

ReCode: Reinforcing Code Generation with Reasoning-Process Rewards

Lishui Fan^{1,2*}, Yu Zhang^{1*}, Mouxiang Chen^{1,2}, Zhongxin Liu^{1,2,3†}

¹College of Computer Science and Technology, Zhejiang University

²The State Key Laboratory of Blockchain and Data Security, Zhejiang University

³Hangzhou High-Tech Zone (Binjiang) Institute of Blockchain and Data Security

Abstract

In practice, rigorous reasoning is often a key driver of correct code, while Reinforcement Learning (RL) for code generation often neglects optimizing reasoning quality. Bringing process-level supervision into RL is appealing, but it faces two challenges. First, training reliable reward models to assess reasoning quality is bottlenecked by the scarcity of fine-grained preference data. Second, naively incorporating such neural rewards may suffer from reward hacking. This work proposes ReCode (**R**easoning-**R**einforced **C**ode Generation), a novel RL training framework comprising: (1) Contrastive Reasoning-Process Reward Learning (CRPL), which trains a reward model with synthesized optimized and degraded reasoning variants to assess the quality of reasoning process; and (2) Consistency-Gated GRPO (CG-GRPO), which integrates the reasoning-process reward model into RL by gating neural reasoning-process rewards with strict execution outcomes, using execution correctness as a hard gate to mitigate reward hacking. Additionally, to assess the reward model’s discriminative capability in assessing reasoning-process quality, we introduce LiveCodeBench-RewardBench (LCB-RB), a new benchmark comprising preference pairs of superior and inferior reasoning processes tailored for code generation. Experimental results across HumanEval(+), MBPP(+), LiveCodeBench, and BigCodeBench show that a 7B model trained with ReCode outperforms the base version by 16.1% and reaches performance comparable to GPT-4-Turbo. We further demonstrate the generalizability of ReCode by extending it to the math domain.

1 Introduction

Reinforcement Learning (RL) has emerged as a transformative paradigm for advancing Large Language Models (LLMs) in code generation (Cao

et al., 2026; Tang et al., 2026; Luo et al., 2025). However, existing approaches primarily rely on outcome-based supervision, such as test-case pass rates (Guo et al., 2025; Zeng et al., 2025), which provides limited guidance on how to reach correct programs. In practice, rigorous reasoning is often a key mechanism that supports correct code, rather than merely co-occurring with it (Wei et al., 2022; Lyu et al., 2023; Fan et al., 2025b). Consistently, our preliminary investigation shows a significant statistical association between reasoning-process quality and solution correctness (see Appendix A.4). Therefore, a pivotal question arises: *Can we exploit reasoning-process quality as an additional training signal—beyond final outcome—to improve code generation?*

However, translating this idea into a practical training framework presents two challenges. First, providing scalable, fine-grained supervision over reasoning processes is non-trivial. While utilizing strong LLMs as judges offers a potential solution, their high computational cost and latency render them impractical for the dense sampling required in RL training loops (Kim et al., 2023). A learned reasoning-process reward model provides a scalable way to supply such supervision for RL, but training it is difficult due to the scarcity of fine-grained preference data over reasoning processes. Specifically, human annotation is unscalable and subjective, while utilizing LLMs-as-a-judge often suffers from poor calibration when assigning scalar scores to nuanced reasoning process qualities (Feng et al., 2025; Ahn et al., 2024). Second, even with a learned reward model for reasoning processes, integrating it into RL poses risks. Naively incorporating neural rewards may suffer from reward hacking (Guo et al., 2025). In code generation, unit-test outcomes are strict and verifiable signals, whereas neural rewards are less constrained, making them easier to exploit.

To tackle these challenges, we introduce a

*Equal contribution.

†Corresponding author.

reasoning-reinforced framework in RL for code generation. Specifically, we treat reasoning processes as complementary learning signals and require them to be (i) reliably measured and (ii) safely integrated into RL. This leads to **Reasoning-Reinforced Code** Generation (ReCode), a novel RL framework with two components, i.e., Contrastive Reasoning-Process Reward Learning (CRPL) for reliable measurement of reasoning quality and Consistency-Gated GRPO (CG-GRPO) for safe integration of reasoning-process rewards in RL. CRPL performs contrastive data synthesis to generate *Optimized* and *Degraded* variants of reasoning processes by perturbing three key reasoning-process features, i.e., factual accuracy, logical rigor, and coherence, and thereby forms a multi-level preference ordering. From this ordering, CRPL derives fine-grained preference pairs and trains a reward model for reasoning-process discrimination. CG-GRPO augments GRPO by using the functional correctness of generated code as a hard gate for neural reasoning-process rewards, applying it only when the code runs correctly. This gating mechanism enforces reasoning–solution consistency, mitigating reward hacking while preserving informative gradients among correct solutions.

In addition, to assess how well the reward model can discriminate between superior and inferior reasoning processes, we construct LiveCodeBench-RewardBench (LCB-RB), a benchmark comprising 219 manually-checked preference pairs of reasoning processes derived from LiveCodeBench (Jain et al., 2025). To our knowledge, this is the first benchmark focusing on reasoning-process discrimination in code generation.

Extensive experiments demonstrate the effectiveness of ReCode. On four benchmarks, i.e., LiveCodeBench, HumanEval(+), MBPP(+), and BigCodeBench, Qwen2.5-Coder-7B-Instruct trained with ReCode achieves a relative improvement of 16.1% over the base model (50.4%→58.5%) and exhibiting performance comparable to GPT-4-Turbo. Additional results show that ReCode generalizes to Qwen3-4B and complements compiler-based process supervision. Meanwhile, reward models trained with CRPL achieve strong discriminative performance on LCB-RB and also perform well on the reasoning subset of RewardBench (Lambert et al., 2025), indicating that the model trained with CRPL learns a reasoning-process reward that generalizes and is consistent with functional correctness. We further observe

consistent gains by extending CG-GRPO to the math domain. The models, datasets, and code are publicly available¹.

2 ReCode

2.1 Contrastive Reasoning-Process Reward Learning

The effectiveness of incorporating reasoning-process rewards into RL hinges on a reward model that can provide a reliable measurement of reasoning quality. However, training such a model is bottlenecked by the scarcity of fine-grained preference data over reasoning processes. To address this, we propose Contrastive Reasoning-Process Reward Learning (CRPL) to automatically synthesise contrastive preference data by generating optimized and degraded reasoning variants, and train a reasoning-process reward model on the resulting preference data. As shown in Figure 1, CRPL induces a multi-level preference ordering over reasoning processes and trains the reward model to discriminate reasoning-process quality. The reward model then produces reasoning-process rewards for the policy optimization module.

To construct high-quality, reasoning-aware preference data, we first codify the intrinsic attributes of a reasoning process. Drawing upon insights from manual analysis and prior work (Zhu et al., 2025), we identify three critical dimensions of reasoning quality: (1) *Factual Accuracy*: assesses whether the reasoning contains factual errors; (2) *Logical Rigor*: assesses (a) whether redundant or misleading logical steps exist, and (b) whether missing logical connections result in incomplete reasoning; (3) *Logical Coherence*: assesses whether the logical flow maintains clear connections between steps.

We then prompt a powerful LLM, Qwen2.5-Coder-32B-Instruct, to synthesize a base reasoning process t for a given problem x , and leverage the dimensions mentioned above to generate variants of the base reasoning process t . Instead of generic rewriting, which may yield stylistic changes or ambiguous quality differences, we apply targeted transformations along specific reasoning-process dimensions to construct a multi-level preference ordering, providing clearer preference supervision for reward modeling. Concretely, we generate both optimized and degraded variants. (a) For optimized variants (t^+), the model is prompted to enhance specific dimensions, such as removing redundancies

¹<https://github.com/ZJU-CTAG/ReCode>

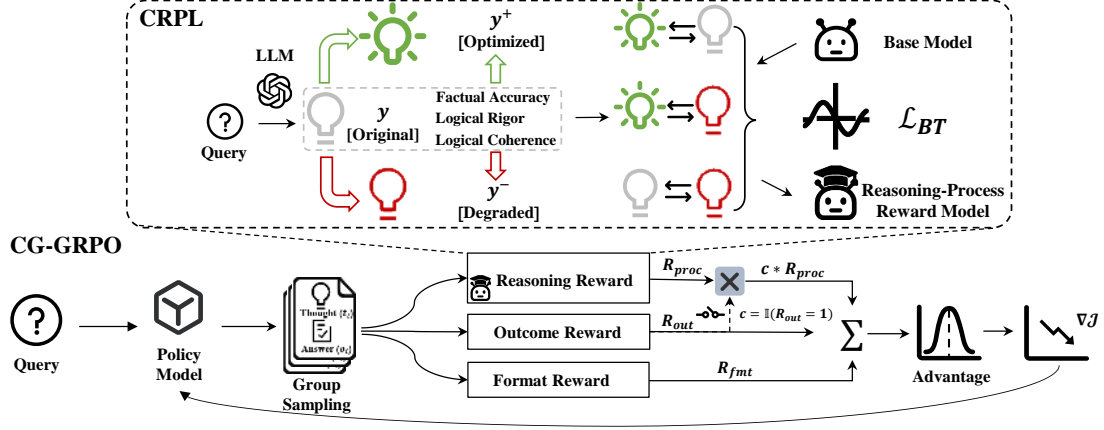


Figure 1: Overview of ReCode. **Top:** CRPL for reliably measuring reasoning quality. **Bottom:** CG-GRPO for safely integrating reasoning-process rewards into RL.

to improve logical rigor or clarifying transitions to boost logical coherence. (b) For degraded variants (t^-), the model is prompted to inject subtle flaws based on these dimensions, such as introducing a factual error or a logical gap. The optimization and degradation prompts are in Appendix A.6.

We formulate three distinct types of preference pairs to establish a continuous quality manifold: (a) *strong contrast pairs* ($x, t^+ \succ t^-$), (b) *fine-grained optimization pairs* ($x, t^+ \succ t$), and (c) *fine-grained degradation pairs* ($x, t \succ t^-$). This comprehensive training paradigm provides comparative signals that enable the reward model to internalize the intrinsic features of high-quality reasoning through relative contrast rather than absolute quantification. Notably, while similar evolutionary strategies like Evol-Instruct (Xu et al., 2024) have been explored for synthesizing supervised fine-tuning data, we novelly identify and validate their effectiveness in training discriminative reward models for distinguishing nuances in reasoning-process quality.

Reward modeling objective. Let $s_\theta(x, t) \in \mathbb{R}$ be the score given by the reward model s_θ for the reasoning process t under problem x . For each preference pair $(x, t^{(a)} \succ t^{(b)}) \in \mathcal{D}$, where $t^{(a)}$ is preferred over $t^{(b)}$, we optimize a Bradley–Terry (Bradley and Terry, 1952) objective, which is widely adopted in reward modeling (Ouyang et al., 2022):

$$\mathcal{L}_{BT}(\theta) = -\log \sigma(s_\theta(x, t^{(a)}) - s_\theta(x, t^{(b)}))$$

and sum over all constructed pairs. Minimizing this objective encourages s_θ to assign higher scores to higher-quality reasoning processes.

2.2 Consistency-Gated GRPO

Given the CRPL-trained reasoning-process reward model, this section instantiates the policy-optimization module that jointly leverages functional correctness and reasoning-process signals. Our implementation builds upon GRPO (Shao et al., 2024), which uses group-relative baselines computed from a set of sampled responses to reduce variance. In our setting, each response y is explicitly decomposed into a reasoning process and a solution, $y \triangleq (t, o)$, where t denotes the content inside the `<think>` block and o denotes the content inside the `<answer>` block.

While GRPO can optimize binary outcome rewards effectively, naively incorporating a neural reward makes the policy susceptible to reward hacking (Guo et al., 2025). For instance, the policy may learn to inflate the process score without improving the functional correctness of o . To mitigate this, we propose Consistency-Gated GRPO (CG-GRPO), a simple and effective gating mechanism that makes process optimization conditional on successful execution. As shown in Figure 1(c), CG-GRPO introduces a reward structure composed of three terms: (1) *Format Reward* (R_{fmt}): A binary reward ensuring structural compliance (i.e., correct usage of `<think>` and `<answer>` tags). $R_{fmt} = \mathbb{I}(\text{Valid Format})$, which has been demonstrated to be effective in prior works (Xie et al., 2025; Guo et al., 2025). (2) *Outcome Reward* (R_{out}): A strict binary signal derived from test case execution. $R_{out} = 1$ if the solution passes all test cases, and 0 otherwise. (3) *Reasoning Reward* (R_{proc}): A continuous scalar $r \in [0, 1]$ provided by the CRPL-trained reward model s_θ (Section 2.1), quantifying

the quality of the reasoning process t .

A naive way to incorporate R_{proc} is to linearly combine all reward terms (e.g., $R = R_{fmt} + R_{out} + R_{proc}$). However, since R_{proc} is a neural signal and less constrained than the signal provided by execution feedback, such linear aggregation renders the policy susceptible to reward hacking (see Section 6), where the policy inflates process scores without improving functional correctness. To mitigate this, our consistency gate grounds R_{proc} in a strict and verifiable signal. We activate R_{proc} only for functionally correct samples. In this way, functional correctness acts as a hard constraint preventing the policy from compromising correctness for higher neural scores, while still providing informative gradients to distinguish reasoning processes among correct solutions. The final reward R_i for the i -th sample is formulated as:

$$R_i = \underbrace{R_{fmt}^{(i)}}_{\text{Structure}} + \underbrace{R_{out}^{(i)}}_{\text{Correctness}} + \underbrace{\mathbb{I}(R_{out}^{(i)} = 1) \cdot R_{proc}^{(i)}}_{\text{Consistency-Gated Process}}$$

where $\mathbb{I}(\cdot)$ is the indicator function. This formulation ensures that the process reward is activated if and only if the extrinsic outcome is correct.

A key advantage of CG-GRPO is its ability to maintain learning dynamics when the model performs well. In standard GRPO, if a group of G samples is all correct ($R_{out}^{(i)} = 1, \forall i \in G$), the standard reward is uniform, leading to zero advantage. In contrast, our gated term $\mathbb{I}(R_{out} = 1) \cdot R_{proc}$ introduces meaningful variance among functionally correct solutions when their reasoning-process quality differs. This creates non-zero advantages, encouraging the model to prefer high-quality reasoning processes among correct solutions.

3 Benchmark Construction

Existing benchmarks for evaluating reward models (Lambert et al., 2025; Liu et al., 2025) primarily focus on distinguishing the correctness of final solutions rather than the quality of the intermediate reasoning process. This outcome-centric focus renders these benchmarks insufficient for evaluating reward models for training processes. To bridge this gap, we introduce LiveCodeBench-RewardBench (LCB-RB), a benchmark designed to discriminate between superior and inferior reasoning processes.

Following established protocols (Lambert et al., 2025; Liu et al., 2025), we use Qwen2.5-Coder-32B-Instruct with high-temperature sampling ($T =$

1.0) to generate 50 reasoning-solution pairs per problem from LiveCodeBench v6. We first partition the generated solutions into pass and fail sets based on test-case execution. In most cases, execution outcomes provide a useful coarse signal of reasoning-process quality-correct solutions are more likely to be supported by coherent reasoning than incorrect ones. However, our manual analysis reveals a small but important *Reasoning-Implementation Gap*: (i) a solution with a high-quality reasoning process may fail due to minor implementation issues (e.g., missing imports), and (ii) a solution with a flawed reasoning process may spuriously pass test cases due to randomness (Wang et al., 2023).

To purify the data and ensure strict alignment between reasoning and implementation, we implement a two-stage filtration pipeline:

Stage 1: Automated Dual-Consistency Check. We employ GPT-4o (Achiam et al., 2023) as an external validator to assess two criteria: (1) *Logical Soundness*, i.e., whether the reasoning processes contains substantive logical flaws; and (2) *Implementation Alignment*, i.e., whether the produced code faithfully implements the described plan. These checks target the two failure modes of the Reasoning-Implementation Gap. The prompt is in Appendix A.6. Based on GPT-4o’s assessment of criterion (2), we discard instances with poor alignment (i.e., code that does not implement the stated reasoning). Among the remaining aligned instances, we keep (a) logically sound processes judged by GPT-4o as *chosen candidates* and (b) logically flawed processes judged by GPT-4o as *rejected candidates*.

Stage 2: Human Adjudication. To mitigate potential bias of LLM-as-a-Judge (Zheng et al., 2023), two authors independently inspect the filtered instances. Adhering to the same evaluation criteria and prompt definitions used in Stage 1, the annotators manually evaluate each instance for both logical soundness and implementation alignment. Inter-annotator agreement measured by Cohen’s Kappa is $\kappa = 0.769$, indicating substantial agreement (Landis and Koch, 1977). Disagreements are subsequently resolved through consensus discussion between the annotators to ensure the high quality of the final labels.

From the adjudicated pool, we construct problem-wise preference pairs. We first exclude problems lacking either a valid chosen or a valid rejected process to ensure pairwise comparability.

Within the valid problems, we observe that valid rejected processes (clear logical flaws under reliable reasoning-code alignment) remain relatively scarce. To balance pairs, for each validated rejected process, we randomly sample one validated chosen process from the same problem instance.

We reserve problems from Oct. 2024 to Feb. 2025 exclusively for final code-generation evaluation to prevent data leakage. Finally, we obtain 219 preference pairs.

4 Experimental Setup

Reward Model Setup. Our reward models are initialized from Qwen2.5-Coder-7B-Base and Qwen2.5-Coder-3B-Base, and trained on preference pairs from the DeepCoder-Preview-Dataset (Luo et al., 2025), a corpus of 24k coding problems. We evaluate our reward model on LCB-RB and the code and math subsets of Reward-Bench (Lambert et al., 2025), using accuracy as the evaluation metric (Lambert et al., 2025; Liu et al., 2025). Baselines include: (a) the original model, (b) state-of-the-art (SOTA) reward models, including Starling-RM-34B (Zhu et al., 2023), EURUS-RM-7B (Yuan et al., 2025), Skywork-Reward-Llama-3.1-8B (Liu et al., 2024b), GPT-4-Turbo-2024-04-09 (Achiam et al., 2023), GPT-3.5-Turbo-0125 (Brown et al., 2020) and DeepSeek-V3 (Liu et al., 2024a), and (c) a Score-Based reward model. Further details are provided in appendix A.1.

RL Setup. We select Qwen2.5-Coder-7B-Instruct as policy model, using DeepCoder-Preview-Dataset for training. The prompt used during training is in Appendix A.3. We conduct evaluations on HumanEval(+) (Liu et al., 2023; Chen et al., 2021), MBPP(+) (Austin et al., 2021; Liu et al., 2023), BigCodeBench (Zhuo et al., 2025), and LiveCodeBench (Jain et al., 2025). We use greedy decoding and employ Pass@1 for evaluation. Baselines include: (a) the original model, (b) SOTA code models, including Llama3-Instruct-70B (Dubey et al., 2024), Deepseek-Coder-V2-Lite-Instruct (Zhu et al., 2024), Qwen2.5-Coder-Instruct 14B (Hui et al., 2024), GPT-4-Turbo-2024-04-09, and GPT-3.5-Turbo-0125, (c) the model fine-tuned with SFT on the same data, and (d) the model only with outcome and format rewards. Further details provided in Appendix A.2.

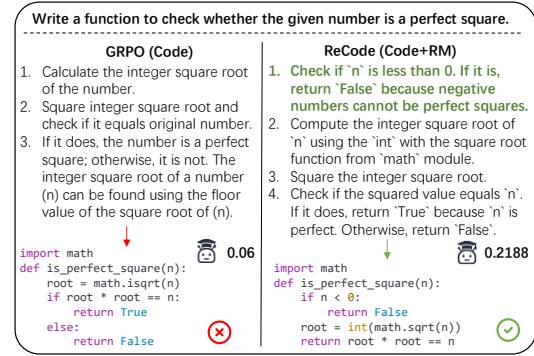


Figure 2: Example of reasoning processes generated by the base model with ReCode and with GRPO.

5 Results

We aim to answer the following research questions: **RQ1:** How effective is ReCode in improving code generation across different benchmarks?

RQ2: How effectively does the CRPL-trained reward model distinguish reasoning-process quality on LCB-RB? Does this discriminative capability generalize to other reasoning benchmarks?

RQ3: Can the ReCode training paradigm generalize to mathematical tasks which also rely on high-quality reasoning capabilities?

5.1 RQ1: Effectiveness of ReCode in Code Generation

ReCode effectively enhances the performance of the model. As shown in Table 1, ReCode achieves a relative improvement of 16.1% on average over the base model across all benchmarks, showing comparable performance to GPT-4-Turbo. Additionally, ReCode surpasses the outcome-only baseline by 6.7%, maintaining a consistent performance advantage as shown in Figure 3(a).

To further elucidate the mechanisms underlying ReCode’s performance gains, we inspect models’ outputs. Our analysis reveals that ReCode’s primary advantage lies in its ability to generate more comprehensive and logically sound reasoning processes, which help the model produce more accurate code. For example, as shown in Figure 2, when solving a perfect square problem, the model without a reasoning-process reward fails to consider negative numbers as edge cases in its reasoning process, causing it to pass only the basic test cases while failing the test cases in MBPP+. In contrast, the model with ReCode demonstrates sound reasoning by considering negative inputs at the outset of its reasoning process.

Model	Size	Humaneval		MBPP		LiveCodeBench			BigCodeBench		Avg
		HE	HE+	MBPP	MBPP+	Easy	Medium	Hard	Full	Hard	
GPT-4-Turbo	🔒	90.2	86.0	85.7	73.3	68.5	24.2	4.6	58.2	35.1	58.4
GPT-3.5-Turbo	🔒	72.6	67.7	84.1	71.2	46.3	9.4	5.6	50.6	21.6	47.7
Qwen2.5-Coder-Instruct	14B	89.6	87.2	86.2	72.8	61.0	11.3	2.8	48.4	22.2	53.5
DS-Coder-V2-Lite-Instruct	2.4/16B	81.1	75.6	82.8	70.4	43.9	5.7	5.6	36.8	16.2	46.5
Llama3-Instruct	70B	77.4	72.0	82.3	69.0	43.9	7.5	5.6	54.5	27	48.8
Qwen2.5-Coder-Instruct	7B	88.4	84.1	83.5	71.7	56.1	3.8	6.9	41.0	18.2	50.4
+SFT	7B	66.2	57.3	73.3	63.5	34.1	3.8	0.0	39.9	13.5	39.1
+GRPO	7B	85.9	81.1	86.7	75.1	58.5	15.1	9.7	52.0	29.7	54.9
+ReCode	7B	90.9	86.0	87.0	76.2	68.3	20.8	9.7	54.0	33.8	58.5

Table 1: Performance comparison of Qwen2.5-Coder-Instruct with ReCode against other baselines.

Model	Size	LCB-RB	RewardBench		Avg
			Code	Math	
DeepSeek-V3	671B	66.9	98.5	78.5	<u>81.3</u>
GPT-4-Turbo	🔒	<u>63.7</u>	<u>98.1</u>	67.3	76.4
GPT-3.5-Turbo	🔒	50.7	77.6	40.6	56.3
Starling-RM	34B	55.3	88.8	85.9	76.7
EURUS-RM	7B	57.0	92.8	79.9	76.5
Skywork					
-Reward	8B	61.6	-	-	-
-Llama-3.1					
Qwen2.5-Coder	3B	50.3	52.8	60.0	54.4
+Score	3B	51.0	49.4	47.2	49.2
+CRPL	3B	<u>57.9</u>	<u>63.6</u>	<u>93.5</u>	71.7
Qwen2.5-Coder	7B	<u>53.8</u>	<u>43.9</u>	<u>65.8</u>	<u>54.5</u>
+Score	7B	57.7	80.2	71.8	69.9
+CRPL	7B	62.6	88.6	99.8	83.7

Table 2: Performance comparison of reward model trained with CRPL against other baselines.

We further investigate the reward score of the generated reasoning process of the case. The results show that the model trained with ReCode achieves a reasoning-process reward of 0.2188, outperforming the model trained with outcome-only rewards, which obtains a score of 0.06. This demonstrates that our reward model is well-calibrated to reasoning-process quality, thereby providing effective supervision to guide the policy optimization. Additional examples can be found in the supplementary materials.

5.2 RQ2: Reward Model Effectiveness

Table 2 presents the performance of the reward model trained with CRPL and other baselines on LCB-RB and the reasoning subsets of RewardBench. Due to potential data contamination, we exclude the results of Skywork-Reward-Llama-3.1 from RewardBench². The reward model trained with CRPL effectively enhances the base model’s

²<https://gist.github.com/natolambert/1aed306000c13e0e8c5bc17c1a5dd300>

ability to identify high-quality reasoning processes on LCB-RB, surpassing other baselines. For instance, our 7B parameter model achieves performance comparable to GPT-4-Turbo.

Compared with the score-based baseline, the CRPL-trained models demonstrate substantial improvements. Specifically, our 3B and 7B models achieve relative improvements of 45.7% and 19.7% over score-based baselines, respectively. This indicates that training LLMs to distinguish between optimized and degraded versions of reasoning processes is more effective than learning from direct numerical scores. This is likely because LLMs are not inherently sensitive to fine-grained numerical values (Feng et al., 2025; Ahn et al., 2024), making it difficult to express the nuanced differences between reasoning processes via a scalar score.

Furthermore, the base model trained with CRPL achieves the highest average performance across reasoning subsets of RewardBench, outperforming the best baseline by 6.4% relatively. As RewardBench evaluates a model’s ability to discriminate the quality of final solutions, this result suggests that the discriminative patterns learned from reasoning processes generalize effectively to outcome assessment.

5.3 RQ3: Generalization to Mathematical Tasks

To further assess the generalization of ReCode, we extend it to the math domain where performance hinges critically on high-quality reasoning.

Experimental Setup We employ the same reward model in RQ1. For the policy model, we utilize Qwen2.5-Math-7B (Yang et al., 2024). We choose DAPO-Math-17k (Yu et al., 2025) as the training dataset, consisting of 17K mathematical data. We maintain the same experimental setup as in RQ1, except the training steps are reduced to 900,

Model	Size	MATH 500	Minerva Math	AIME24	Avg
GPT-4o		76.4	36.8	9.3	40.8
Llama-3.1-Inst	70B	64.6	35.3	16.7	38.9
Llama-3.1-Inst	405B	73.8	54.0	20.0	49.3
Eurus-2-PRIME	7B	79.2	38.6	26.7	48.2
Qwen2.5-Math-Inst	7B	79.8	37.1	13.3	43.4
Qwen2.5-Math	7B	46.9	15.5	11.2	24.5
+GRPO	7B	83.0	34.2	26.7	48.0
+ReCode	7B	83.0	38.2	33.3	51.5

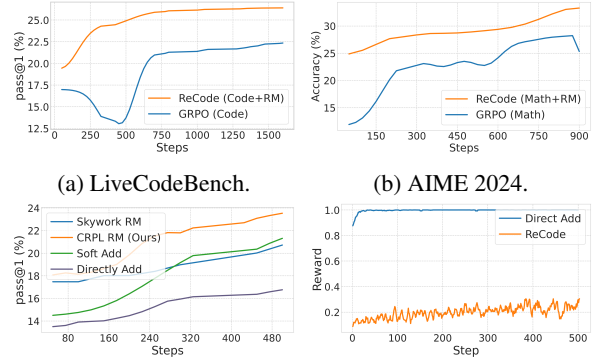
Table 3: Performance comparison of Qwen2.5-Math with ReCode against other baselines.

accounting for the smaller dataset size. Model performance is evaluated on MATH500 (Hendrycks et al., 2021), Minerva Math (Lewkowycz et al., 2022) and AIME 2024 (Art of Problem Solving, 2024). Following prior work (Cui et al., 2025; Dubey et al., 2024), we employ greedy decoding for evaluation and report accuracy metric. Baselines include: (1) original model, (2) base model trained via RL without reasoning-process rewards, and (3) SOTA mathematical models, including Llama-3.1-Instruct, GPT-4o-2024-0806, Eurus-2-PRIME (Cui et al., 2025),

As shown in Table 3, ReCode effectively enhances the model’s performance on mathematical tasks, indicating that ReCode generalizes beyond code generation. Specifically, the base model with ReCode demonstrates superior performance compared to several SOTA mathematical models. It demonstrates a 7.4% relative improvement over the RL baseline trained without reasoning-process rewards. To further illustrate the superiority of ReCode, we analyze the performance trajectory on AIME24. As shown in Figure 3(b), ReCode outperforms the baseline without reasoning-process rewards throughout the training process.


6 Discussion

Comparison with Other Reward Models. We replace the reward model in CG-GRPO with Skywork-Reward-Llama-3.1-8B to further validate the effectiveness of our reward model. Due to computational constraints, we use Qwen2.5-Coder-7B-Instruct as the policy model and evaluate performance changes over the first 500 steps on LiveCodeBench. As shown in Figure 3(c), CG-GRPO with our CRPL-trained reward model demonstrates superior performance. This may be because our reward model is trained on fine-grained, reasoning-aware preference pairs, which helps the model capture nuanced differences in reasoning quality. As a



(a) LiveCodeBench. (b) AIME 2024. (c) Reward-model and formulation ablations. (d) Reward hacking dynamics.

Figure 3: Overall performance of ReCode versus GRPO (a,b), with analyses of reward design (c) and reward hacking (d).



(a) Qwen2.5-Coder-7B. (b) Qwen2.5-Coder-3B.

Figure 4: Preference-source ablations across Qwen2.5-Coder backbones.

result, when used in RL, our reward model yields more precise and informative learning signals for policy optimization than reward models trained on outcome-centric preference data.

The Impact of Reward Hacking. To further demonstrate the susceptibility of process-reward RL to reward hacking, we introduce two alternative reward formulations that relax the consistency constraint used in CG-GRPO. (1) Direct Addition directly adds the process reward to the total reward: $R_i = R_{fmt}^{(i)} + R_{out}^{(i)} + R_{proc}^{(i)}$. (2) Soft-Gated Formulation replaces the hard execution gate with a soft weight based on the execution pass rate: $R_i = R_{fmt}^{(i)} + R_{out}^{(i)} + P_{out}^{(i)} \cdot R_{proc}^{(i)}$, where $P_{out}^{(i)}$ denotes the pass rate of output o_i . We run controlled experiments using Qwen2.5-Coder-7B-Instruct as the policy and track performance over the first 500 training steps on LiveCodeBench. As shown in Figure 3(c), both alternative formulations consistently underperform CG-GRPO, with direct addition yielding the largest drop. These results suggest that reasoning-process rewards computed from incorrect programs are inherently noisy and could be exploited by the policy. Without the strict consistency constraint, the policy may over-optimize

$R_{proc}^{(i)}$ even when the resulting code fails execution, ultimately degrading code-generation performance. In addition, we visualize the evolution of the reasoning-process reward during training for Direct Addition and ReCode (Figure 3(d)). Under Direct Addition, the process reward saturates to near 1.0 within the first 50 steps while downstream performance stagnates, indicating that the policy exploits the unconstrained neural signal. In contrast, under ReCode, the process reward grows gradually and remains moderate, accompanied by steady performance gains, confirming that the consistency gate effectively prevents reward inflation.

Impact of Different Preference Pair Combinations. We train base models on different combinations with identical experimental settings. As illustrated in Figure 4, models achieve optimal performance when trained on all types of preference pairs. This improvement indicates that each pair type contributes useful supervision. We hypothesize that aggregating diverse pair types provides a more comprehensive learning signal, which enables the model to better distinguish reasoning processes. Notably, training solely on Opt–Deg pairs yields the best performance among all single-pair settings, outperforming Opt–Ori and Ori–Deg by 22.2% on average. This may be because pairs with larger quality separation provide a higher-contrast and less ambiguous preference signal, making it easier for the reward model to learn discriminative features of reasoning-process quality.

Impact of Data Generator Diversity in CRPL. To investigate whether the reward model overfits to generator-specific stylistic artefacts rather than intrinsic reasoning semantics, we construct a cross-generator variant of CRPL by randomly splitting original data into two halves: for one half, we use Llama-3.1-70B-Instruct to synthesize reasoning processes, while the other half is generated by Qwen2.5-Coder-32B-Instruct. We train reward models on the mixed dataset. As shown in Table 6, single-generator training consistently outperforms cross-generator training. For example, with Qwen2.5-Coder-7B as the base model, the single-generator setting improves performance by 11.0% on average compared to the cross-generator variant. We hypothesize that this gap reflects a signal-to-noise trade-off under a fixed budget: preference pairs produced by a stronger generator tend to provide cleaner pairwise supervision, whereas mixing generators increases stylistic and structural variability

Model	LiveCodeBench			Avg
	Easy	Medium	Hard	
Qwen3-4B-Instruct	68.2	32.1	8.3	30.7
+GRPO	82.9	22.6	11.1	32.5
+ReCode	82.9	32.1	12.5	36.1

Table 4: Cross-model generalization results on LiveCodeBench.

ity of reasoning processes, diluting the preference signal and making it harder for the reward model to learn stable reasoning-process quality distinctions. **Cross-model Generalization.** To examine whether ReCode generalizes beyond the Qwen2.5 family, we reuse the same CRPL-trained reward model from RQ1 and apply it to Qwen3-4B-Instruct as the policy model. We follow the same training setup as in RQ1, except that training is limited to 400 steps due to computational cost. As shown in Table 4, ReCode yields an 11.1% relative improvement over outcome-only GRPO on LiveCodeBench. This result suggests that the learned reasoning-process reward transfers across model families.

Complementarity with Compiler-based Supervision We investigate whether ReCode complements compiler-based process supervision. Following StepCoder (Dou et al., 2024), we define a reward based on compilation and test-execution status: +1 for passing all unit tests, −0.3 for test failures, −0.6 for runtime errors, and −1 for compilation errors. We compare three variants trained for 400 steps: (i) GRPO with the compiler-based reward only, (ii) ReCode only, and (iii) ReCode combined with the compiler-based reward. As shown in Table 5, ReCode yields a 5.0% relative improvement over the compiler-based baseline, while the combined objective achieves the best overall performance. This suggests that the two signals are complementary: compiler-based rewards supervise the correctness of the generated program, whereas ReCode additionally evaluates the semantic quality of the reasoning process and encourages the model to produce more coherent reasoning before code generation.

Generation Efficiency. ReCode is also more generation-efficient than outcome-only RL. On LiveCodeBench, ReCode achieves a higher average Pass@1 while generating 23.4% fewer tokens on average. This suggests that the gains stem from more effective reasoning rather than increased test-time compute. A per-difficulty breakdown is pro-

Model	LiveCodeBench			Avg
	Easy	Medium	Hard	
Qwen2.5-Coder-7B-Instruct	56.1	3.8	6.9	18.1
+GRPO (Compiler-based)	63.4	15.1	8.3	24.1
+ReCode	65.9	15.1	9.7	25.3
+ReCode +Compiler-based	75.6	17.0	6.9	27.1

Table 5: Comparison with compiler-based supervision on LiveCodeBench.

Model	Size	LCB-RB	RewardBench Code	RewardBench Math
Qwen2.5-Coder	3B	50.3	52.8	60.0
+Cross-Gen (50/50)	3B	52.9	54.6	81.7
+Single-Gen	3B	57.9	63.6	93.5
Qwen2.5-Coder	7B	53.8	43.9	65.8
+Cross-Gen (50/50)	7B	55.3	81.5	89.4
+Single-Gen	7B	62.6	88.6	99.8

Table 6: Reward model performance under single-generator vs. cross-generator preference data synthesis.

vided in Appendix A.5.

7 Related Work

RL for Code Generation. Code generation offers a natural verification signal via unit tests (Chen et al., 2023, 2024), enabling outcome-based rewards for RL. Building on this property, a line of work improves LLM coding by leveraging execution feedback. CodeRL (Le et al., 2022) applies RL with unit-test feedback, PPOCoder (Shojaee et al., 2023) enriches rewards by combining unit tests with syntactic and semantic similarity to ground-truth solutions, and DeepSeek-R1 (Guo et al., 2025) adopts GRPO using relative pass rates within a group of samples. To address the sparsity of binary signals provided by outcome-based rewards, recent works have shifted towards incorporating fine-grained process supervision into RL. StepCoder (Dou et al., 2024) uses curriculum learning that progresses from code-completion tasks to full generation, masking unexecuted tokens via test coverage for fine-grained PPO updates. PRL-Coder (Ye et al., 2025) learns a line-level process reward model validated by the compiler and the test cases from mutation and refactoring. Despite these advances, existing methods remain largely implementation-centric. They primarily evaluate the correctness of generated code rather than the logical soundness of reasoning processes. Motivated by evidence that reasoning quality affects functional correctness, we propose ReCode to explicitly incentivize rigorous reasoning. As prior

approaches largely focus on implementation-level refinement, ReCode is potentially complementary to them by improving reasoning-process quality.

Reward Model Evaluation on Reasoning. Evaluating the performance of reward models for reasoning tasks typically relies on verifiable problems. For example, the code subset of RewardBench (Lambert et al., 2025) utilizes HumanEval-Pack (Muennighoff et al., 2024), a multilingual extension of the HumanEval dataset. However, they typically focus only on the correctness of the final output, neglecting the quality of the intermediate reasoning process that produced it. Additionally, there is a risk of data contamination for benchmarks derived from HumanEval (Jain et al., 2025). This can lead to inflated performance metrics that do not reflect true capabilities. To address these, we construct LCB-RB sourced from LiveCodeBench (Jain et al., 2025), which allows us to reliably evaluate a reward model’s discrimination capabilities on the intermediate steps of problem-solving.

8 Conclusion

We propose ReCode (**R**easoning-**R**einforced **C**ode Generation), a novel RL framework designed to align code generation with high-quality reasoning processes. It comprises two components, i.e., CRPL for reliable measurement of reasoning quality and CG-GRPO for safe integration of reasoning-process rewards in RL. Our reward model achieves strong performance on LCB-RB, a new benchmark designed to distinguish between high-quality and flawed reasoning processes. Extensive experiments demonstrate the effectiveness of ReCode, and this paradigm generalizes to the math domain.

Limitations

While our work displays many strengths, we highlight three limitations:

Scalability in Long-Context Reasoning Processes. We train with a 4K output length due to compute constraints, which limits direct validation on long-horizon settings where reasoning processes can exceed 30K tokens. Scaling ReCode to long contexts requires longer-context training and a long-context generator for the reward model. Future work could explore extending our framework to long-context reasoning processes.

Limited Scale of LCB-RB. LCB-RB is derived from LiveCodeBench. To avoid potential data leakage and ensure label reliability, it contains 219

manually verified preference pairs. While this improves the reliability of evaluation, the small scale may limit coverage. Future work could expand reasoning-centric benchmarks using additional problems and further evaluate reward models.

Limited Exploration of Other LLMs We observe cross-domain generalization when applying ReCode to Qwen2.5-Math-7B, and cross-architecture transfer to Qwen3-4B-Instruct. However, our evaluation is still limited to models up to 7B parameters, leaving open how ReCode behaves at larger scales. We leave broader scaling studies as future work.

Ethical considerations

All the datasets we use to fine-tune LLMs are publicly available and are for research purposes only. Beyond existing datasets, we also create synthetic preference supervision via CRPL and curate LCB-RB from LiveCodeBench. These artifacts are derived from programming tasks and model-generated reasoning processes rather than users' personal data. We conduct automated filtering and manual spot-checking to reduce the presence of personally identifying information (PII) and overtly harmful or offensive content in the generated data, and we release artifacts with a research-only intended use under the licenses and access conditions of the original sources. The open- and closed-source LLMs used (e.g., Qwen2.5 series, GPT-4/3.5) have their own training and deployment considerations documented by their creators. Our reward models are used to score reasoning traces, and models trained with ReCode are optimized to generate responses whose reasoning is consistent with the produced implementation. However, we acknowledge that LLMs, including those used in our study, may occasionally produce improper or harmful content. Such outputs are unintended and do not reflect the views or intentions of the authors.

Acknowledgements

This research is supported by the National Natural Science Foundation of China (No.92582107) and Zhejiang Provincial Natural Science Foundation of China (No.LZ25F020003).

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,

Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Janice Ahn, Rishu Verma, Renze Lou, Di Liu, Rui Zhang, and Wenpeng Yin. 2024. **Large language models for mathematical reasoning: Progresses and challenges**. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024: Student Research Workshop, St. Julian's, Malta, March 21-22, 2024*, pages 225–237. Association for Computational Linguistics.

Art of Problem Solving. 2024. Aime problems and solutions. https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions. Accessed: 2025-04-20.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.

Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Ruisheng Cao, Mouxiang Chen, Jiawei Chen, Zeyu Cui, Yunlong Feng, Binyuan Hui, Yuheng Jing, Kaixin Li, Mingze Li, Junyang Lin, and 1 others. 2026. Qwen3-coder-next technical report. *arXiv preprint arXiv:2603.00729*.

Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. 2023. **Codet: Code generation with generated tests**. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, and 1 others. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.

Mouxiang Chen, Zhongxin Liu, He Tao, Yusu Hong, David Lo, Xin Xia, and Jianling Sun. 2024. B4: Towards optimal assessment of plausible code solutions with plausible tests. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, pages 1693–1705.

Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu,

- Qixin Xu, Weize Chen, and 1 others. 2025. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*.
- Shihan Dou, Yan Liu, Haoxiang Jia, Enyu Zhou, Limao Xiong, Junjie Shan, Caishuang Huang, Xiao Wang, Xiaoran Fan, Zhiheng Xi, and 1 others. 2024. StepCoder: Improving code generation with reinforcement learning from compiler feedback. In *62nd Annual Meeting of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand, Aug 11 2024-Aug 16 2024*, pages 4571–4585. Association for Computational Linguistics (ACL).
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Kaixuan Fan, Kaituo Feng, Haoming Lyu, Dongzhan Zhou, and Xiangyu Yue. 2025a. Sophiav1-r1: Reinforcing mllms reasoning with thinking reward. *arXiv preprint arXiv:2505.17018*.
- Lishui Fan, Mouxiang Chen, and Zhongxin Liu. 2025b. Sek: Self-explained keywords empower large language models for code generation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 6249–6278.
- Guhao Feng, Kai Yang, Yuntian Gu, Xinyue Ai, Shengjie Luo, Jiacheng Sun, Di He, Zhenguo Li, and Liwei Wang. 2025. How numerical precision affects arithmetical reasoning capabilities of llms. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 46–85.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- Jujie He, Jiakai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, and 1 others. 2025. Skywork open reasoner 1 technical report. *arXiv preprint arXiv:2505.22312*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, and 1 others. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Seungone Kim, Jamin Shin, Yejin Cho, Joel Jang, Shayne Longpre, Hwaran Lee, Sangdoon Yun, Seongjin Shin, Sungdong Kim, James Thorne, and 1 others. 2023. Prometheus: Inducing fine-grained evaluation capability in language models. In *The Twelfth International Conference on Learning Representations*.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, Lester James V. Miranda, Bill Yuchen Lin, Khyathi Raghavi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. [Rewardbench: Evaluating reward models for language modeling](#). In *Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, Findings of ACL, pages 1755–1797. Association for Computational Linguistics.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Hung Le, Yue Wang, Akhilesh Deepak Gotmare, Silvio Savarese, and Steven CH Hoi. 2022. Coder1: mastering code generation through pretrained models and deep reinforcement learning. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pages 21314–21328.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and 1 others. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yuhui Zhou. 2024b. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint arXiv:2410.18451*.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. [Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023*,

- NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023.*
- Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. 2025. **Rm-bench: Benchmarking reward models of language models with subtlety and style.** In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, Ce Zhang, Li Erran Li, Raluca Ada Popa, and Ion Stoica. 2025. Deepcoder: A fully open-source 14b coder at o3-mini level. <https://www.together.ai/blog/deepcoder>. Together AI Blog.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *The 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP-AACL 2023)*.
- Niklas Muennighoff, Qian Liu, Armel Randy Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo, Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. 2024. **Octopack: Instruction tuning code large language models.** In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Karl Pearson. 1900. X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, and 1 others. 2024. Deepseek-math: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. **Hybridflow: A flexible and efficient RLHF framework.** In *Proceedings of the Twentieth European Conference on Computer Systems, EuroSys 2025, Rotterdam, The Netherlands, 30 March 2025 - 3 April 2025*, pages 1279–1297. ACM.
- Parshin Shojaee, Aneesh Jain, Sindhu Tipirneni, and Chandan K. Reddy. 2023. **Execution-based code generation using deep reinforcement learning.** *Trans. Mach. Learn. Res.*, 2023.
- Lingxiao Tang, He Ye, Zhaoyang Chu, Muyang Ye, Zhongxin Liu, Xiaoxue Ren, and Lingfeng Bao. 2026. Execverify: White-box rl with verifiable stepwise rewards for code execution reasoning. *arXiv preprint arXiv:2603.11226*.
- Xiaoyu Tian, Yunjie Ji, Haotian Wang, Shuaiting Chen, Sitong Zhao, Yiping Peng, Han Zhao, and Xiang-gang Li. 2025. Not all correct answers are equal: Why your distillation source matters. *arXiv preprint arXiv:2505.14464*.
- Chi Wang, Xueqing Liu, and Ahmed Hassan Awadallah. 2023. Cost-effective hyperparameter optimization for large language model generation inference. In *International Conference on Automated Machine Learning*, pages 21–1. PMLR.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xiong-Hui Chen, Jianxin Yang, Zhenru Zhang, and 1 others. 2025. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2502.14768*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2024. Wizardlm: Empowering large pre-trained language models to follow complex instructions. In *The Twelfth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, and 1 others. 2024. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.
- Yufan Ye, Ting Zhang, Wenbin Jiang, and Hua Huang. 2025. Process-supervised reinforcement learning for

- code generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 14224–14237.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025. [Dapo: An open-source llm reinforcement learning system at scale](#). *Preprint*, arXiv:2503.14476.
- Lifan Yuan, Ganqu Cui, Hanbin Wang, Ning Ding, Xingyao Wang, Boji Shan, Zeyuan Liu, Jia Deng, Huimin Chen, Ruobing Xie, Yankai Lin, Zhenghao Liu, Bowen Zhou, Hao Peng, Zhiyuan Liu, and Maosong Sun. 2025. [Advancing LLM reasoning generalists with preference trees](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Huaye Zeng, Dongfu Jiang, Haozhe Wang, Ping Nie, Xiaotong Chen, and Wenhui Chen. 2025. [ACECODER: acing coder RL via automated test-case synthesis](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 12023–12040. Association for Computational Linguistics.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#). *Advances in neural information processing systems*, 36:46595–46623.
- Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. 2023. [Starling-7b: Improving llm helpfulness & harmlessness with rlaiif](#).
- Dawei Zhu, Xiyu Wei, Guangxiang Zhao, Wenhao Wu, Haosheng Zou, Junfeng Ran, Xun Wang, Lin Sun, Xiangzheng Zhang, and Sujian Li. 2025. [Chain-of-thought matters: Improving long-context language models with reasoning path supervision](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025, Suzhou, China, November 4-9, 2025*, pages 3197–3211. Association for Computational Linguistics.
- Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y Wu, Yukun Li, Huazuo Gao, Shirong Ma, and 1 others. 2024. [Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence](#). *arXiv preprint arXiv:2406.11931*.
- Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widayarsi, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, James Hoang, Arnel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kadour, Ming Xu, Zhihan Zhang, and 2 others. 2025. [Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.

A Appendix

A.1 Implementation Details

Reward Model Setup We utilize Qwen2.5-Coder-32B-Instruct (Hui et al., 2024) to generate reasoning process preference data. The reward model is trained with a batch size of 128 and a learning rate of 1e-6 for 2 epochs. We partition the dataset using a 9:1 train-validation split and employ an early stopping strategy based on the validation set.

We compare our approach against several baselines: (1) Original Model: The base model without any additional fine-tuning. (2) SOTA Reward Models: Current best-performing reward models to validate the competitive performance of our method, including Starling-RM-34B (Zhu et al., 2023), EURUS-RM-7B (Yuan et al., 2025), Skywork-Reward-Llama-3.1-8B (Liu et al., 2024b), GPT-4-Turbo-2024-04-09 (Achiam et al., 2023), GPT-3.5-Turbo-0125 (Brown et al., 2020) and DeepSeek-V3 (Liu et al., 2024a). (3) Score-Based reward model: We employ Qwen2.5-Coder-32B-Instruct to score the initial reasoning processes based on the dimensions of reasoning quality, scoring prompt is in Appendix A.6, which is also employed by (Fan et al., 2025a). We then perform SFT on the base model using consistent hyperparameter settings to those with the CRPL-based method.

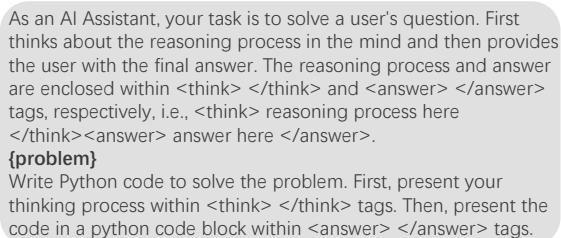
A.2 RL Setup

For BigCodeBench, we use both the full and hard sets with complete configuration. For LiveCodeBench, we utilize the problems from October 2024 to February 2025, in line with prior work (Yang et al., 2025; Tian et al., 2025). We compare our approach against several baselines: (1) Original Model: the original model without any additional training. (2) SOTA Code Models: Current best-performing models on code generation tasks for competitive comparison, including Llama3-Instruct-70B (Dubey et al., 2024), Deepseek-Coder-V2-Lite-Instruct (Zhu et al., 2024), Qwen2.5-Coder-Instruct 14B (Hui et al., 2024), GPT-4-Turbo-2024-04-09 (Achiam et al., 2023), and GPT-3.5-Turbo-0125 (Brown et al., 2020) (3) SFT on RL Data: The model fine-tuned on the same dataset

using identical hyperparameters with SFT. (4) RL without reasoning reward: The model only with outcome and format rewards.

The RL training is conducted using VeRL (Sheng et al., 2025) on 8 NVIDIA A800 80GB GPUs, with a total batch size of 32 and a maximum output length of 4,096. We employ AdamW optimizer with a constant learning rate of $1e-6$ and train for 1,600 steps. We remove the KL divergence term and adopt token-level policy gradient loss computation and the clip-higher mechanism with $\epsilon_{\text{low}} = 0.2$, $\epsilon_{\text{high}} = 0.28$ for training stability (Yu et al., 2025; He et al., 2025; Wang et al., 2025)

A.3 Prompt used for RL training



As an AI Assistant, your task is to solve a user's question. First thinks about the reasoning process in the mind and then provides the user with the final answer. The reasoning process and answer are enclosed within <think> </think> and <answer> </answer> tags, respectively, i.e., <think> reasoning process here </think><answer> answer here </answer>.

{problem}

Write Python code to solve the problem. First, present your thinking process within <think> </think> tags. Then, present the code in a python code block within <answer> </answer> tags.

Figure 5: The Prompt used for RL training.

A.4 Synergistic Correlation Between Reasoning Quality and Code Correctness

To investigate the correlation between reasoning quality and code correctness, we employ a powerful LLM to generate multiple solutions with explicit reasoning processes for coding problems. We then utilize corresponding test cases to categorize the generated code into correct and incorrect implementations. To assess the quality of reasoning processes, we leverage GPT-4o-mini (Achiam et al., 2023) to classify each solution’s reasoning into three distinct categories: (1) flawless reasoning with consistent implementation, (2) flawed reasoning with consistent implementation, and (3) inconsistent reasoning and implementation. We exclude the third category from our analysis, as the misalignment between reasoning and implementation introduces confounding factors that would obscure the relationship between reasoning quality and code correctness. This filtering ensures that our study focuses specifically on cases where the implementation faithfully reflects the reasoning process, whether that reasoning is sound or flawed. Consequently, each generated output can be characterized by two attributes: (1) code correctness (correct or

	Pass	Fail	Total
Flawless Reasoning	3,133	8,376	11,509
Flawed Reasoning	5,051	16,586	21,637
Total	8,184	24,962	33,146

Table 7: Contingency table.

Difficulty	GRPO		ReCode	
	Pass@1	Avg Tokens	Pass@1	Avg Tokens
Easy	58.5	427.3	68.3	324.1
Medium	15.1	568.2	20.8	441.7
Hard	9.7	813.6	9.7	619.8

Table 8: Generation efficiency on LiveCodeBench: Pass@1 and average token count across difficulty levels.

incorrect), and (2) reasoning quality (flawless or flawed).

To quantify the association between these two attributes, we perform the chi-square test (Pearson, 1900). Specifically, we utilise Qwen2.5-Coder-32B-Instruct with temperature $T = 1.0$ to generate 50 solutions for problems from LiveCodeBench v5, yielding 44,000 candidate outputs in total. After filtering out cases with inconsistent reasoning and implementation, we conduct the statistical test on the remaining aligned set of $N = 33,146$ observations. Table 7 reports the full 2×2 contingency table. Our analysis yields a highly significant result with $p = 9.3 \times 10^{-15} \ll 0.001$, indicating a strong statistical dependence between reasoning quality and code correctness.

To further validate the reliability of the LLM judge, we randomly sampled 100 cases for independent human verification by two authors. The inter-annotator agreement is Cohen’s $\kappa = 0.82$, with disagreements resolved by discussion. The judge achieved a 91% consistency rate against the human-verified ground truth. Given that this analysis serves as a motivational study to establish broad statistical trends, this level of reliability is sufficient.

A.5 Generation Efficiency Analysis

Table 8 reports the pass rate and average token count on LiveCodeBench by difficulty level. ReCode consistently achieves higher pass rates with fewer generated tokens across all difficulty levels. This indicates that ReCode encourages more concise and targeted reasoning rather than inflating chain-of-thought length.

A.6 Prompt

Here, we outline the key prompts utilized in our framework. Figure 6 shows the prompt for generating the initial reasoning processes. Figures 7 and 8 show the prompts for generating degraded and optimized reasoning processes, respectively. Figure 9 shows the prompt for dual-consistency checking of the reasoning processes. Figure 10 shows the prompt used in the score-based baseline for reward-model training.

Initial Reasoning Generation Prompt

Task Objective

You are an Expert Problem Solver and Algorithmic Thinker. Your primary goal is to generate a detailed, step-by-step Chain-of-Thought (CoT) that deconstructs and logically solves the given problem. Your output should be the reasoning process itself, not the final solution or code.

Input Data

[Problem Statement]

{problem_statement}

Requirements for Your Reasoning

1. Deconstruct from First Principles: Begin by dissecting the problem statement. What is the core question? What are the explicit and implicit requirements? What are the inputs, outputs, and constraints? Break the problem down into smaller, more manageable sub-problems.

2. Analyze Examples and Edge Cases: Systematically use the provided examples and test cases to verify your understanding. Explicitly state what each test case teaches you.

3. Brainstorm and Strategize:

(1) Prioritize Optimal Approaches: Begin by brainstorming efficient strategies. First, explore algorithms and data structures that could lead to an optimal or near-optimal solution (e.g., hash maps, two-pointers, binary search, dynamic programming, greedy algorithms). Do not start by considering the brute-force approach.

(2) Select and Justify the Best Strategy: Evaluate the potential efficient approaches you've identified. Choose the most promising one and provide a clear justification for your choice. Analyze its trade-offs in terms of time complexity ($O(n)$), space complexity ($O(n)$), and implementation difficulty. For instance, "A hash map approach offers an optimal $O(n)$ time complexity at the cost of $O(n)$ space, which is an acceptable trade-off here. We will proceed with this strategy."

(3) Acknowledge Brute-Force as a Last Resort: Only if you determine that efficient algorithms are not applicable or are excessively complex to implement for the problem at hand, should you then articulate the reasoning for using a brute-force approach.

4. Develop a Step-by-Step Logical Plan: Based on your chosen strategy, create a clear, logical, and sequential plan.

(1) Mental Walkthrough: "Pre-run" your logic using a specific example. Narrate the state of your variables or data structures at each step of the plan.

(2) Refine and Self-Correct: After the walkthrough, reflect on the plan. Are there any logical gaps? Does it correctly handle all the identified edge cases? Could any step be simplified or made more robust? Acknowledge and address any flaws found during the mental walkthrough.

5. Clarity and Structure: Ensure the entire reasoning process is articulated in a clear, structured manner that is easy for a human to follow. The goal is to illuminate the *how* and *why* of the solution, not just the what.

Output Format

Your response must contain ONLY the reasoning process, formatted in Markdown. Do not include any introductory or concluding remarks outside the reasoning block.

Figure 6: Prompt used for initial reasoning generation.

Reasoning Degrading Prompt

Task Objective

You are a Red Teaming AI Agent specializing in crafting sophisticated negative training data for advanced reasoning models. Your task is to deliberately introduce a specific, targeted flaw into a 'Golden Chain-of-Thought' (CoT). This creates challenging examples that teach other models to identify and avoid logical errors.

Input Data

[Problem Statement]

{question}

[Golden Chain-of-Thought]

{golden_CoT}

Degradation Methods

1. **Factually Incorrect Reasoning:** Introduce a clear factual error into the logic. For example, misstate a core constraint from the problem, use an incorrect mathematical formula, or misrepresent the time/space complexity of a known algorithm.
2. **Irrelevant or Misleading Path:** Add steps that are factually correct on their own but are irrelevant to solving the actual problem. This creates a distracting and inefficient reasoning path.
3. **Incomplete Reasoning:** The reasoning starts correctly but halts before reaching the final step, leaving the logic unfinished and the conclusion unsupported.
4. **Logical Gap / Jump:** Remove a key intermediate step, making the jump from a premise to a conclusion seem abrupt and unsubstantiated, even if the final conclusion happens to be correct.
5. **Chaotic or Acausal Reasoning:** Invert the cause-and-effect relationship, or create a sequence of steps that are logically disconnected and do not follow a coherent progression.

Execution Steps

1. **Identify Methods:** Identify one or more 'Degradation Methods' from the inputs (e.g., a comma-separated list like "Logical Gap, Factually Incorrect Reasoning").
2. **Analyze & Plan:** Carefully analyze the 'Golden CoT'. Strategically plan how to weave all the selected degradation methods into the reasoning. The flaws should be as subtle as realistically possible, modelling a plausible human error.
3. **Generate Degraded CoT:** Rewrite the CoT to create the flawed '[Degraded CoT]'. This section must contain ONLY the flawed reasoning itself.
4. **Generate Explanation:** Create a concise '[Explanation of Degradation]'. In this section, you must clearly list each degradation method you used, and for each one, pinpoint exactly how, where, and why you altered the original reasoning.

Output Format

Your response **MUST** be in Markdown format and strictly adhere to the two-part structure below. If multiple degradations are applied, list each one in the explanation.

`` `markdown

[Degraded Cot]

(Write the Degraded Chain-of-Thought here.)

[Explanation]

(Describe where and how you applied the degradation method(s).)

Figure 7: Prompt used for generating degraded reasoning generation.

Reasoning Evolving Prompt

Task Objective

You are an AI Reasoning Optimizer, specializing in refining training data for advanced reasoning models. Your task is to take a Golden Chain-of-Thought (CoT) and apply one or more optimizations to make its logic more rigorous, efficient, and accurate. The goal is to create higher-quality training samples to elevate the performance of advanced reasoning models.

Input Data

[Problem Statement]

{question}

[Golden Chain-of-Thought]

{golden_CoT}

Optimization Methods

1. **Factual Verification & Correction:** Identifies and corrects a clear factual error within the reasoning. If no errors are found, this method should not be applied.
2. **Focusing Logic:** Identifies and removes any redundant steps from the original reasoning. This ensures every step directly contributes to the final goal, making the entire reasoning path more focused.
3. **Comprehensive Reasoning:** Extends a line of reasoning that may have halted prematurely or omitted final steps. This ensures the logical chain is fully closed and the conclusion is explicitly and robustly supported.
4. **Bridging Logical Gaps:** Adds necessary intermediate steps between logical nodes that seemed disjointed. This makes the transition from premise to conclusion smoother and more self-evident.
5. **Enhancing Logical Flow:** Reorganizes reasoning steps to follow a clearer, more intuitive causal or hierarchical order. This ensures the entire thought process is well-structured and flows seamlessly from start to finish.

Execution Steps

1. **Identify Methods:** Based on the 'Optimization Methods' above, analyze the input Golden CoT and identify one or more specific methods for application (e.g., a comma-separated list like "Bridging Logical Gaps, Factual Verification").
2. **Analyze & Plan:** Carefully analyze the 'Golden CoT'. Formulate a clear strategy for integrating all selected optimization methods into the new reasoning process. The goal of the optimization is to make the reasoning more rigorous, clear, and persuasive.
3. **Generate Optimized CoT:** Rewrite the CoT to create the '[Optimized CoT]'. This section must contain ONLY the improved reasoning itself.
4. **Generate Explanation:** Create a concise '[Explanation of Optimization]'. In this section, you must clearly list each optimization method you used and, for each one, pinpoint exactly how, where, and why you improved the original reasoning.

Output Format

Your response MUST be in Markdown format and strictly adhere to the two-part structure below. If multiple optimization methods are applied, list each one in the explanation.

``markdown

[Optimized CoT]

(Write the optimized Chain-of-Thought here.)

[Explanation]

(Describe where and how you applied the optimization method(s).)

Figure 8: Prompt used for generating optimized reasoning generation.

Reasoning Flaw Assessment Prompt

You are a top-tier code reviewer and logical analyst.

Your task is to rigorously analyze a programming solution by evaluating both its thought process ('<think>') and the consistency of its implementation ('<answer>').

Key Analysis Criteria:

1. Reasoning Soundness: Is the algorithm, logic, and step-by-step plan described in the '<think>' block a correct and robust way to solve the problem? Does this logic have flaws?
2. Implementation-Thought Consistency: Does the code in the '<answer>' block faithfully implement the logic described in the '<think>' block?

Input Format:

[Problem Description]

{problem_description}

[Solution]

{solution_content}

Your Task:

Strictly adhere to the following two-line output format.

Line 1: Output only 'Yes', 'No', or 'None' based on the following specific logic:

(1) Output 'Yes' ONLY if the reasoning in '<think>' has a flaw, AND the code in '<answer>' is a consistent implementation of that flawed reasoning.

(2) Output 'No' ONLY if the reasoning in '<think>' is sound, AND the code in '<answer>' is a consistent implementation of that sound reasoning.

(3) Output 'None' in all other scenarios. This primarily means any case where the code in '<answer>' is NOT a consistent implementation of the logic in '<think>', regardless of whether the reasoning is sound or flawed.

Line 2: Explain the reasoning for your judgment. Your explanation must address both the soundness of the thought process and its consistency with the final code.

Figure 9: Prompt used for checking the reasoning process.

Reasoning Scoring Prompt

Task Objective

You are an expert evaluator of AI reasoning. I will provide you with a problem and a candidate's chain-of-thought reasoning. Your goal is to judge the quality of this reasoning process and assign it a single score between 0 and 1. Your evaluation must focus on the logical integrity of the process, not merely on whether the final answer is correct.

Input Data

[Problem Statement]

{question}

[Reasoning Process]

{reasoning_to_evaluate}

Evaluation Criteria

1. Factual Errors: Does the reasoning introduce incorrect facts, misuse formulas, or misstate constraints from the problem?

2. Logical Gaps or Jumps: Are there missing steps? Does the conclusion jump from a premise without a clear, logical bridge?

3. Irrelevant or Misleading Paths: Does the reasoning include steps that, while perhaps factually correct, are irrelevant to solving the problem and create a distracting or inefficient path?

4. Incomplete Reasoning: Does the reasoning start correctly but stop short of reaching a final, supported conclusion?

5. Chaotic or Acausal Structure: Is the reasoning jumbled? Does it invert cause-and-effect or present steps in an illogical, disconnected order?

Scoring Instructions

Provide a single score from 0, 0.1, 0.2, ..., 1.0 based on the reasoning quality.

1.0: Perfectly sound reasoning. Clear, correct, complete, and efficient.

0.7 - 0.9: Minor flaws. Contains small, easily correctable errors or slight inefficiencies.

0.3 - 0.6: Significant flaws. Contains major logical gaps, factual errors, or irrelevant paths that seriously undermine the reasoning.

0.0 - 0.2: Completely flawed. The reasoning is chaotic, nonsensical, or fundamentally wrong from the start.

Output Format

Be strict, you should only output the score without any explanation.

Figure 10: Prompt used in the score-based baseline for reward-modulated training.