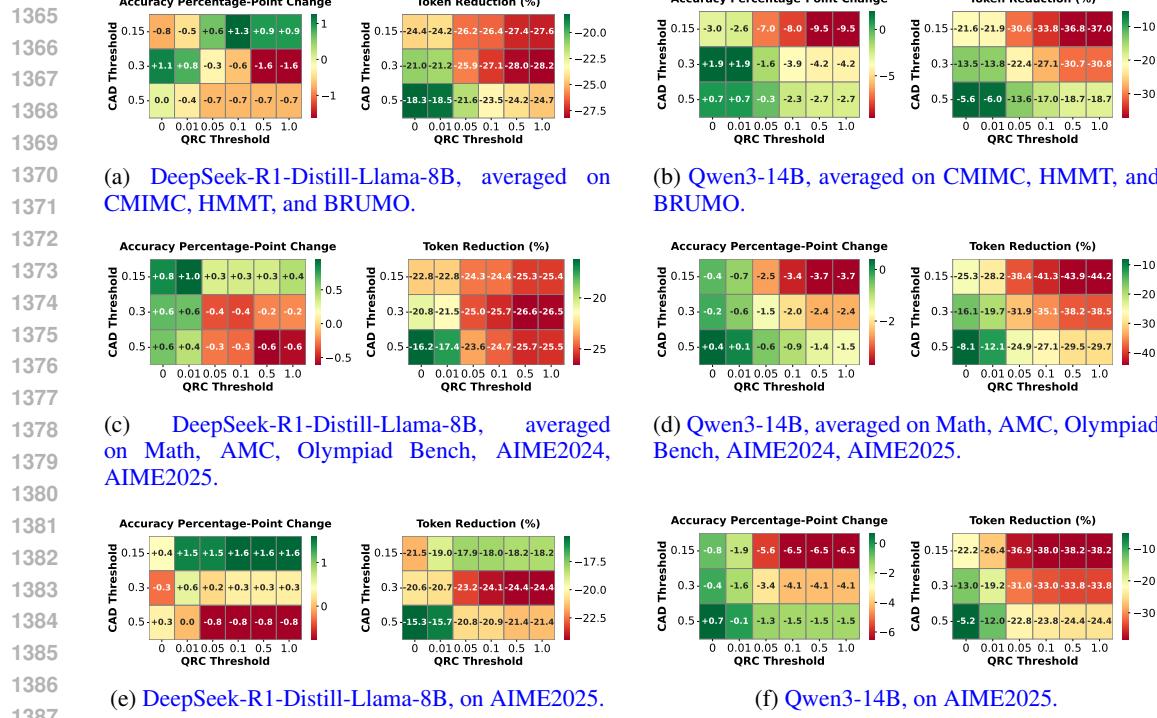


Figure 17: Transferability of CAD and QRC trained on Qwen3-8B and applied to DeepSeek-R1-Distill-Llama-8B and Qwen3-14B.

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403