

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 SELF-DUAL: UNIFYING NATURAL LANGUAGE AND PROGRAMMATIC THINKING FOR ENHANCED MATHE- MATICAL REASONING IN LLMs

Anonymous authors

Paper under double-blind review

## ABSTRACT

Large language models (LLMs) have made significant progress in mathematical reasoning. However, the methods that rely on a single reasoning paradigm exhibit clear limitations. This has motivated recent studies to combine multiple paradigms, but existing studies often fail to systematically exploit their complementary strengths. In this study, we first examine the complementary relationship between natural language (NL) and programmatic language (PL) reasoning, and show that their integration leads to consistent improvements in mathematical reasoning performance. Building on this analysis, we introduce Self-Dual, a framework that unifies the two paradigms within a single inference process by generating complementary reasoning trajectories and combining them through structured self-reflection. Beyond inference, we extend this principle to training by adopting the Self-Dual data format to construct complementary reasoning datasets and evaluate its effectiveness in model training. We conduct comprehensive evaluations of Self-Dual in both inference and training contexts. During inference, Self-Dual consistently surpasses NL-only, PL-only, and hybrid baseline methods across multiple benchmarks. DeepSeek-V3-0324 integrated with Self-Dual attains 47.8% accuracy on the AIME25 dataset, outperforming Chain-of-Thought (CoT) at 39.2% and Program-Aided Language (PAL) at 35.6%. In the training experiments, we apply the Self-Dual framework to further train Qwen2.5-7B-Instruct with only 7.5K MATH samples and construct Qwen2.5-7B-SD. The new model improves performance on MATH500 by more than 4% over the base model Qwen2.5-7B-Instruct. It also surpasses Qwen2.5-Math-7B-Instruct on AIME25. These results demonstrate that the Self-Dual framework effectively exploits complementary reasoning paradigms and substantially enhances the mathematical reasoning ability of large language models in both inference and training.

## 1 INTRODUCTION

Large Language Models (LLMs) (Achiam et al., 2024; Guo et al., 2025; Qwen et al., 2025) make remarkable strides across a diverse array of tasks through the application of prompting techniques. Nevertheless, complex reasoning tasks (Xie et al., 2024; Sprague et al., 2024) such as mathematical reasoning continue to present significant challenges. NL-based prompting methods, such as Chain-of-Thought (CoT), utilize natural language to decompose complex reasoning tasks into multiple intermediate steps and solve them step by step (Wei et al., 2022; Zhou et al., 2023; Li et al., 2023). These methods offer high interpretability but are prone to calculation and logical errors (Madaan & Yazdanbakhsh, 2022; Nogueira et al., 2021; Qian et al., 2022). In contrast, PL-based prompting methods like Program-aided Language (PAL) (Gao et al., 2023b) generate programs as intermediate reasoning steps and offload the solution steps to a runtime such as a Python interpreter. They are more logically robust and computationally accurate (Chen et al., 2023; Gou et al., 2024b). However, such methods often struggle to represent certain steps for solving natural language problems in programming language.

Prior work on combining NL-based and PL-based reasoning can be grouped into two categories. At inference time, methods such as Automatic (Zhao et al., 2023) select the final answer from CoT and PAL outputs, but are sensitive to candidate ordering and often unstable. Other methods,

like CRITIC (Gou et al., 2024a), refine programs through NL reflection, yet remain limited by the expressiveness of PL. At training time, approaches such as MathCoder (Wang et al., 2023) construct datasets by deriving formulas and code with PL reasoning, while SAAS (Kim et al., 2024) and Qwen2.5-Math (Yang et al., 2024a) mix NL and PL data for joint training. CoR (Yu et al., 2025b) adopts progressive paradigm training, where the model is first trained on NL-reasoning data and then on PL-reasoning data. All these methods empirically combine NL-based reasoning and PL-based reasoning, but they overlook the central principle of complementarity between them, leaving much of their combined potential unexplored.

In this paper, our observation confirms that NL reasoning and PL reasoning are inherently complementary: their union consistently outperforms either alone, while their intersection reflects convergence across reasoning attempts, as shown in Figure 1. This suggests that hybrid reasoning should explicitly leverage complementarity, rather than treat NL and PL as independent or sequential signals. Building on this insight, we propose Self-Dual, a unified framework that exploits the complementarity of NL and PL reasoning to enhance inference, while also serving as a data construction format to support model training. At inference time, Self-Dual produces two complementary reasoning paths within a single forward pass: one in NL and one in programmatic form. These paths are then compared and integrated through a structured self-reflection process consisting of three steps: Look Back, Decomposition, and Resolution. At training time, Self-Dual generates complementary trajectories that are used as cold-start data. These data are then integrated into Group Relative Policy Optimization (GRPO) (Shao et al., 2024) with a tailored reward function to assess whether the Self-Dual format can improve model performance in mathematical reasoning.

For inference-time experiments, we evaluate the empirical performance of Self-Dual across multiple series of LLMs, including DeepSeek-V3-0324 (DeepSeek-AI, 2024) and Gemma 3 (Team, 2025) variants (4B, 12B, 27B), on diverse mathematical reasoning benchmarks. Our findings show that Self-Dual consistently surpasses prior techniques without supplementary data or training. For example, DeepSeek-V3-0324 with Self-Dual BAND achieves 4% improvement in accuracy on the MATH500 benchmark and 16.66% gain on AIME25. For training-time experiments, we use Qwen2.5-7B-Instruct (Qwen et al., 2025) as the base model and train it with Self-Dual data, obtaining Qwen2.5-7B-SD. On AIME25, Qwen2.5-7B-SD surpasses Qwen2.5-Math-7B-Instruct (Yang et al., 2024a), which is trained with large-scale hybrid data and GRPO. This result shows that the Self-Dual format can deliver competitive improvements in mathematical reasoning with much smaller training data.

## 2 RELATED WORK

### 2.1 ENHANCING MATHEMATICAL REASONING

There are two main techniques for enhancing the mathematical reasoning capabilities of LLMs: prompting methods and fine-tuning-based methods. The in-context few-shot learning through prompting is simple and broadly applicable, like popular CoT (Wei et al., 2022) and PAL (Gao et al., 2023b). Automatic-CoT (Zhang et al., 2023), Many-Shot (Agarwal et al., 2024) and Synthetic Prompting (Sprague et al., 2024) explore improving model performance by enhancing the quality or quantity of few-shot examples. ToT (Yao et al., 2023), GoT (Besta et al., 2024) and BoT (Yang et al., 2024b) design different ways of structuring thoughts to further enhance LLMs performance. TIR (Gou et al., 2024b), CRITIC (Gou et al., 2024a) and Recursive-Introspection (Qu et al., 2024) perform multi-round self-evaluation based on previous results to iteratively improve the reasoning process. The most related to our work are Mathprompter (Kim et al., 2023a) and Automatic (Zhao et al., 2023), which use simple selection or majority voting to select the results of PAL and CoT. Unlike these methods, our approach decomposes the complementary reasoning paths to extract their respective strengths and synthesizes a new solution during the refine stage.

Beyond prompt-based approaches, fine-tuning improves mathematical reasoning by training on curated datasets (Tong et al., 2024a; Zhou et al., 2024; Luo et al., 2025a) or large-scale synthetic data, such as MetaMath (Yu et al., 2024) and DeepSeekMath (Shao et al., 2024). Building on this line, recent work explores multi-paradigm reasoning by mixing natural language, code, and symbolic forms (Yu et al., 2025a; Zheng et al., 2025) or by integrating reinforcement learning and external tools (Luo et al., 2025a; Gou et al., 2024b). Yet these methods treat NL-based and PL-based rea-

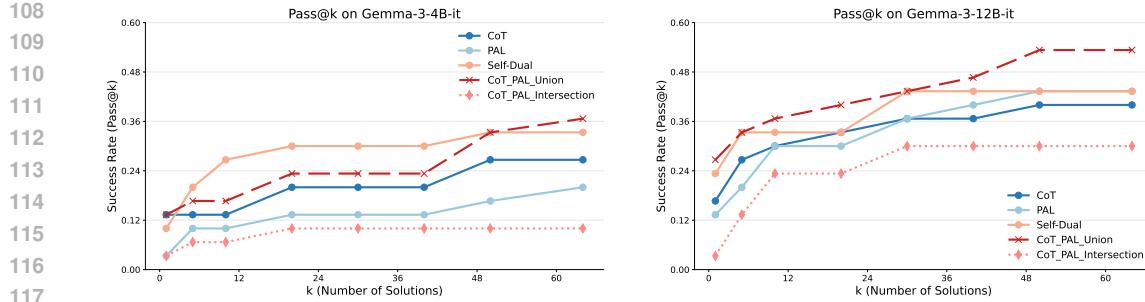


Figure 1: Pass@k curves of Gemma3-4B-it and Gemma3-12B-it on CoT, PAL, Self-Dual, as well as the union and intersection of CoT and PAL, evaluated on the AIME25 dataset. Setup is given in Section 4.1.1.

sioning as simple combinations, overlooking their complementarity. Our work instead unifies both paradigms within a single inference pass and reconciles them through self-reflection.

## 2.2 SELF-REFINE

Self-Refine refers to the ability of LLMs to improve or revise their previous responses based on feedback (Madaan et al., 2023; Shinn et al., 2023). Automated feedback has primarily been used for training-time correction (Yan et al., 2023) or for guiding LLMs during inference through prompting (Weng et al., 2023). One crucial aspect of self-refinement is that the source of feedback provides reliable and effective additional information. A vast amount of research leverage external feedback for optimization, including external knowledge sources (Gao et al., 2023a), external tools (Chern et al., 2024), program executors (Chen et al., 2024), symbolic solvers (Pan et al., 2023), or even trained models (Paul et al., 2024). Some approaches, such as Self-Refine (Madaan et al., 2023) and RCI (Kim et al., 2023b), further optimize responses through multiple iterations by leveraging the model’s inherent reflection capabilities. While the primary goal of our method is to generate complementary solutions that inspire reflection, external feedback from the Python interpreter is also integrated into the refine stage as a component of the PL-based approach.

## 3 METHODOLOGY

### 3.1 COMPLEMENTARITY BETWEEN REASONING MODES

NL-based methods are adept at capturing semantic and intuitive reasoning, while PL-based methods primarily emphasize symbolic and structured reasoning. Intuitively, integrating these two paradigms is expected to enhance reasoning capabilities, much like the human cognitive process that combines intuition with logic. Although several recent studies (Yu et al., 2025b) have explored the integration of NL and PL to improve the reasoning ability of LLMs, to the best of our knowledge, a systematic understanding of their complementary interplay is still missing.

In this work, we explicitly examine the notion of complementarity and conceptualize it as a central principle underlying the effectiveness of hybrid NL-PL reasoning. Concretely, we position CoT as the representative paradigm of NL-based reasoning and PAL as the representative paradigm of PL-based reasoning. By analyzing these two paradigms, we provide empirical evidence supporting their inherent complementarity.

To evaluate performance, we adopt the Pass@K metric, which measures the probability of obtaining at least one correct solution when sampling  $K$  diverse outputs from a model. Formally, let  $S_i$  denote the set of solutions generated by method  $i$  ( $i \in \text{CoT, PAL}$ ), and let  $\mathcal{C}$  denote the set of correct solutions. Then Pass@K is defined as:

$$\text{Pass@K}(S_i) = \mathbb{P}[S_i \cap \mathcal{C} \neq \emptyset]. \quad (1)$$

In addition to evaluating each method individually, we also analyze the union and intersection of the solutions produced by CoT and PAL. The union reflects the expanded solution space when

162 both paradigms are combined, while the intersection reveals their agreement. Examining these sets  
 163 allows us to better understand the complementary and consistent aspects of NL-based and PL-based  
 164 reasoning. They can be expressed as:

$$S_{\cup} = S_{\text{CoT}} \cup S_{\text{PAL}}, \quad S_{\cap} = S_{\text{CoT}} \cap S_{\text{PAL}}. \quad (2)$$

168 As illustrated in Fig. 1, our key observations are as follows.  
 169

170 **Union as upper bound.**  $\text{Pass}@K(S_{\cup})$  increases steadily with  $K$  and consistently surpasses both  
 171 CoT and PAL alone. This demonstrates that NL-based and PL-based reasoning are complemen-  
 172 tary. Since their outputs are not fully overlapping, the union expands the effective solution set.  
 173  $\text{Pass}@K(S_{\cup})$  therefore serves as an upper bound for hybrid reasoning.

174 **Intersection as consistency.**  $\text{Pass}@K(S_{\cap})$  increases and gradually converges as  $K$  grows. This  
 175 behavior indicates that CoT and PAL exhibit increasing consistency in their outputs. At the same  
 176 time,  $\text{Pass}@K(S_{\cap}) \leq \min(\text{Pass}@K(S_{\text{CoT}}), \text{Pass}@K(S_{\text{PAL}}))$ , which confirms that their solution  
 177 sets are not fully overlapping. This motivates the use of the symmetric difference to characterize  
 178 their non-shared components:

$$S_{\Delta} = (S_{\text{CoT}} \setminus S_{\text{PAL}}) \cup (S_{\text{PAL}} \setminus S_{\text{CoT}}) \quad (3)$$

181 **Symmetric difference as complementarity.** The non-empty symmetric difference further supports  
 182 the claim that CoT and PAL capture distinct reasoning capacities. Their complementarity provides  
 183 opportunities for hybrid methods to exploit.

184 In a nutshell, the union highlights the upper bound of hybrid reasoning, the intersection reflects  
 185 growing consistency, and the symmetric difference confirms complementarity. Together, these ob-  
 186 servations provide a foundation for designing methods that explicitly exploit NL–PL synergy.  
 187

### 188 3.2 INFERENCE-TIME COMPLEMENTARITY 189

190 To exploit the complementary strengths of NL and PL reasoning, we propose Self-Dual, a hybrid  
 191 framework that integrates both paradigms within a single inference. The key idea is to generate dual  
 192 reasoning trajectories in one forward pass and then combine them through structured self-reflection.  
 193 This design preserves the diversity of multiple reasoning paths while maintaining the efficiency of  
 194 single-shot inference.

195 Self-Dual operates in two stages: Dual-Path Generation and Refinement. In the first stage, the  
 196 model produces an NL-based reasoning path  $R_{\text{NL}} = \text{LLM}(Q; \theta_{\text{NL}})$ , followed by a PL-based path  
 197  $R_{\text{PL}} = \text{LLM}(Q, R_{\text{NL}}; \theta_{\text{PL}})$  with  $R_{\text{NL}}$  as context. This sequential generation improves consistency  
 198 between the two paradigms, as discussed in Section 4.1.3. During the generation of  $R_{\text{PL}}$ , external  
 199 feedback can be incorporated via stop tokens “`output`”, allowing partial code execution and  
 200 verification  $f(R_{\text{PL}}) \rightarrow \text{Execution Feedback}$  to strengthen the reliability of the PL trajectory.

201 In the second stage, Refinement, the two paths are compared and integrated in three structured steps.  
 202 Look Back prompts the model to revisit the original problem statement  $Q$ , extract key information,  
 203 and mitigate attention drift. Decomposition contrasts  $R_{\text{NL}}$  and  $R_{\text{PL}}$  to identify agreements and con-  
 204 flicts between the two paths. Resolution then consolidates  $R_{\text{NL}}$ ,  $R_{\text{PL}}$ , and the execution feedback  
 205  $f(R_{\text{PL}})$  into a refined trajectory  $R$ , from which the final answer  $A = \text{Answer}(R)$  is derived and  
 206 presented as “boxed{A}”.

### 208 3.3 TRAINING-TIME COMPLEMENTARITY 209

210 At training time, we use the Self-Dual framework as a data generator to construct complementary  
 211 reasoning trajectories. These trajectories serve as cold-start data for reinforcement learning and  
 212 allow us to examine whether the Self-Dual format can enhance mathematical reasoning through  
 213 data construction. We adopt standard GRPO for training and design a reward function that preserves  
 214 the complementary reasoning pattern of Self-Dual:

$$R = \delta \cdot (\text{format reward}), \quad (4)$$

216 where  $\delta \in \{0, 1\}$  is a binary coefficient, and `format_reward` is provided by an external reward model.  
 217

218 The value of  $\delta$  is determined by two conditions. The first condition enforces a complementarity  
 219 constraint: the model must output both an NL-based and a PL-based answer, preserving the reasoning  
 220 format of Self-Dual during training. We do not evaluate the correctness of these intermediate  
 221 answers, so the model can explore diverse reasoning paths. The second condition is the correctness  
 222 of the final answer  $A$ , which determines the success of the solution. The coefficient  $\delta$  is set to 1  
 223 only when both conditions are satisfied. The format reward term provides a fine-grained evaluation  
 224 of reasoning quality, covering the completeness of the three-stage Self-Dual process, the correctness  
 225 of the reasoning trajectory, and the correctness of the final answer.  
 226

## 227 4 EXPERIMENTS

### 229 4.1 INFERENCE-TIME

#### 231 4.1.1 SETUP

232 **Benchmarks.** We evaluate our approach on four widely-used benchmarks in the field of mathematical  
 233 reasoning: SVAMP (Patel et al., 2021), GSM8K (Cobbe et al., 2021), MATH500 (Lightman  
 234 et al., 2024) and AIME25. Their difficulty increases progressively in the given order. Please see  
 235 Table 7 for details in Appendix.

236 **LLMs.** We conduct experiments using DeepSeek-V3-0324 (DeepSeek-AI, 2024), as well as the  
 237 Gemma-3 Instruct model (Team, 2025) across three parameter scales: 4B, 12B, and 27B. To ensure  
 238 a fair comparison, we set the temperature of all LLMs to 0.7 and the maximum token limit to 8192  
 239 for all experiments.  
 240

241 **Baselines.** We group the baselines into three categories. For NL-based methods, we consider  
 242 CoT (Wei et al., 2022), which solves problems step by step in natural language, and Reflexion (Shinn  
 243 et al., 2023), which follows the Response–Evaluation–Revision paradigm. For PL-based meth-  
 244 ods, we include PAL (Gao et al., 2023b), which generates executable programs. For hybrid-based  
 245 methods, we evaluate Automatic (Zhang et al., 2023), a simple selection between CoT and PAL,  
 246 Self-Dual-Automatic, which integrates independently executed CoT and PAL under the Self-Dual  
 247 framework, CRITIC (Gou et al., 2024a), which iteratively refines programs using natural language  
 248 feedback, and TIR (Gou et al., 2024b), which generates a rationale before producing and refining  
 249 code. We set the maximum number of iterations to  $n = 4$  for iterative methods, while for single-pass  
 250 methods such as CoT and Self-Dual we report the average over three runs.  
 251

#### 251 4.1.2 MAIN RESULTS

253 In Table 1, we report both the accuracy and the number of API/LLM calls (#Call), which serves as  
 254 an approximate measure of computational cost. Table 6 shows remaining results of GSM8K and  
 255 SVAMP in Appendix. Hybrid-based methods consistently outperform single-paradigm methods  
 256 across both benchmarks. While NL-based methods such as CoT and Reflexion capture semantic  
 257 reasoning effectively, and PL-based methods such as PAL and TIR emphasize symbolic reasoning,  
 258 their performance remains limited when used in isolation. By contrast, hybrid-based approaches  
 259 achieve clear improvements, highlighting the intrinsic complementarity between NL-based and PL-  
 260 based paradigms.  
 261

262 The proposed Self-Dual framework establishes better performance across all model scales, surpass-  
 263 ing both automatic switching methods and their tool-free variants. Notably, Self-Dual achieves  
 264 47.78% on AIME25 with DeepSeek-V3-0324 and 90.04% on MATH500, demonstrating that explic-  
 265 itly generating and reconciling dual reasoning paths yields stronger gains than heuristic integration.  
 266 These improvements are particularly pronounced for smaller models, where Self-Dual narrows the  
 267 gap with much larger baselines. Tool feedback plays a crucial role in amplifying the benefits of  
 268 programmatic reasoning. As shown by the “w/o tool” settings, removing Python interpreter feed-  
 269 back substantially degrades performance for PAL, CRITIC, and Self-Dual, confirming that external  
 270 verifiability strengthens the PL-based reasoning pathway. Importantly, even without tool feedback,  
 271 Self-Dual still maintains competitive accuracy, underscoring the robustness of its complementary  
 272 reasoning design.  
 273

270 Table 1: Solve rates on AIME25 and MATH500. The "w/o tool" setting indicates that the method  
 271 described in the previous line does not utilize the execution results from the Python interpreter as  
 272 feedback. D-V3 refers to DeepSeek-V3-0324, while G-4B, G-12B, and G-27B denote different  
 273 sizes of the Gemma-3 family. SD-Auto denotes the Self-Dual-Automatic method. The best results  
 274 are highlighted in bold, and the second-best are underlined.

Methods	AIME25				MATH500				
	D-V3	G-4B	G-12B	G-27B	D-V3	G-4B	G-12B	G-27B	#Call
<i>NL-based methods</i>									
CoT	27.78	8.9	17.78	20	87.67	62.2	<u>79.87</u>	85.53	1
Reflextion	<u>36.67</u>	<u>10</u>	20	26.67	<u>88</u>	64.8	79.6	85.2	3
<i>PL-based methods</i>									
PAL	16.67	0	13.33	16.67	68	35.2	58.4	68.6	1
<i>Hybrid-based methods</i>									
TIR	13.33	0	6.67	23.33	81	43.4	63	77.2	1-5
w/o tool	13.33	0	3.33	10	76	18.8	32.8	45.8	1
CRITIC	13.33	3.33	10	10	68.4	39.8	60.6	68.6	9
w/o tool	16.67	3.33	10	16.67	65.6	37.6	60.6	65.2	9
Automatic	40	6.67	<u>23.33</u>	26.67	81.8	60.2	75.8	82.6	3
SD-Auto	40	<u>10</u>	<u>23.33</u>	26.67	<u>88</u>	<u>67</u>	82	85.6	3
Self-Dual	<b>47.78</b>	<b>15.56</b>	<b>24.44</b>	<b>30</b>	<b>90.04</b>	<b>70.33</b>	<b>81.6</b>	<b>86.33</b>	2
w/o tool	26.67	10	16.67	16.67	87.4	68	79.6	<u>86</u>	1

#### 4.1.3 ANALYSIS

296 To further investigate the proposed framework in depth, we randomly sample 100 examples from  
 297 MATH500. All sampled problems are kept consistent across methods to ensure a fair comparison.  
 298 A comparative analysis of all methods is shown in Table 2, based on the consistency of the two  
 299 initial reasoning paths and the correctness of the final outputs in the refine stage. We define four  
 300 main categories: *Concordant Correctness (CC)*, *Concordant Error (CE)*, *Discordant Correctness*  
 301 (*DC*), and *Discordant Error (DE)*. *CC* denotes cases where both NL- and PL-based paths give the  
 302 same correct answer, and refinement preserves it. *CE* corresponds to both paths producing the same  
 303 wrong answer, which refinement follows. When the two paths are different, the case is *Discordant*.  
 304 *DC* indicates that refinement yields the correct answer, while *DE* means it remains incorrect. We  
 305 further distinguish *DC<sub>NL</sub>* and *DE<sub>NL</sub>* when the NL path is correct, and *DC<sub>PL</sub>* and *DE<sub>PL</sub>* when  
 306 the PL path is correct. *Badcode* refers to PL programs that fail to execute, *FN (False Negative)* to  
 307 correct answers mismatched with the ground truth, and *No2Paths* to cases where dual solutions are  
 308 not generated.

309 Table 2: Results of Self-Dual, Self-Dual-Auto, Auto, and DCT on MATH500, categorized by the  
 310 consistency of the initial reasoning paths and the correctness of the final outputs in the refine stage.

Models	Methods	CC	CE	DC	DE	Badcode	FN	No2Paths
DeepSeek-V3-0324	Self-Dual	77%	4%	6%	4%	4%	1%	4%
	Self-Dual-Auto	67%	1%	8%	9%	15%	0%	-
	Auto	67%	1%	4%	13%	15%	0%	-
Gemma-3-12B-it	Self-Dual	67%	6%	8%	5%	4%	1%	9%
	Self-Dual-Auto	53%	4%	16%	13%	12%	2%	-
	Auto	53%	4%	11%	18%	12%	2%	-
	DCT	80%	9%	4%	1%	-	2%	4%

322 **Complementary Analysis.** In this paper, we regard the divergence between the NL-based and  
 323 PL-based methods in the initial stage as the main concrete manifestation of their underlying  
 324 complementarity. As shown in Table 2, in Gemma-3-12B-it, the proportion of Discordant outputs in

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

Table 3: Confidence of NL-based and PL-based methods when a divergence occurs in the initial stage.

Models	Methods	$C_{NL}$	$C_{PL}$
DeepSeek-V3-0324	Self-Dual	50%	30%
	Self-Dual-Auto	41.18%	27.65%
Gemma-3-12B-it	Self-Dual	61.54%	15.38%
	Self-Dual-Auto	51.72%	20.69%

Table 4: Selection accuracy of each method in the refine stage.

Models	Methods	SCR	DCR
DeepSeek-V3-0324	Self-Dual	95.4%	60%
	Self-Dual-Auto	89.29%	47.06%
	Auto	84.52%	23.53%
Gemma-3-12B-it	Self-Dual	93.75%	61.54%
	Self-Dual-Auto	84.15%	55.17%
	Auto	78.05%	37.93%

Self-Dual is higher than that of DCT. This indicates that employing NL and PL reasoning separately does indeed increase the likelihood of generating divergent outputs.

The comparison between Self-Dual and Self-Dual-Auto can be viewed as an evaluation of generating complementary reasoning paths in in-context versus out-of-context settings, respectively. The results demonstrate that the in-context setting reduces the occurrence of divergence and improves the accuracy of selection. Moreover, the divergence observed between the two methods suggests that LLMs may exhibit an inherent "duality" in their reasoning behaviors.

**Analysis of Divergent Reasoning Path.** To better analyze the reliability of PL-based and NL-based methods when they produce divergent outputs, we propose a new set of evaluation metrics for comparison. The confidence score  $C$  is computed using the following formula:

$$C_x = \frac{DC_x + DE_x}{DC + DE}, x \in [NL, PL]. \quad (5)$$

Table 3 shows the main results of confidence score. From a methodological perspective, Self-Dual demonstrates higher reliability than Self-Dual-Auto when divergences occur. For example, in the Deepseek model, Self-Dual achieves 80% confidence compared to Self-Dual-Auto's 69.45%. This indicates that the in-context setting in Self-Dual simultaneously enhances the confidence of both reasoning methods while reducing the likelihood of divergence (from 81% to 68%). Besides, both models consistently show that the NL-based method exhibits higher reliability; however, the contributions of the PL-based method remain significant and should not be overlooked.

**Refinement Strategy Evaluation** In this section, we propose new metrics to evaluate the selection accuracy of each method during the refine stage. Since CE indicates that both solutions are incorrect, we do not consider this case when evaluating the impact on the second stage. The *Selection Correctness Ratio (SCR)* measures the proportion of correct selections made among avoidable errors. It is defined by the following formula:

$$SCR = \frac{CC + DC}{CC + DC + DE}. \quad (6)$$

The *Discordant Correctness Ratio (DCR)* is calculated using the following Equation 7 and is designed to measure the proportion of making the correct selection when the two methods produce divergent answers. It serves as a reliable indicator of the effectiveness of the refine stage.

$$DCR = \frac{DC}{DC + DE}. \quad (7)$$

As shown in Table 4, it can be concluded that Self-Dual performs the best, followed by Self-Dual-Auto, while Auto shows the weakest performance based on the overall data. The difference between Self-Dual-Auto and Auto can be entirely attributed to the design of the refine stage, demonstrating that our proposed refinement strategy is significantly superior to the simple selection mechanism used in Auto. The results on Self-Dual and Self-Dual-Auto also merit further discussion. This may result from the in-context setting, which enhances consistent correctness and supports better decision-making when divergences arise.

378 Table 5: Main training results on MATH500, AMC23, AIME24 and AIME25. Bold numbers indicate  
 379 the best accuracy, and underline indicates the second best.  
 380

381 <b>Model</b>	382 <b>Backbone Model</b>	383 <b>RLVR</b>	384 <b>DataSize</b>	385 <b>Training Format</b>
386      Qwen2.5-7B-Ins	387      Qwen2.5-7B-Base	388      No	389      -	390      -
388 <i>w/ SFT</i>	389      Qwen2.5-7B-Ins	390      No	391      3.1k	392      Self-Dual
390 <i>w/ SFT+GRPO</i>	391      Qwen2.5-7B-Ins	392      Yes	393      7.5k	394      Self-Dual
391      Qwen2.5-Math-7B-Ins	392      Qwen2.5-Math-7B-Base	393      Yes	394      3000k	395      CoT & TIR
<b>MATH500</b>		<b>AMC23</b>	<b>AIME24</b>	<b>AIME25</b>
<i>Pass@1</i>		<i>Best@64</i>	<i>Best@64</i>	<i>Best@64</i>
395      Qwen2.5-7B-Ins	396      72.2	397 <b>65.0</b>	398      20.0	399      16.67
399 <i>w/ SFT</i>	400      74.4	401 <u>62.5</u>	402 <u>26.67</u>	403      16.67
403 <i>w/ SFT+GRPO</i>	404 <u>76.8</u>	405 <u>62.5</u>	406 <b>30.0</b>	407 <b>26.67</b>
407      Qwen2.5-Math-7B-Ins	408 <b>77.8</b>	409 <u>62.5</u>	410      20.0	411 <u>20.0</u>

393  
 394  
 395      **4.2 TRAINING-TIME**

396  
 397      **4.2.1 EXPERIMENTAL SETUP**

398  
 399      **Training Setup.** We adopt Qwen2.5-7B-Instruct as the backbone model, which is one of the most  
 400      widely used open-source instruction models. To evaluate the quality of model outputs, we employ  
 401      Qwen3-32B as the reward model, responsible for computing the format reward. All training data  
 402      are drawn from the 7.5K training set of the MATH benchmark.

403      **SFT Data Construction.** We follow the data generation paradigm established in prior work (Guan  
 404      et al., 2025). Specifically, we use DeepSeek-V3.1 as the generator and score the outputs with  
 405      Qwen3-32B under the same prompts used in the RL stage. Starting with 3.5K samples as the  
 406      candidate pool, a sample is included in the SFT set if it satisfies two conditions: the final answer is  
 407      correct and the format reward exceeds 0.8. Otherwise, the sample remains in the pool and a new  
 408      answer is generated. We allow up to three iterations and collect approximately 3.1K high-quality  
 409      SFT samples. During this process, we apply the Self-Dual reasoning template to ensure consistent  
 410      formatting. In both SFT and RL training, we only provide the problem statement as input to reduce  
 411      inference latency.

412      **RLVR Data Construction and Training.** Samples that fail the filtering process are combined with  
 413      the remaining 4K examples to form a 4.4K RLVR dataset, which encourages exploration in the RL  
 414      stage. For training, we use the LLaMA-Factory framework for SFT and the VERL framework to  
 415      perform reinforcement learning with standard GRPO. For SFT, we adopt LoRA with rank 32 and  
 416       $\alpha = 64$ , using a batch size of 32 and a learning rate of  $5 \times 10^{-5}$ . For GRPO, we train the full model  
 417      with a batch size of 64 and a learning rate of  $1 \times 10^{-6}$ . We set the group number  $n = 6$  and train for  
 418      3 epochs, which corresponds to about 200 steps. All training is conducted on NVIDIA H200 GPUs.

419  
 420      **4.2.2 MAIN RESULTS**

421  
 422      As shown in Table 5, we evaluate the effectiveness of Self-Dual during training by comparing it with  
 423      Qwen2.5-Math-7B-Instruct, a strong math-specific baseline. Qwen2.5-Math-7B-Instruct is trained  
 424      on large-scale CoT-TIR mixed data and RLAR, making it a competitive reference point. We focus  
 425      on the AIME25 benchmark for two reasons. First, AIME25 has high difficulty and can better re-  
 426      flect a model’s reasoning ability. Second, both Qwen2.5-Math-7B-Instruct and Qwen2.5-7B-Instruct  
 427      were released after the creation of AIME25, while our training data come exclusively from the 7.5K  
 428      MATH training set, ensuring no data contamination. Unlike prior work, Self-Dual does not rely  
 429      on large hybrid datasets. Instead, it enhances reasoning by explicitly combining complementary  
 430      paradigms. For this reason, we directly use Qwen2.5-7B-Instruct as the backbone to leverage its  
 431      existing natural language reasoning and coding capabilities. This setup allows us to isolate the con-  
 432      tribution of the Self-Dual framework during training. We further compare Qwen2.5-7B-SD against  
 433      prior methods on MATH-500 and GSM8K, with results presented in Table 6.

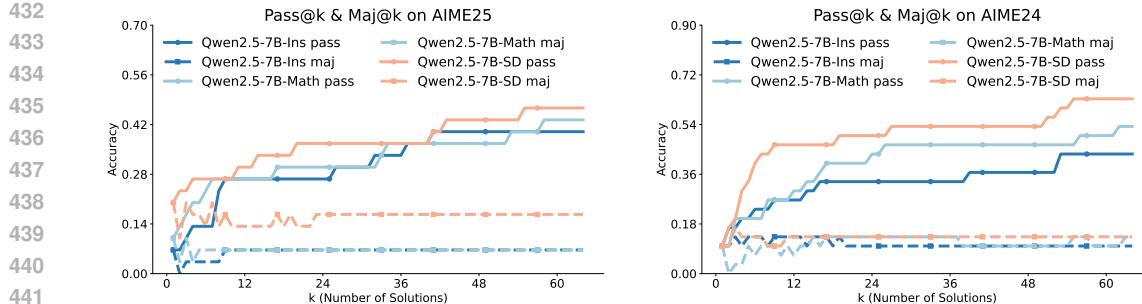


Figure 2: Pass@K and Maj@K curves of Qwen2.5-7B-SD, Qwen2.5-7B-Instruct, and Qwen2.5-7B-Math-Instruct on the AIME24 and AIME25 benchmarks.

We observe that Self-Dual training consistently improves performance across multiple math benchmarks. On MATH500, Qwen2.5-7B-Instruct achieves 72.2, while adding Self-Dual SFT raises the score to 74.4. With GRPO, the score further increases to 76.8, which is close to Qwen2.5-Math-7B-Instruct at 77.8. On AIME24 and AIME25, Self-Dual with SFT and GRPO, denoted as Qwen2.5-7B-SD, achieves 30.0 on AIME24 and 26.67 on AIME25, significantly higher than both the instruct and math baselines. These results show that Self-Dual not only improves single-answer accuracy but also provides stronger generalization on challenging reasoning tasks.

The comparison also highlights the efficiency of Self-Dual. Qwen2.5-Math-7B-Instruct relies on about 3000k mixed samples combining CoT and TIR, whereas Qwen2.5-7B-SD uses only the 7.5k MATH dataset. Despite this gap in scale, Qwen2.5-7B-SD reaches similar accuracy on MATH500 and even surpasses the math-specific model on AIME. This demonstrates that combining complementary reasoning with reinforcement learning can achieve strong results with much less data. We hope that the data construction approach of Self-Dual can inspire further efforts to enhance reasoning ability through data-centric methods.

### 4.3 PASS@K AND MAJ@K

As shown in Fig. 2, Self-Dual training yields clear performance gains. On both AIME24 and AIME25, the trained model achieves higher Pass@k than the baselines, including Qwen2.5-7B-Ins and Qwen2.5-7B-Math-Ins. The improvement is especially strong on AIME24, where the model shows higher accuracy at small  $k$  and maintains the lead as  $k$  increases. This indicates that Self-Dual improves not only single-solution quality but also overall exploration. Another interesting observation comes from the Maj@k curves. On AIME25, the Self-Dual model achieves a stable improvement, reaching about twice the accuracy of the baselines. This indicates that Self-Dual consistently increases the reliability of reasoning by integrating diverse paths. On AIME24, the gains are smaller, but the curve is more stable compared to both instruct and math-instruct baselines. This suggests that Self-Dual improves robustness and consistency, even in settings where performance improvements are less pronounced.

## 5 CONCLUSION

We presented Self-Dual, a framework that unifies NL and PL reasoning by generating complementary trajectories and integrating them through structured self-reflection. At inference time, Self-Dual consistently improves reasoning performance across benchmarks, showing that complementarity can be effectively exploited within a single forward pass. At training time, we further evaluate Self-Dual as a data construction strategy. On the challenging AIME25 benchmark, the resulting model Qwen2.5-7B-SD surpasses both the general instruct and math-specialized baselines, demonstrating the potential of Self-Dual to enhance reasoning under limited data conditions.

486 ETHICS STATEMENT  
487488 This work focuses on improving the mathematical reasoning capabilities of large language models  
489 through the integration of complementary reasoning paradigms. All experiments are conducted  
490 on publicly available benchmark datasets such as MATH500, AIME, and GSM8K, which contain  
491 no sensitive or personally identifiable information. Our methods do not involve the collection of  
492 new human data, nor do they raise direct privacy concerns. The proposed framework is intended to  
493 advance fundamental research in reasoning and transparency, and has no foreseeable immediate neg-  
494 ative societal impact. Nevertheless, as with all improvements to language models, there remains the  
495 potential risk of misuse in high-stakes applications. We encourage responsible deployment, includ-  
496 ing careful evaluation of robustness and fairness, before applying these methods beyond academic  
497 research.  
498499 REPRODUCIBILITY STATEMENT  
500501 All datasets used in this work are publicly available. Detailed experimental settings, including model  
502 configurations, training procedures, and evaluation protocols, are provided in the paper. We will  
503 release our code, data construction scripts, and trained checkpoints to facilitate full reproducibility.  
504505 REFERENCES  
506507 Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Ale-  
508 man, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical  
509 report, 2024. URL <https://arxiv.org/abs/2303.08774>.510 Rishabh Agarwal, Avi Singh, Lei M Zhang, Bernd Bohnet, Luis Rosias, Stephanie C.Y. Chan, Biao  
511 Zhang, Aleksandra Faust, and Hugo Larochelle. Many-shot in-context learning. In ICML 2024  
512 Workshop on In-Context Learning, 2024. URL <https://openreview.net/forum?id=goi7DFHlqS>.  
513514 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michał Podstawski, Lukas Gi-  
515 ainazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten  
516 Hoefler. Graph of thoughts: solving elaborate problems with large language models. AAAI  
517 Press, 2024. doi: 10.1609/aaai.v38i16.29720. URL <https://doi.org/10.1609/aaai.v38i16.29720>.  
518519 Wenhui Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting:  
520 Disentangling computation from reasoning for numerical reasoning tasks. Transactions on  
521 Machine Learning Research, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=YfZ4ZPt8zd>.  
522524 Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. Teaching large language models  
525 to self-debug. In The Twelfth International Conference on Learning Representations, 2024. URL  
526 <https://openreview.net/forum?id=KuPixIqPiq>.  
527528 I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian  
529 He, Graham Neubig, and Pengfei Liu. Factool: Factuality detection in generative AI - a  
530 tool augmented framework for multi-task and multi-domain scenarios, 2024. URL <https://openreview.net/forum?id=jolYuxpVn1>.  
531532 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,  
533 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John  
534 Schulman. Training verifiers to solve math word problems, 2021. URL <https://arxiv.org/abs/2110.14168>.  
535536 DeepSeek-AI. Deepseek-v3 technical report, 2024. URL <https://arxiv.org/abs/2412.19437>.  
537538 Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan,  
539 Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. RARR: Researching

540 and revising what language models say, using language models. In Anna Rogers, Jordan Boyd-  
 541 Gruber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association*  
 542 *for Computational Linguistics (Volume 1: Long Papers)*, pp. 16477–16508, Toronto, Canada, July  
 543 2023a. Association for Computational Linguistics. URL <https://aclanthology.org/2023.acl-long.910/>.

544

545 Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and  
 546 Graham Neubig. Pal: program-aided language models. In *Proceedings of the 40th International*  
 547 *Conference on Machine Learning*, ICML’23, pp. 10764 – 10799. JMLR.org, 2023b.

548

549 Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujiu Yang, Nan Duan, and Weizhu Chen.  
 550 CRITIC: Large language models can self-correct with tool-interactive critiquing. In *The Twelfth*  
 551 *International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=Sx038qxjek>.

552

553 Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujiu Yang, Minlie Huang, Nan Duan, and  
 554 Weizhu Chen. ToRA: A tool-integrated reasoning agent for mathematical problem solving. In  
 555 *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=Ep0TtjVoap>.

556

557 Melody Y. Guan, Manas Joglekar, Eric Wallace, Saachi Jain, Boaz Barak, Alec Helyar, Rachel Dias,  
 558 Andrea Vallone, Hongyu Ren, Jason Wei, Hyung Won Chung, Sam Toyer, Johannes Heidecke,  
 559 Alex Beutel, and Amelia Glaese. Deliberative alignment: Reasoning enables safer language  
 560 models, 2025. URL <https://arxiv.org/abs/2412.16339>.

561

562 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,  
 563 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms  
 564 via reinforcement learning, 2025. URL <https://arxiv.org/abs/2501.12948>.

565

566 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
 567 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In  
 568 *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks*  
 569 *Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.

570

571 Yiming Huang, Xiao Liu, Yeyun Gong, Zhibin Gou, Yelong Shen, Nan Duan, and Weizhu Chen.  
 572 Key-point-driven data synthesis with its enhancement on mathematical reasoning, 2024. URL  
 573 <https://arxiv.org/abs/2403.02333>.

574

575 Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks.  
 576 In *Proceedings of the 37th International Conference on Neural Information Processing Systems*,  
 577 NIPS ’23, pp. 39648 – 39677, Red Hook, NY, USA, 2023a. Curran Associates Inc.

578

579 Geunwoo Kim, Pierre Baldi, and Stephen McAleer. Language models can solve computer tasks.  
 580 In *Proceedings of the 37th International Conference on Neural Information Processing Systems*,  
 581 NIPS ’23, pp. 39648 – 39677, Red Hook, NY, USA, 2023b. Curran Associates Inc.

582

583 Hyeonwoo Kim, Gyoongjin Gim, Yungi Kim, Jihoo Kim, Byungju Kim, Wonseok Lee, and Chanjun  
 584 Park. Saas: Solving ability amplification strategy for enhanced mathematical reasoning in large  
 585 language models, 2024. URL <https://arxiv.org/abs/2404.03887>.

586

587 Chen Li, Weiqi Wang, Jingcheng Hu, Yixuan Wei, Nanning Zheng, Han Hu, Zheng Zhang, and  
 588 Houwen Peng. Common 7b language models already possess strong math capabilities, 2024a.  
 589 URL <https://arxiv.org/abs/2403.04706>.

590

591 Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif  
 592 Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, et al. Numinamath: The largest public dataset in  
 593 ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*,  
 594 13(9):9, 2024b.

595

596 Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. Mak-  
 597 ing language models better reasoners with step-aware verifier. In Anna Rogers, Jordan Boyd-  
 598 Gruber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for*  
 599 *Computational Linguistics (Volume 1: Long Papers)*, pp. 5315–5333, Toronto, Canada, July 2023.

594 Association for Computational Linguistics. URL [https://aclanthology.org/2023.](https://aclanthology.org/2023.acl-long.291/)  
 595 [acl-long.291/](https://aclanthology.org/2023.acl-long.291/).

596

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

601

602 Haoxiong Liu, Yifan Zhang, Yifan Luo, and Andrew Chi-Chih Yao. Augmenting math word  
 603 problems via iterative question composing, 2024. URL [https://arxiv.org/abs/2401.](https://arxiv.org/abs/2401.09003)  
 604 [09003](https://arxiv.org/abs/2401.09003).

605

606 Zimu Lu, Aojun Zhou, Ke Wang, Houxing Ren, Weikang Shi, Junting Pan, Mingjie Zhan, and Hong-  
 607 sheng Li. Mathcoder2: Better math reasoning from continued pretraining on model-translated  
 608 mathematical code, 2024. URL <https://arxiv.org/abs/2410.08196>.

609

610 Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jian-Guang Lou, Chongyang Tao, Xiubo Geng,  
 611 Qingwei Lin, Shifeng Chen, Yansong Tang, and Dongmei Zhang. Wizardmath: Empowering  
 612 mathematical reasoning for large language models via reinforced evol-instruct. In The Thirteenth  
 613 International Conference on Learning Representations, 2025a. URL [https://openreview.](https://openreview.net/forum?id=mMPMHW0dOy)  
 614 [net/forum?id=mMPMHW0dOy](https://openreview.net/forum?id=mMPMHW0dOy).

615

616 Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng,  
 617 Qingwei Lin, Shifeng Chen, Yansong Tang, and Dongmei Zhang. Wizardmath: Empowering  
 618 mathematical reasoning for large language models via reinforced evol-instruct, 2025b. URL  
 619 <https://arxiv.org/abs/2308.09583>.

620

621 Aman Madaan and Amir Yazdanbakhsh. Text and patterns: For effective chain of thought, it takes  
 622 two to tango, 2022. URL <https://arxiv.org/abs/2209.07686>.

623

624 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri  
 625 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad  
 626 Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. Self-refine:  
 627 iterative refinement with self-feedback. In Proceedings of the 37th International Conference on  
 628 Neural Information Processing Systems, NIPS '23, pp. 46534 – 46594, Red Hook, NY, USA,  
 629 2023. Curran Associates Inc.

630

631 Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. Investigating the limitations of transformers with  
 632 simple arithmetic tasks, 2021. URL <https://arxiv.org/abs/2102.13019>.

633

634 Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. Logic-LM: Empowering large  
 635 language models with symbolic solvers for faithful logical reasoning. In Houda Bouamor, Juan  
 636 Pino, and Kalika Bali (eds.), Findings of the Association for Computational Linguistics: EMNLP  
 637 2023, pp. 3806–3824, Singapore, December 2023. Association for Computational Linguistics.  
 638 doi: 10.18653/v1/2023.findings-emnlp.248. URL [https://aclanthology.org/2023.](https://aclanthology.org/2023.findings-emnlp.248)  
 639 [findings-emnlp.248/](https://aclanthology.org/2023.findings-emnlp.248).

640

641 Arkil Patel, Satwik Bhattacharya, and Navin Goyal. Are NLP models really able to solve sim-  
 642 ple math word problems? In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek  
 643 Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou  
 644 (eds.), Proceedings of the 2021 Conference of the North American Chapter of the Association  
 645 for Computational Linguistics: Human Language Technologies, pp. 2080–2094, Online, June  
 646 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.168. URL  
 647 <https://aclanthology.org/2021.naacl-main.168/>.

648

649 Debjit Paul, Mete Ismayilzada, Maxime Peyrard, Beatriz Borges, Antoine Bosselut, Robert West,  
 650 and Boi Faltings. REFINER: Reasoning feedback on intermediate representations. In Yvette  
 651 Graham and Matthew Purver (eds.), Proceedings of the 18th Conference of the European Chapter  
 652 of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1100–1126,  
 653 St. Julian's, Malta, March 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.eacl-long.67/>.

648 Jing Qian, Hong Wang, Zekun Li, Shiyang Li, and Xifeng Yan. Limitations of language models in  
 649 arithmetic and symbolic induction, 2022. URL <https://arxiv.org/abs/2208.05051>.  
 650

651 Yuxiao Qu, Tianjun Zhang, Naman Garg, and Aviral Kumar. Recursive introspection: Teaching lan-  
 652 guage model agents how to self-improve. In *The Thirty-eighth Annual Conference on Neural*  
 653 *Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=DRC9pZwBwR>.  
 654

655 Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan  
 656 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,  
 657 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin  
 658 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li,  
 659 Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang,  
 660 Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025.  
 661 URL <https://arxiv.org/abs/2412.15115>.  
 662

663 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 664 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathe-  
 665 matical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.  
 666

667 Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion:  
 668 language agents with verbal reinforcement learning. In *Proceedings of the 37th International*  
 669 *Conference on Neural Information Processing Systems, NIPS '23*, pp. 8634 – 8652, Red Hook,  
 670 NY, USA, 2023. Curran Associates Inc.

671 Zayne Rea Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. MuSR: Test-  
 672 ing the limits of chain-of-thought with multistep soft reasoning. In *The Twelfth International*  
 673 *Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=jenyYQzue1>.  
 674

675 Gemma Team. Gemma 3. 2025. URL <https://goo.gle/Gemma3Report>.  
 676

677 Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. DART-math: Difficulty-  
 678 aware rejection tuning for mathematical problem-solving. In *The Thirty-eighth Annual*  
 679 *Conference on Neural Information Processing Systems*, 2024a. URL <https://openreview.net/forum?id=zLU21oQjD5>.  
 680

682 Yuxuan Tong, Xiwen Zhang, Rui Wang, Ruidong Wu, and Junxian He. Dart-math: Difficulty-  
 683 aware rejection tuning for mathematical problem-solving, 2024b. URL <https://arxiv.org/abs/2407.13690>.  
 684

685 Ke Wang, Houxing Ren, Aojun Zhou, Zimu Lu, Sichun Luo, Weikang Shi, Renrui Zhang, Linqi  
 686 Song, Mingjie Zhan, and Hongsheng Li. Mathcoder: Seamless code integration in llms for en-  
 687 hanced mathematical reasoning, 2023. URL <https://arxiv.org/abs/2310.03731>.  
 688

689 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi,  
 690 Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language  
 691 models. In *Proceedings of the 36th International Conference on Neural Information Processing*  
 692 *Systems, NIPS '22*, pp. 24824 – 24837, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN  
 693 9781713871088.

694 Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun  
 695 Zhao. Large language models are better reasoners with self-verification. In *The 2023 Conference*  
 696 *on Empirical Methods in Natural Language Processing*, 2023. URL <https://openreview.net/forum?id=s4xIeYimGQ>.  
 697

699 Jian Xie, Kai Zhang, Jiangjie Chen, Tinghui Zhu, Renze Lou, Yuandong Tian, Yanghua Xiao,  
 700 and Yu Su. Travelplanner: a benchmark for real-world planning with language agents. In  
 701 *Proceedings of the 41st International Conference on Machine Learning, ICML'24*, pp. 54590  
 – 54613. JMLR.org, 2024.

702 Hao Yan, Saurabh Srivastava, Yintao Tai, Sida I. Wang, Wen-tau Yih, and Ziyu Yao. Learning to sim-  
 703 ulate natural language feedback for interactive semantic parsing. In Anna Rogers, Jordan Boyd-  
 704 Gruber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for*  
 705 *Computational Linguistics (Volume 1: Long Papers)*, pp. 3149–3170, Toronto, Canada, July 2023.  
 706 Association for Computational Linguistics. URL <https://aclanthology.org/2023.acl-long.177/>.

707

708 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu,  
 709 Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu,  
 710 Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical ex-  
 711 pert model via self-improvement, 2024a. URL <https://arxiv.org/abs/2409.12122>.

712

713 Ling Yang, Zhaochen Yu, Tianjun Zhang, Shiyi Cao, Minkai Xu, Wentao Zhang, Joseph E. Gonza-  
 714 lez, and Bin CUI. Buffer of thoughts: Thought-augmented reasoning with large language models.  
 715 In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL  
 716 <https://openreview.net/forum?id=AN01i9JPtb>.

717

718 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R  
 719 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In  
 720 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=5Xc1ecx01h>.

721

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

725

726 Yiyao Yu, Yuxiang Zhang, Dongdong Zhang, Xiao Liang, Hengyuan Zhang, Xingxing Zhang,  
 727 Mahmoud Khademi, Hany Hassan Awadalla, Junjie Wang, Yujiu Yang, and Furu Wei. Chain-  
 728 of-reasoning: Towards unified mathematical reasoning in large language models via a multi-  
 729 paradigm perspective. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Moham-  
 730 mad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for*  
 731 *Computational Linguistics (Volume 1: Long Papers)*, pp. 24914–24937, Vienna, Austria, July  
 732 2025a. Association for Computational Linguistics. ISBN 979-8-89176-251-0. URL <https://aclanthology.org/2025.acl-long.1213/>.

733

734 Yiyao Yu, Yuxiang Zhang, Dongdong Zhang, Xiao Liang, Hengyuan Zhang, Xingxing Zhang, Ziyi  
 735 Yang, Mahmoud Khademi, Hany Awadalla, Junjie Wang, Yujiu Yang, and Furu Wei. Chain-  
 736 of-reasoning: Towards unified mathematical reasoning in large language models via a multi-  
 737 paradigm perspective, 2025b. URL <https://arxiv.org/abs/2501.11110>.

738

739 Liang Zeng, Liangjun Zhong, Liang Zhao, Tianwen Wei, Liu Yang, Jujie He, Cheng Cheng, Rui  
 740 Hu, Yang Liu, Shuicheng Yan, Han Fang, and Yahui Zhou. Skywork-math: Data scaling laws  
 741 for mathematical reasoning in large language models – the story goes on, 2024. URL <https://arxiv.org/abs/2407.08348>.

742

743 Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. Automatic chain of thought prompting in  
 744 large language models. In *The Eleventh International Conference on Learning Representations*,  
 745 2023. URL <https://openreview.net/forum?id=5NTt8GFjUHkr>.

746

747 James Zhao, Yuxi Xie, Kenji Kawaguchi, Junxian He, and Michael Xie. Automatic model  
 748 selection with large language models for reasoning. In Houda Bouamor, Juan Pino, and  
 749 Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP*  
 750 2023, pp. 758–783, Singapore, December 2023. Association for Computational Linguistics.  
 751 doi: 10.18653/v1/2023.findings-emnlp.55. URL <https://aclanthology.org/2023.findings-emnlp.55/>.

752

753 Tong Zheng, Lichang Chen, Simeng Han, R. Thomas McCoy, and Heng Huang. Learning to reason  
 754 via mixture-of-thought for logical reasoning, 2025. URL <https://arxiv.org/abs/2505.15817>.

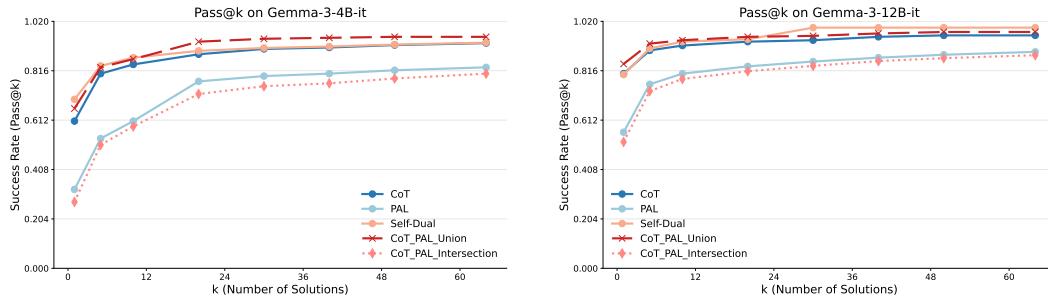
755

756 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,  
 757 Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables complex  
 758 reasoning in large language models, 2023. URL <https://arxiv.org/abs/2205.10625>.

760 Kun Zhou, Beichen Zhang, jiapeng wang, Zhipeng Chen, Xin Zhao, Jing Sha, Zhichao Sheng, Shijin  
 761 Wang, and Ji-Rong Wen. Jiuzhang3.0: Efficiently improving mathematical reasoning by training  
 762 small data synthesis models. In *The Thirty-eighth Annual Conference on Neural Information  
 763 Processing Systems*, 2024. URL <https://openreview.net/forum?id=ujDKXWTbJX>.

## 766 A COMPLEMENTARY EXPERIMENTS

### 768 A.1 PASS@K ON MATH500



780 Figure 3: Pass@k curves of Gemma3-4B-it and Gemma3-12B-it on CoT, PAL, Self-Dual, as  
 781 well as the union and intersection of CoT and PAL, evaluated on the MATH500 dataset.

783 The results on MATH500 show clear differences among CoT, PAL, Self-Dual, and their  
 784 combinations. CoT and PAL achieve moderate Pass@k, but both plateau early, indicating limited  
 785 exploration. Self-Dual consistently outperforms CoT and PAL across different values of  $k$ , which  
 786 demonstrates its ability to integrate complementary reasoning paths within a single inference.

788 The union of CoT and PAL sets provides the highest Pass@k across all  $k$ , confirming that the two  
 789 paradigms cover different solution spaces. The intersection curve is lower but stable, reflecting  
 790 their shared consistency. Importantly, Self-Dual tracks close to the union curve, suggesting that the  
 791 framework approximates the exploration benefit of multiple samples with only one inference. These  
 792 results highlight the effectiveness of Self-Dual in leveraging complementarity while maintaining  
 793 efficiency.

### 794 A.2 ABLATION STUDY

796 **Contribution of the Initial Stage.** The method DCT in Figure 4 and Table 2 refers to Double CoT  
 797 in Self-Dual, where two identical CoT-based reasoning paths are used in the initial stage instead of  
 798 complementary thinking modes. We designed this method to maintain the similar reasoning stages  
 799 and token consumption as Self-Dual, ensuring a fair comparison. As shown in Figure 4, Self-Dual  
 800 consistently outperforms both CoT and DCT across the two models. For example, on MATH-  
 801 500, DCT slightly outperforms CoT by 2.4% on Gemma-3-4B-it, suggesting that iterative reasoning  
 802 and reflection over dual outputs provide modest performance gains. However, Self-Dual achieves a  
 803 5.6% improvement over DCT, highlighting the effectiveness of leveraging complementary reasoning  
 804 modes across natural and programming languages.

805 **The Impact of Refine Stage.** Since both Self-Dual-Auto and Auto are built upon existing NL-  
 806 based and PL-based methods, we ensure that they utilize the same underlying methods and results  
 807 for consistency. On the diverse MATH-500 dataset, Self-Dual-Auto significantly outperforms Auto  
 808 in terms of accuracy, as shown in Table 1. For example, on Gemma-3-4B-it, Auto yields a -3.4%  
 809 change while Self-Dual-Auto achieves a +4.6% improvement. These results suggest that the refine-  
 810 ment stage in Self-Dual-Auto is more stable and adaptable than the simple selection mechanism

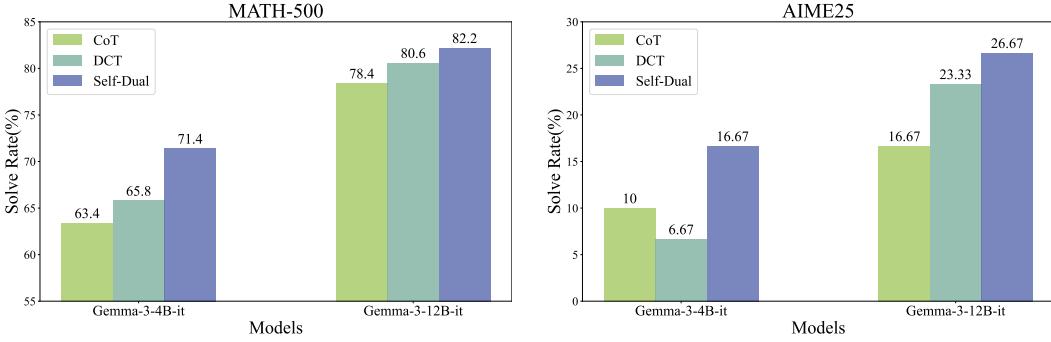


Figure 4: Results of CoT, DCT, and Self-Dual with Gemma-4-4B-it and Gemma-3-12B-it on MATH-500 (a) and AIME (b).

employed by Auto. On the AIME benchmark, except for Gemma-3-4B-it which outperforms Auto, all other models maintain comparable performance. This may suggest that the complementary potential between independently applied CoT and PAL has reached its upper bound on the AIME2025 dataset. The result that BAND outperforms both Self-Dual-Auto and Auto further demonstrates that the Self-Dual format is a more efficient way to combine complementary thinking modes.

### A.3 REMAINING RESULTS

Table 6 presents the remaining results on the GSM8K and SVAMP benchmarks. Consistent with the findings on MATH-500 and AIME, Self-Dual continues to demonstrate superior performance on GSM8K. Notably, even when using Gemma-3-27B-it, Self-Dual without external tools outperforms all other methods, including BAND itself. This advantage may be attributed to inherent randomness and the relatively lower difficulty of the GSM8K dataset. Furthermore, it can be observed that once accuracy surpasses 93%, the performance gains from prompt-based methods over standard approaches become marginal. This plateau may be constrained by the model’s inherent capabilities and limitations.

Table 6: The remaining results on SVAMP and GSM8K.

Methods	GSM8K			SVAMP			#Call
	G-4B	G-12B	G-27B	G-4B	G-12B	G-27B	
<i>NL-based methods</i>							
CoT	77.33	91.96	95.15	87.44	94.11	<b>94.67</b>	1
Reflextion	78.24	92.12	94.09	81.33	93	<u>93.67</u>	3
<i>PL-based methods</i>							
PAL	70.2	91.05	93.93	83.33	91.33	92	1
TIR	66.34	87.95	91.58	85.33	93.33	91.33	1-5
<i>w/o tool</i>	13.19	26.54	41.85	65	79.33	84	1
CRITIC	72.4	91.28	93.4	86	92	92	9
<i>w/o tool</i>	70.96	90.9	92.95	85.67	92	91.33	9
<i>Hybrid-based methods</i>							
Automatic	77.86	91.48	95.25	85.67	<u>94.33</u>	<u>93.67</u>	3
SD-Auto	<u>82.71</u>	<u>93.71</u>	<u>95.68</u>	<b>90</b>	<b>94.67</b>	93.33	3
Self-Dual	<b>83.07</b>	93.61	95.43	<u>88</u>	92	92.33	2
<i>w/o tool</i>	82.26	<b>93.93</b>	<b>95.83</b>	87.67	92	91	1

In contrast to previous results, Self-Dual Auto achieves the best performance on the SVAMP dataset, although its advantage gradually diminishes as model size increases. This trend may be attributed to two main factors: 1) SVAMP is relatively simple and imposes minimal demands on complex computation; and 2) existing models already possess strong capabilities for solving such problems. Together, these factors limit the effectiveness of complementary thinking modes in further improving



918 Table 8: The most detailed categorization of sampled information.  
919

920 Models	921 Methods	922 DC	923 CE	924 D2E	925 DC <sub>NL</sub>	926 DE <sub>NL</sub>	927 DC <sub>PL</sub>	928 DE <sub>PL</sub>	929 Badcode	930 FN	931 No2Paths
922 Deepseek -V3	BAND	77%	4%	4%	5%	0%	1%	2%	4%	1%	4%
	LEGO	67%	1%	1%	7%	0%	1%	2%	15%	0%	-
	Auto	67%	1%	1%	3%	4%	1%	2%	15%	0%	-
924 Gemma-3 -12B-it	BAND	67%	6%	6%	8%	0%	0%	2%	4%	1%	9%
	LEGO	53%	4%	4%	14%	1%	2%	4%	12%	2%	-
	Auto	53%	4%	4%	10%	5%	1%	5%	12%	2%	-
	DCT	80%	9%	9%	3%	0%	-	-	-	2%	4%

928 Table 9: Pass@1 accuracy of models in the zero-shot setting on MATH and GSM8k.  
929

931	932 Model	933 Base	934 MATH	935 GSM8k
			ZS Pass@1	ZS Pass@1
933	Skywork-Math (Zeng et al., 2024)	Llama-2	47.7	72.9
934	Xwin-Math(Li et al., 2024a)	Llama-2	40.6	82.6
935	WizardMath-Qwen (Luo et al., 2025b)	Qwen2.5-Math	77.8	93.9
936	WizardMath-Qwen (Luo et al., 2025b)	Qwen2.5	74.5	94.0
937	WizardMath-DeepSeek (Luo et al., 2025b)	DeepSeekMath	64.6	91.0
938	DART-Math (Tong et al., 2024b)	DeepSeekMath	52.9	88.2
939	DeepSeekMath-Instruct-7B (Shao et al., 2024)	DeepSeekMath	46.8	73.6
940	DeepSeekMath-RL-7B (Shao et al., 2024)	DeepSeekMath	51.7	88.2
941	MMIQC (Liu et al., 2024)	DeepSeekMath	45.3	79.0
942	KPMath-Plus (Huang et al., 2024)	DeepSeekMath	52.9	88.2
943	MetaMath-Llemma-7B (Yu et al., 2024)	Llemma-7B	30.0	69.2
944	MetaMath-Mistral-7B (Yu et al., 2024)	Mistral-7B	28.2	77.7
945	MathCoder2-DeepSeekMath (Lu et al., 2024)	DeepSeekMath	38.6	68.8
946	MathCoder2-Code-Llama (Lu et al., 2024)	Code-Llama	28.8	52.3
947	MathCoder2-Mistral (Lu et al., 2024)	Mistral	36.7	68.2
948	TORA (Gou et al., 2024b)	LLaMA-2	40.1	68.8
949	ToRA-CODE (Gou et al., 2024b)	CodeLLaMA	44.6	72.6
950	NuminaMath-CoT (Li et al., 2024b)	DeepSeekMath	55.2	75.4
951	CoR-Math-7B (Yu et al., 2025b)	DeepSeekMath	66.7	88.7
952	Qwen2.5-7B-Math-Instruct (Yang et al., 2024a)	Qwen2.5-7B-Math	<b>83.6</b>	<b>95.2</b>
953	Qwen-2.5-7B-SD	Qwen2.5-7B-Instruct	<u>78.12</u>	92.4

955  
956 A.6 INFERENCE-TIME CASE STUDY957  
958 We further discuss notable and illustrative examples in the process of reviewing the experimental  
959 results in this section.960  
961 **Overcoming the Consensus Fallacy.** In general, if both methods reach a consensus during the  
962 initial stage, the refine stage typically accepts this agreed-upon answer as the final correct solution.  
963 This observation is also evident from the sampled results presented in Table 2. But we also discover  
964 an interesting case worth discussing. The two solutions in the first stage consistently arrived at the  
965 incorrect answer of 1. However, during the refine stage, the model reanalyzed the problem and  
966 successfully produced the correct answer of  $\frac{1}{4}$ . We believe that the key component that triggers the  
967 realization of the initial error lies in the *Look Back* phase—specifically, the prompt segment: “*Let’s*  
968 *analyze the problem.*” This reflective step encourages the model to revisit and reassess the reasoning  
969 process, ultimately enabling it to identify and correct earlier mistakes.970  
971 **Reasoning, Verification and Reflection Format.** We also identify a typical case pattern with po-  
972 tential for further utilization: the PL-based method serves to verify the correctness of the result  
973 generated by the NL-based method, rather than producing an entirely separate reasoning path. This

972 format can improve the computational accuracy of the NL-based method, but it comes at the cost of  
 973 losing the potential benefits of independent PL-mode reasoning.  
 974

975 **Diverse Reasoning, Shared Conclusion.** Another meaningful Case is selected from the CC-type  
 976 results. Concretely, both solutions correctly identify 27 as the smallest positive perfect cube. The  
 977 NL-based solution offers a clear and logically structured mathematical explanation, while the PL-  
 978 based solution validates the answer through straightforward brute-force computation. This case also  
 979 exemplifies how complementary reasoning modes can manifest in practice.  
 980

981 **Hallucination in Self-Dual.** The sampled results from the Self-Dual method reveal several cases of  
 982 hallucination, which we report herein. A case illustrates a hallucination in which the LLM fabricates  
 983 a programming solution and evaluates it as if it were valid, despite the absence of real code. An-  
 984 other case demonstrates that the model mistakenly treats numerically equivalent values expressed in  
 985 LaTeX format and decimal notation as unequal. A third case involves an incorrect prediction of the  
 986 code execution result. We hope these cases may inspire future research on hallucination.  
 987

#### 988 A.7 COMPARE WITH OTHER METHODS

989 Table 9 summarizes pass@1 results on GSM8k and MATH in the zero-shot setting. All compared  
 990 models are standardized to 7B parameters to ensure fairness. Our proposed model, Qwen-2.5-7B-  
 991 SD, achieves 78.1% on MATH and 92.4% on GSM8k. This performance surpasses the expert model  
 992 on arithmetic reasoning and is second only to the proprietary Qwen-2.5-7B-Math-Instruct on MATH.  
 993 On GSM8k, our method also outperforms most fine-tuned and RL-based approaches. These results  
 994 demonstrate that combining NL and PL supervision, even with only 3k samples for SFT initial-  
 995 ization and 4k samples for RL training, can substantially enhance reasoning performance under a  
 996 lightweight training regime.  
 997

## 998 B THE USAGE OF LARGE LANGUAGE MODELS

999 LLMs were used in a limited capacity for brainstorming, minor phrasing suggestions, and partial  
 1000 writing refinement. All generated text was carefully reviewed and substantially revised by the au-  
 1001 thors. No LLMs were involved in research design, model development or experiments. Therefore,  
 1002 we confirm that LLMs did not play a significant role and should not be regarded as contributors.  
 1003

## 1004 C PROMPTS

1005 **Listing 1: The Self-Dual prompt designed for MATH-500 benchmark.**

1006 You need to follow the following process to solve the math problem  
 1007 1. Natural language based solution process, i.e., think step by step  
 1008 using natural language  
 1009 2. Programming language based solution process, i.e. writing python code  
 1010 to accomplish step-by-step thinking  
 1011 3. Then you need to scrutinize the solving step of each solution and  
 1012 compare them with each other. Give a new correct solution. You need  
 1013 to pay attention to a few points:  
 1014 1. please note that you need to ensure that the final answer is the  
 1015 correct answer to the user's problem  
 1016 2. The answers to the two solutions may not be the same, and you need  
 1017 to double-check to get the right answer.  
 1018 3. the output of the program is through the python interpreter, so you  
 1019 can trust the results of the code execution. So you need to pay  
 1020 attention to whether the code solves the problem or not.  
 1021 4. Give your answer by '\boxed{}'.

1022 ---

1023 Question: Suppose  $d \neq 0$ . We can write  $\left(12d+13+14d^2\right)$   
 1024  $+\left(2d+1\right)$ , in the form  $ad+bd+cd^2$ , where  $a$ ,  $b$ , and  
 1025  $c$  are integers. Find  $a+b+c$ .  
 1026

```

1026 Answer:
1027 Natural language based solution process:
1028  $12d + 13 + 14d^2 + 2d + 1 = 14d^2 + 14d + 14$ . Then  $a=14$ ,  $b=14$ ,  $c=0$ .  

1029  $b+c = 14+14+0 = 28$ . The answer is 28.
1030
1031 Programming language based solution process:
1032 '''python
1033 from sympy import symbols, simplify
1034
1035 d = symbols('d')
1036 expr = (12*d + 13 + 14*d**2) + (2*d + 1)
1037 simplified_expr = simplify(expr)
1038
1039 a = simplified_expr.coeff(d)
1040 b = simplified_expr.coeff('d')
1041 c = simplified_expr.coeff(d**2)
1042
1043 result = a + b + c
1044 print(result)
1045 '''
1046
1047 Scrutinize:
1048 ### Analysis of Solutions
1049 ###### Natural Language Solution:
1050 1. **Combining Like Terms**:
1051   - Correctly combines:  $(12d + 13 + 14d) + (2d + 1)$ 
1052   - Results in:  $14d + (12d + 2d) + (13 + 1) = 14d + 14d + 14$ 
1053 2. **Identifying Coefficients**:
1054   - Correctly identifies:
1055     - a (coefficient of d) = 14
1056     - b (constant term) = 14
1057     - c (coefficient of d) = 14 (not 0 as stated)
1058 3. **Calculation Error**:
1059   - Incorrectly states  $c = 0$  (should be 14)
1060   - Incorrect sum:  $14 + 14 + 0 = 28$  (should be  $14 + 14 + 14 = 42$ )
1061
1062 ###### Programming Solution:
1063 1. **Symbolic Computation**:
1064   - Correctly uses sympy to simplify the expression
1065   - Gets correct simplified form:  $14d + 14d + 14$ 
1066 2. **Coefficient Extraction**:
1067   - Correctly extracts coefficients:
1068     - a (d term) = 14
1069     - b (constant) = 14
1070     - c (d term) = 14
1071 3. **Correct Sum**:
1072   - Computes  $14 + 14 + 14 = 42$ 
1073
1074 ###### Key Differences:
1075 1. The natural language solution made an error in identifying c (coefficient of d) as 0 when it should be 14
1076 2. The programming solution correctly handles all coefficients through symbolic computation
1077
1078 ###### Corrected Solution:
1079 1. Combine like terms:
1080    $(12d + 13 + 14d) + (2d + 1) = 14d + 14d + 14$ 
1081 2. Identify coefficients:
1082   - a (d) = 14
1083   - b (constant) = 14
1084   - c (d) = 14

```

```

1080 3. Compute sum:
1081   a + b + c = 14 + 14 + 14 = 42
1082
1083 ### Final Answer:
1084 The correct answer is  $\boxed{42}$ .
1085 The programming solution was correct, while the natural language solution
1086 contained an error in identifying the d coefficient. The proper sum
1087 of coefficients is indeed 42.
1088
1089 ---
1090 Question: How many combinations of pennies (1 cent), nickels (5 cents)
1091 and/or dimes (10 cents) are there with a total value of 25 cents?
1092
1093 Answer:
1094 Natural language based solution process:
1095 1. Let  $p$  be the number of pennies,  $n$  the number of nickels and  $d$  the
1096 number of dimes. We want to find the number of non-negative
1097 integer solutions to  $p + 5n + 10d = 25$ .
1098
1099 We can consider cases based on the number of dimes  $d$ .
1100 \begin{itemize}
1101 \item  $d=0$ :  $p+5n=25$ . Then  $5n \leq 25$ , so  $n \leq 5$ .  $n$  can take
1102 values from 0 to 5, so there are 6 solutions.
1103 \item  $d=1$ :  $p+5n=15$ . Then  $5n \leq 15$ , so  $n \leq 3$ .  $n$  can take
1104 values from 0 to 3, so there are 4 solutions.
1105 \item  $d=2$ :  $p+5n=5$ . Then  $5n \leq 5$ , so  $n \leq 1$ .  $n$  can take
1106 values from 0 to 1, so there are 2 solutions.
1107 \end{itemize}
1108 The total number of solutions is  $6+4+2=12$ .
1109 Alternatively, we can express  $p = 25 - 5n - 10d$ . Since  $p \geq 0$ , we
1110 have  $25 - 5n - 10d \geq 0$ , so  $5n + 10d \leq 25$ , which simplifies
1111 to  $n + 2d \leq 5$ .
1112 For  $d=0$ ,  $n \leq 5$ , so  $n=0,1,2,3,4,5$  (6 solutions).
1113 For  $d=1$ ,  $n \leq 3$ , so  $n=0,1,2,3$  (4 solutions).
1114 For  $d=2$ ,  $n \leq 1$ , so  $n=0,1$  (2 solutions).
1115 For  $d=3$ ,  $n \leq -1$ , which is impossible.
1116 Thus the total number of solutions is  $6+4+2=12$ .
1117
1118 The answer is 12.
1119
1120 Programming language based solution process:
1121 '''python
1122 def solve():
1123     count = 0
1124     for p in range(26):
1125         for n in range(5):
1126             for d in range(5):
1127                 if 1*p + 5*n + 10*d == 25:
1128                     count += 1
1129     print(count)
1130 solve()
1131 '''
1132 '''output
1133 11
1134 '''
1135
1136 Scrutinize:
1137 ### Analysis of Solutions
1138
1139 #### Problem Interpretation:
1140 We need to count all non-negative integer solutions to the equation:
1141  $p + 5n + 10d = 25$ 
1142 where:
```

Listing 2: The reward prompt designed for filtering and GRPO training.

You are tasked with evaluating a response to a math problem to determine if it adheres to the specified format and requirements. You must assign a score between 0 and 1 based on the response's compliance with the format, focusing on completeness and clarity/structure. Your evaluation must check each component, identify deviations, provide suggestions for improvement, and output the results in a structured format that includes the natural language (NL) and programming language (PL) answers, as well as the assigned score. If the natural language solution is missing, report '[NL-Based Answer]' as 'None'.

### Expected Format and Requirements

The response must solve a math problem using the following structure:

1. Natural Language-Based Solution Process:
  - Provide a step-by-step explanation in natural language, solving the problem as a human would.
  - Clearly combine terms, derive coefficients, or perform calculations as needed.
  - Identify the final values (e.g., coefficients a, b, c) and compute the required answer.

1188 2. Programming Language-Based Solution Process:  
 1189   - Provide Python code that solves the problem step-by-step using  
 1190     symbolic computation (e.g., with 'sympy') or numerical methods.  
 1191   - The code must include a print statement to output the final result.  
 1192   - The code must be executable through a Python interpreter.  
 1193 3. Scrutiny and Comparison:  
 1194   - Analyze both the natural language and programming solutions.  
 1195   - Compare the two solutions to determine if they align or differ in  
 1196     their approach or results.  
 1197   - If discrepancies exist, explain why and provide a reconciled solution  
 1198 4. Final Answer:  
 1199   - Present the final answer in a LaTeX box format: '\boxed{}'.  
 1200   - The response must terminate after the '\boxed{}' answer, with no  
 1201     additional or redundant content.  
 1202 5. Additional Notes:  
 1203   - The response must explicitly label sections (e.g., "Natural Language-  
 1204     Based Solution Process," "Programming Language-Based Solution  
 1205     Process," "Scrutiny," "Final Answer").  
 1206   - The programming solution must be verifiable via Python execution.  
 1207   - The scrutiny section must compare the two approaches and explain any  
 1208     differences.  
 1209   - The response must be coherent, with logical and organized reasoning  
 1210     in both the natural language and scrutiny sections.  
 1211  
 1212   **### Scoring Criteria**  
 1213   Assign a score between 0 and 1 based on the following criteria:  
 1214  
 1215   - Completeness (60%):  
 1216     - 15% for including a natural language solution.  
 1217     - 15% for including a complete, executable Python code solution.  
 1218     - 15% for including a scrutiny section that compares both solutions.  
 1219     - 15% for including a final answer in '\boxed{}' format.  
 1220   - Clarity and Structure (40%):  
 1221     - 20% for clear, step-by-step, and coherent explanations in the natural  
 1222       language solution (if present), with logical progression and no  
 1223       rambling or disorganized reasoning.  
 1224     - 10% for proper labeling of sections (e.g., "Natural Language-Based  
 1225       Solution Process").  
 1226     - 10% for a thorough scrutiny section that clearly compares solutions  
 1227       and explains differences in an organized manner.  
 1228   - Penalties:  
 1229     - Deduct up to 15% per missing section (e.g., no natural language  
 1230       solution, no scrutiny section).  
 1231     - Deduct up to 10% for minor issues (e.g., unclear steps, improper  
 1232       labeling).  
 1233     - Deduct up to 20% for major issues (e.g., incomplete code, vague  
 1234       scrutiny section).  
 1235     - Format Penalty: Deduct up to 20% if the response continues with  
 1236       redundant or repeated content after the '\boxed{}' answer (e.g.,  
 1237       restarting the solution, adding unnecessary explanations, or  
 1238       reiterating the answer).  
 1239     - Clarity Penalty: Deduct up to 25% for incoherent, disorganized, or  
 1240       rambling reasoning in the natural language or scrutiny sections (e.g.  
 1241       ., illogical jumps, contradictory statements, excessive verbosity,  
 1242       or lack of clear progression).  
 1243     - A score of 0 is assigned if the response is entirely absent or  
 1244       irrelevant.  
 1245  
 1246   **### Evaluation Task**  
 1247   Given a response to a math problem, evaluate its compliance with the  
 1248     format and assign a score. Provide your assessment in the following  
 1249     structure:  
 1250  
 1. Compliance Check:

1242 - Does the response include all required sections (Natural Language-  
 1243 Based Solution, Programming Language-Based Solution, Scrutiny,  
 1244 Final Answer)?  
 1245 - Does the natural language solution (if present) follow a clear, step-  
 1246 by-step, and coherent process?  
 1247 - Does the programming solution include executable Python code?  
 1248 - Does the scrutiny section compare solutions and explain differences  
 1249 in an organized manner?  
 1250 - Is the final answer presented in '\boxed{{}}', with no redundant  
 1251 content afterward?  
 1252 2. Issues Identified:  
 1253 - List any deviations from the required format (e.g., missing sections,  
 1254 incorrect labeling, unclear explanations, non-executable code).  
 1255 - Note if the natural language solution is missing or incomplete.  
 1256 - Note if the scrutiny section fails to compare solutions or explain  
 1257 differences.  
 1258 - Note if the response continues with redundant or repeated content  
 1259 after '\boxed{{}}'.  
 1260 - Note if the natural language or scrutiny sections are incoherent,  
 1261 disorganized, or rambling.  
 1262 3. Suggestions for Correction:  
 1263 - For each issue, suggest how the response could be revised to meet the  
 1264 requirements (e.g., add missing section, clarify steps, ensure  
 1265 code is executable, terminate after '\boxed{{}}', improve coherence  
 1266 ).  
 1267 4. Score Breakdown:  
 1268 - Provide a breakdown of the score based on Completeness (60%) and  
 1269 Clarity and Structure (40%).  
 1270 - Justify deductions for any missing components, issues, or penalties (including  
 1271 format and clarity penalties).  
 1272 5. Overall Assessment:  
 1273 - State the final score (0 to 1, rounded to two decimal places).  
 1274 - Summarize whether the response complies with the format.  
 1275 **### Output Format**  
 1276 Provide your evaluation in the following format:  
 1277  
 1278 **Evaluation of Response**  
 1279  
 1280 **Compliance Check:**  
 1281 - [ ] All required sections included  
 1282 - [ ] Natural language solution is step-by-step, clear, and coherent (if  
 1283 present)  
 1284 - [ ] Programming solution includes executable Python code  
 1285 - [ ] Scrutiny section compares solutions and explains differences in an  
 1286 organized manner  
 1287 - [ ] Final answer in '\boxed{{}}' with no redundant content afterward  
 1288  
 1289 **Issues Identified:**  
 1290 - [List specific issues, e.g., "Missing natural language solution," "  
 1291 Improper section labeling," "Scrutiny section lacks comparison," "  
 1292 Redundant content after '\boxed{{}}','" "Incoherent reasoning in  
 1293 natural language solution."]  
 1294  
 1295 **Suggestions for Correction:**  
 1296 - [For each issue, provide a specific suggestion, e.g., "Include a  
 1297 natural language solution with step-by-step reasoning," "Label  
 1298 sections clearly as specified," "Terminate response after '\boxed  
 1299 {{}}','" "Organize reasoning to avoid rambling."]  
 1300  
 1301 **Score Breakdown:**  
 1302 - Completeness (60%): [Score, e.g., 45/60, with justification, e.g., "  
 1303 Missing natural language solution (-15%)"]  
 1304

1296 - Clarity and Structure (40%): [Score, e.g., 20/40, with justification, e  
 1297 .g., "Clear code but improper section labeling (-10%), redundant  
 1298 content after '\boxed{}' (-20%), incoherent reasoning (-20%)]  
 1299 - Total Score: [Score, e.g., 0.65]

1300 [NL-Based Answer]: [Extract the final answer from the natural language  
 1301 solution, e.g., "28," or "None" if missing]  
 1302 [PL-Based Answer]: [Extract the final answer from the programming  
 1303 solution, e.g., "42"]  
 1304 [Format Reward]: [Final score, e.g., 0.65]

1305 Overall Assessment:  
 1306 - [Summarize compliance, e.g., "The response partially complies due to a  
 1307 missing natural language solution, improper section labeling,  
 1308 redundant content after '\boxed{}', and disorganized reasoning.  
 1309 The programming solution and scrutiny section are present but need  
 1310 clearer comparison and more coherent reasoning."]

1311 ### Example Response to Evaluate  
 1312 [Insert the response to be evaluated here, or provide a placeholder if  
 1313 none is given.]

1314 ### Notes for the Model  
 1315 - If no response is provided, indicate that a response must be submitted  
 1316 for evaluation and assign a score of 0.  
 1317 - If the natural language solution is missing, report '[NL-Based Answer]'  
 1318 as 'None'.  
 1319 - Ensure the score is calculated systematically based on the criteria and  
 1320 rounded to two decimal places.  
 1321 - Extract the '[NL-Based Answer]' and '[PL-Based Answer]' directly from  
 1322 the respective sections of the response, if available.  
 1323 - If the problem involves symbolic computation (e.g., finding  
 1324 coefficients), ensure the programming solution uses appropriate tools  
 1325 like 'sympy'.  
 1326 - Do not evaluate the correctness of the answers (e.g., whether the  
 1327 coefficients or final answer are mathematically correct); focus  
 1328 solely on format compliance, completeness, and clarity/structure.  
 1329 - Do not modify the response or provide a new solution unless explicitly  
 1330 asked; focus on evaluating compliance and scoring.  
 1331 - Apply the format penalty (up to 20%) for responses that include  
 1332 redundant or repeated content after the '\boxed{}' answer, such as  
 1333 restarting the solution, reiterating the answer, or adding  
 1334 unnecessary explanations.  
 1335 - Apply the clarity penalty (up to 25%) for responses with incoherent,  
 1336 disorganized, or rambling reasoning in the natural language or  
 1337 scrutiny sections, such as illogical jumps, contradictory statements,  
 1338 excessive verbosity, or lack of clear progression.

1339 ### Example Application  
 1340 Evaluation of Response

1341 Compliance Check:  
 1342 - [x] All required sections included  
 1343 - [ ] Natural language solution is step-by-step, clear, and coherent  
 1344 - [x] Programming solution includes executable Python code  
 1345 - [x] Scrutiny section compares solutions and explains differences in an  
 1346 organized manner  
 1347 - [ ] Final answer in '\boxed{}' with no redundant content afterward

1348 Issues Identified:  
 1349 - The natural language solution is present but disorganized, with  
 1350 rambling explanations and illogical jumps (e.g., contradictory steps  
 1351 in deriving coefficients).  
 - Section labeling is correct, but the scrutiny section could be more  
 1352 concise in explaining differences.

1350 - The response includes redundant content after '\\boxed{{}}', restarting  
 1351 the explanation with an alternative approach.  
 1352

1353 Suggestions for Correction:  
 1354 - Revise the natural language solution to follow a clear, logical, and  
 1355 concise step-by-step process, eliminating rambling or contradictory  
 1356 statements.  
 1357 - Streamline the scrutiny section to focus on the key discrepancy without  
 1358 redundant explanation.  
 1359 - Remove the redundant content after '\\boxed{{}}' to adhere to the  
 1360 format requirement of terminating the response.

1360 Score Breakdown:  
 1361 - Completeness (60%): 60/60  
 1362 - Natural language solution included (15%).  
 1363 - Programming solution included and executable (15%).  
 1364 - Scrutiny section included (15%).  
 1365 - Final answer in '\\boxed{{}}' (15%).  
 1366 - Clarity and Structure (40%): 15/40  
 1367 - Natural language solution is disorganized and rambling (-15%, clarity  
 1368 penalty).  
 1369 - Sections are properly labeled (10%).  
 1370 - Scrutiny section is thorough but slightly verbose (10%).  
 1371 - Redundant content after '\\boxed{{}}' (-15%, format penalty).  
 1372 - Total Score: (60 + 15) / 100 = 0.75

1372 [NL-Based Answer]: 28  
 1373 [PL-Based Answer]: 42  
 1374 [Format Reward]: 0.75

1375 Overall Assessment:  
 1376 - The response complies with most format requirements, including all  
 1377 required sections. However, the natural language solution is  
 1378 disorganized and rambling, reducing clarity. The programming solution  
 1379 is executable, and the scrutiny section compares approaches but is  
 1380 slightly verbose. A significant deduction is made for redundant  
 1381 content after '\\boxed{{}}' and incoherent reasoning in the natural  
 1382 language solution.  
 1383 """

1384

1385 FILTER\_PROMPT = """

1386 You are tasked with evaluating a response to a math problem to determine  
 1387 if it adheres to the specified format and requirements. You must  
 1388 assign a score between 0 and 1 based on the response's compliance  
 1389 with the format, focusing on completeness and clarity/structure. Your  
 1390 evaluation must check each component, identify deviations, provide  
 1391 suggestions for improvement, and output the results in a structured  
 1392 format that includes the natural language (NL) and programming  
 1393 language (PL) answers, as well as the assigned score. If the natural  
 1394 language solution is missing, report '[NL-Based Answer]' as 'None'.

1394  
 1395 #### Expected Format and Requirements  
 1396 The response must solve a math problem using the following structure:  
 1397 1. Natural Language-Based Solution Process:  
 1398 - Provide a step-by-step explanation in natural language, solving the  
 1399 problem as a human would.  
 1400 - Clearly combine terms, derive coefficients, or perform calculations  
 1401 as needed.  
 1402 - Identify the final values (e.g., coefficients a, b, c) and compute  
 1403 the required answer.  
 1402 2. Programming Language-Based Solution Process:  
 1403 - Provide Python code that solves the problem step-by-step using  
 1404 symbolic computation (e.g., with 'sympy') or numerical methods.

1404     - The code must include a `print` statement to output the final result.  
 1405     - The code must be executable through a Python interpreter.  
 1406 3. Scrutiny and Comparison:  
 1407     - Analyze both the natural language and programming solutions.  
 1408     - Compare the two solutions to determine if they align or differ in  
       their approach or results.  
 1409     - If discrepancies exist, explain why and provide a reconciled solution  
 1410 .  
 1411 4. Final Answer:  
 1412     - Present the final answer in a LaTeX box format: `\boxed{}}`.  
 1413 5. Additional Notes:  
 1414     - The response must explicitly label sections (e.g., "Natural Language-  
       Based Solution Process," "Programming Language-Based Solution  
       Process," "Scrutinize," "Final Answer").  
 1415     - The programming solution must be verifiable via Python execution.  
 1416     - The scrutiny section must compare the two approaches and explain any  
       differences.  
 1417  
 1418 1419     ### Scoring Criteria  
 1420 Assign a score between 0 and 1 based on the following criteria:  
 1421 - Completeness (60%):  
 1422     - 15% for including a natural language solution.  
 1423     - 15% for including a complete, executable Python code solution.  
 1424     - 15% for including a scrutiny section that compares both solutions.  
 1425     - 15% for including a final answer in `\boxed{}}` format.  
 1426 - Clarity and Structure (40%):  
 1427     - 20% for clear, step-by-step explanations in the natural language  
       solution (if present).  
 1428     - 10% for proper labeling of sections (e.g., "Natural Language-Based  
       Solution Process").  
 1429     - 10% for a thorough scrutiny section that clearly compares solutions  
       and explains differences.  
 1430 - Penalties:  
 1431     - Deduct up to 15% per missing section (e.g., no natural language  
       solution, no scrutiny section).  
 1432     - Deduct up to 10% for minor issues (e.g., unclear steps, improper  
       labeling).  
 1433     - Deduct up to 20% for major issues (e.g., incomplete code, vague  
       scrutiny section).  
 1434     - A score of 0 is assigned if the response is entirely absent or  
       irrelevant.  
 1435  
 1436 1437     ### Evaluation Task  
 1438 Given a response to a math problem, evaluate its compliance with the  
 1439 format and assign a score. Provide your assessment in the following  
 1440 structure:  
 1441  
 1442 1. Compliance Check:  
 1443     - Does the response include all required sections (Natural Language-  
       Based Solution, Programming Language-Based Solution, Scrutiny,  
       Final Answer)?  
 1444     - Does the natural language solution (if present) follow a clear, step-  
       by-step process?  
 1445     - Does the programming solution include executable Python code?  
 1446     - Does the scrutiny section compare solutions and explain differences?  
 1447     - Is the final answer presented in `\boxed{}}`?  
 1448 2. Issues Identified:  
 1449     - List any deviations from the required format (e.g., missing sections,  
       incorrect labeling, unclear explanations, non-executable code).  
 1450     - Note if the natural language solution is missing or incomplete.  
 1451     - Note if the scrutiny section fails to compare solutions or explain  
       differences.  
 1452 3. Suggestions for Correction:  
 1453

1458 - For each issue, suggest how the response could be revised to meet the  
 1459 requirements (e.g., add missing section, clarify steps, ensure  
 1460 code is executable).  
 1461 4. Score Breakdown:  
 1462 - Provide a breakdown of the score based on Completeness (60%) and  
 1463 Clarity and Structure (40%).  
 1464 - Justify deductions for any missing components or issues.  
 1465 5. Overall Assessment:  
 1466 - State the final score (0 to 1, rounded to two decimal places).  
 1467 - Summarize whether the response complies with the format.  
 1468 **### Output Format**  
 1469 Provide your evaluation in the following format:  
 1470 **Evaluation of Response**  
 1471 **Compliance Check:**  
 1472 - [ ] All required sections included  
 1473 - [ ] Natural language solution is step-by-step and clear (if present)  
 1474 - [ ] Programming solution includes executable Python code  
 1475 - [ ] Scrutiny section compares solutions and explains differences  
 1476 - [ ] Final answer in '\boxed{{}}'  
 1477 **Issues Identified:**  
 1478 - [List specific issues, e.g., "Missing natural language solution," "  
 1479 Improper section labeling," "Scrutiny section lacks comparison."]  
 1480 **Suggestions for Correction:**  
 1481 - [For each issue, provide a specific suggestion, e.g., "Include a  
 1482 natural language solution with step-by-step reasoning," "Label  
 1483 sections clearly as specified."]  
 1484 **Score Breakdown:**  
 1485 - Completeness (60%): [Score, e.g., 45/60, with justification, e.g., "  
 1486 Missing natural language solution (-15%)"]  
 1487 - Clarity and Structure (40%): [Score, e.g., 30/40, with justification, e  
 1488 .g., "Clear code but improper section labeling (-10%)"]  
 1489 - Total Score: [Score, e.g., 0.75]  
 1490 **[NL-Based Answer]:** [Extract the final answer from the natural language  
 1491 solution, e.g., "28," or "None" if missing]  
 1492 **[PL-Based Answer]:** [Extract the final answer from the programming  
 1493 solution, e.g., "42"]  
 1494 **[Format Reward]:** [Final score, e.g., 0.75]  
 1495 **Overall Assessment:**  
 1496 - [Summarize compliance, e.g., "The response partially complies due to a  
 1497 missing natural language solution and improper section labeling. The  
 1498 programming solution and scrutiny section are present but need  
 1499 clearer comparison."]  
 1500 **### Example Response to Evaluate**  
 1501 [Insert the response to be evaluated here, or provide a placeholder if  
 1502 none is given.]  
 1503 **### Notes for the Model**  
 1504 - If no response is provided, indicate that a response must be submitted  
 1505 for evaluation and assign a score of 0.  
 1506 - If the natural language solution is missing, report '[NL-Based Answer]'  
 1507 as 'None'.  
 1508 - Ensure the score is calculated systematically based on the criteria and  
 1509 rounded to two decimal places.  
 1510 - Extract the '[NL-Based Answer]' and '[PL-Based Answer]' directly from  
 1511 the respective sections of the response, if available.

```

1512 - If the problem involves symbolic computation (e.g., finding
1513   coefficients), ensure the programming solution uses appropriate tools
1514   like 'sympy'.
1515 - Do not evaluate the correctness of the answers (e.g., whether the
1516   coefficients or final answer are mathematically correct); focus
1517   solely on format compliance, completeness, and clarity/structure.
1518 - Do not modify the response or provide a new solution unless explicitly
1519   asked; focus on evaluating compliance and scoring.
1520 ---
1521 ### Example Application
1522 Compliance Check:
1523 - [x] All required sections included
1524 - [x] Natural language solution is step-by-step and clear
1525 - [x] Programming solution includes executable Python code
1526 - [x] Scrutiny section compares solutions and explains differences
1527 - [x] Final answer in '\boxed{{}}'
1528 Issues Identified:
1529 - The natural language solution is present but contains a minor error in
1530   clarity (incorrectly states  $c = 0$ , though steps are clear).
1531 - Section labeling is correct, but the scrutiny section could be more
1532   concise in explaining differences.
1533 Suggestions for Correction:
1534 - Clarify the coefficient  $c = 14$  in the natural language solution to
1535   avoid confusion.
1536 - Streamline the scrutiny section to focus on the key discrepancy ( $c = 0$ 
1537   vs.  $c = 14$ ) without redundant explanation.
1538 Score Breakdown:
1539 - Completeness (60%): 60/60
1540   - Natural language solution included (15%).
1541   - Programming solution included and executable (15%).
1542   - Scrutiny section included (15%).
1543   - Final answer in '\boxed{{}}' (15%).
1544 - Clarity and Structure (40%): 35/40
1545   - Natural language solution is clear but has a minor error in
1546     coefficient  $c$  (-5%).
1547   - Sections are properly labeled (10%).
1548   - Scrutiny section is thorough but slightly verbose (10%).
1549   - Total Score:  $(60 + 35) / 100 = 0.95$ 
1550
1551 [NL-Based Answer]: 28
1552 [PL-Based Answer]: 42
1553 [Format Reward]: 0.95
1554
1555 Overall Assessment:
1556 - The response fully complies with the format, including all required
1557   sections. The natural language solution is clear but has a minor
1558   error in stating  $c = 0$ . The programming solution is executable and
1559   clear, and the scrutiny section effectively compares the two
1560   approaches. A slight deduction is made for the clarity issue in the
1561   natural language solution.
1562 ---
1563 ### Notes on Modifications
1564 - Scoring Structure: Changed to Completeness (60%) and Clarity and
1565   Structure (40%), removing Correctness as requested. Each section
1566   under Completeness is weighted equally (15%), and Clarity and
1567   Structure is split to emphasize the natural language explanation
1568   (20%) while maintaining weight for labeling and scrutiny (10% each).

```

1566 - NL Missing Case: Explicitly instructs to report '[NL-Based Answer]' as  
1567 'None' if the natural language solution is absent.  
1568 - Correctness Removed: The prompt no longer evaluates the mathematical  
1569 accuracy of the answers, focusing solely on format compliance,  
1570 completeness, and clarity.  
1571 - Penalties: Adjusted to reflect the new scoring weights, with deductions  
1572 scaled to the 60/40 split.  
1573 ---  
1574 Question: {question}  
1575 Model's Response: {response}  
1576 Ground Truth: {ground\_truth}  
1577  
1578 Give your detailed Evaluation.  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593  
1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619