

# CoVerRL: Breaking the Consensus Trap in Label-Free Reasoning via Generator-Verifier Co-Evolution

Teng Pan<sup>1,2</sup>, Yuchen Yan<sup>1</sup>, Zixuan Wang<sup>1,2</sup>, Ruiqing Zhang<sup>2</sup>, Guiyang Hou<sup>1</sup>,  
Wenqi Zhang<sup>1</sup>, Weiming Lu<sup>1\*</sup>, Jun Xiao<sup>1</sup>, Yongliang Shen<sup>1\*</sup>,

<sup>1</sup>Zhejiang University, <sup>2</sup>Baidu Inc.  
{pt6, syl}@zju.edu.cn

## Abstract

Label-free reinforcement learning enables large language models to improve reasoning capabilities without ground-truth supervision, typically by treating majority-voted answers as pseudo-labels. However, we identify a critical failure mode: as training maximizes self-consistency, output diversity collapses, causing the model to confidently reinforce systematic errors that evade detection. We term this the consensus trap. To escape it, we propose CoVerRL, a framework where a single model alternates between generator and verifier roles, with each capability bootstrapping the other. Majority voting provides noisy but informative supervision for training the verifier, while the improving verifier progressively filters self-consistent errors from pseudo-labels. This co-evolution creates a virtuous cycle that maintains high reward accuracy throughout training. Experiments across Qwen and Llama model families demonstrate that CoVerRL outperforms label-free baselines by 4.7-5.9% on mathematical reasoning benchmarks. Moreover, self-verification accuracy improves from around 55% to over 85%, confirming that both capabilities genuinely co-evolve.

## 1 Introduction

Large Language Models (LLMs) (DeepSeek-AI et al., 2025; OpenAI et al., 2024) have demonstrated remarkable capabilities in complex reasoning tasks, especially with recent advances in Reinforcement Learning with Verifiable Rewards (RLVR) that further amplify these abilities by optimizing models against ground-truth signals. However, the standard RLVR paradigm is heavily based on large amounts of high-quality labeled data or ground-truth verifiers, which are often unavailable or expensive to obtain. Consequently, enabling LLMs to self-evolve in a *label-free* setting has emerged as a critical frontier in current research.

\*Corresponding author

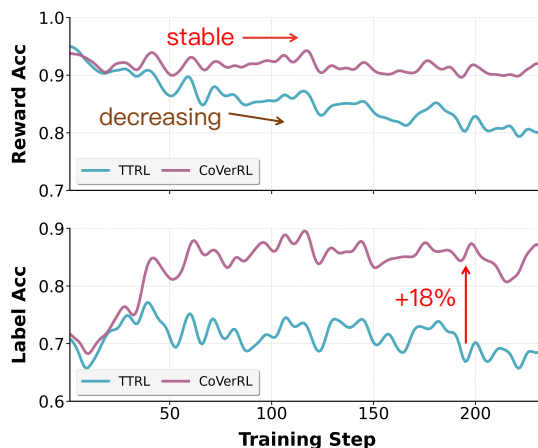


Figure 1: Training dynamics of TTRL and CoVerRL on Qwen3-1.7B-Base. CoVerRL maintains high reward accuracy while steadily improving label accuracy, whereas TTRL suffers from reward accuracy degradation as output diversity collapses.

The dominant paradigm for label-free training employs majority voting as a proxy for correctness (Prasad et al., 2025; Zuo et al., 2025; Liu et al., 2025a). Given a problem, the model generates multiple reasoning trajectories, and the most frequent answer becomes the pseudo-label for reward computation. The underlying assumption is that consistent answers across diverse reasoning paths likely reflect correct solutions. Notably, test-time reinforcement learning (TTRL) (Zuo et al., 2025) reveal that majority voting remains effective even when pseudo-labels are incorrect through a “lucky hit” mechanism: samples whose answers differ from an incorrect pseudo-label still receive correct negative rewards. This explains why reward accuracy can remain high despite low label accuracy, providing viable training signals even on challenging benchmarks.

However, we identify a critical instability in this paradigm. Since rewards are computed against the majority answer, the training objective inherently

encourages the model to produce more consistent outputs. As training progresses, output diversity systematically decreases and predictions concentrate around the majority answer. When this answer is incorrect, a destructive feedback loop emerges: the model confidently reinforces its own errors, diversity collapses further, and reward accuracy degrades. Figure 1 provides empirical evidence: under TTRL, reward accuracy begins near 95% but steadily drops to approximately 80%, while label accuracy stagnates around 70% with no improvement. Once the model converges to systematic errors with high confidence, pure consensus methods offer no mechanism for recovery. This reveals a fundamental limitation: majority voting captures consistency but cannot detect correctness. When models repeatedly produce identical wrong answers, a phenomenon known as self-consistent errors (Tan et al., 2025), consensus-based rewards actively reinforce these mistakes.

A natural remedy is self-verification: if the model could judge correctness, self-consistent errors could be filtered before corrupting training. However, prior work has shown that LLMs without specific training cannot reliably detect their own reasoning errors (Huang et al., 2024; Tyen et al., 2024). Expecting a model to verify solutions it cannot correctly generate creates a circular dependency that severely limits self-correction without external supervision. The situation appears intractable: consensus provides training signal but cannot escape systematic errors, while verification could detect errors but cannot be trained without ground truth.

We resolve this tension through a key observation: majority voting and self-verification, while individually insufficient, can bootstrap each other’s improvement. We propose **CoVerRL** (Co-evolving Verifier and generator via Reinforcement Learning), a framework where a single model alternates between generator and verifier roles within a unified multi-turn RL process. The generator produces candidate solutions and establishes pseudo-labels via majority voting; the verifier scrutinizes reasoning processes to identify errors independent of answer frequency. The critical insight is bidirectional supervision: consensus-derived pseudo-labels provide contrastive signal for verifier training, while the improving verifier filters self-consistent errors from pseudo-labels. This creates a virtuous cycle where better verification yields cleaner labels, producing more accurate rewards that further improve both generation and verifica-

tion. As shown in Figure 1, CoVerRL maintains reward accuracy above 90% throughout training, while label accuracy steadily increases from 75% to nearly 88%, demonstrating the verifier’s growing ability to escape the failure mode that plagues pure consensus methods.

We evaluate CoVerRL across three models on four reasoning benchmarks. CoVerRL consistently outperforms TTRL, achieving average improvements of 5.7%, 5.9%, and 4.7% respectively across model scales. Beyond generation accuracy, CoVerRL produces capable verifiers as a byproduct: self-verification accuracy improves dramatically, from 55.8% to 81.2% on Qwen3-1.7B and from 54.0% to 86.5% on Qwen2.5-7B. This confirms that the two capabilities genuinely co-evolve rather than one merely supporting the other.

We summarize our contributions as follows: (1) We identify the consensus trap in majority voting based label-free RL, where diversity collapse causes reward accuracy degradation as models become overconfident in systematic errors, explaining why such methods eventually stagnate. (2) We propose CoVerRL, a co-evolution framework that unifies generation and verification into a multi-turn RL process, enabling mutual bootstrapping where each capability supervises improvement of the other without external labels. (3) We validate CoVerRL across Qwen and Llama model families, demonstrating 4-6% improvements over label-free baselines on mathematical reasoning benchmarks while producing verifiers that generalize well to held-out evaluation.

## 2 Related Work

**LLM Self-Evolution** To achieve label-free evolution, existing research has exploited the intrinsic confidence of LLMs as a supervision signal (Prabhudesai et al., 2025; Zhao et al., 2025b; Anonymous, 2025; Li et al., 2025), leveraged self-consistency through majority voting to establish pseudo-labels (Prasad et al., 2025; Zuo et al., 2025; Liu et al., 2025a; Zhou et al., 2025; Yu et al., 2025b), or employed self-verification mechanisms (Yuan et al., 2025). Another paradigm is self-play, where a single model simulates multiple roles and generates its own training data (Zhao et al., 2025a; Huang et al., 2025; Xia et al., 2025; Yang et al., 2025b; Guo et al., 2025; Wei et al., 2025). In this work, we leverage both majority voting and self-verification to co-evolve the generator

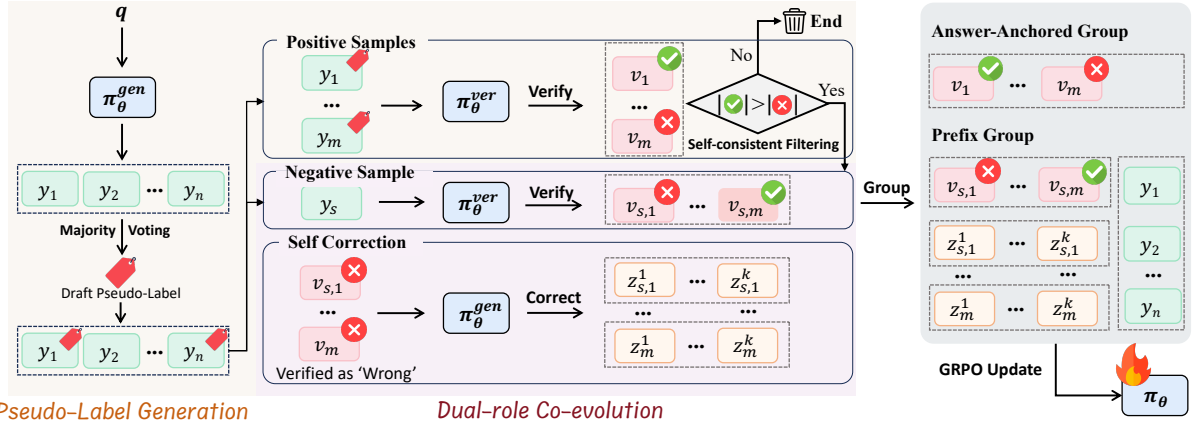


Figure 2: Overview of the CoVerRL framework. The pipeline illustrates the Pseudo-Label Generation process and the Online Dual-Role Co-Evolution strategy, facilitating a mutual bootstrapping process where generation and verification capabilities jointly evolve.

and verifier, transcending the limitations of consensus trap and unreliable standalone verification.

**LLM Self-Verification and Self-Correction** Recent advances have demonstrated that LLMs are capable of self-refinement when provided with external feedback (Chen et al., 2023; Liu et al., 2025c; Jain et al., 2024; Liu et al., 2025b). Concurrently, Several studies train standalone verifier models to bridge the verification gap (Cobbe et al., 2021; Lightman et al., 2023; Zhang et al., 2025a,b). Beyond relying on external signals, many studies focus on intrinsic self-correction (Miao et al., 2023; Lee et al., 2025; Ma et al., 2025; Xiong et al., 2025), which predominantly relies on either prompt engineering or two-stage paradigms combining Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL). With the rapid progress in multi-turn RL (Feng et al., 2025a; Jin et al., 2025; Wang et al., 2025; Feng et al., 2025b), recent studies have investigated training self-verification and self-correction capabilities directly through a single reinforcement learning stage (Liu et al., 2025d; Zha et al., 2025; Jiang et al., 2025). In our work, we focus on enhancing self-verification and self-correction capabilities in a label-free paradigm, which has been rarely explored in previous studies.

### 3 Method

In this section, we present CoVerRL, a framework that enables label-free reinforcement learning through the co-evolution of generation and verification capabilities. The key principle of CoVerRL is that a single policy model  $\pi_\theta$  acts as two roles: a generator  $\pi_\theta^{\text{gen}}$  and a verifier  $\pi_\theta^{\text{ver}}$ . Through

their interaction, the model produces high-quality pseudo-labels (Section 3.1) and achieves coordinated dual-role co-evolution (Section 3.2).

#### 3.1 Pseudo-Label Generation

The key challenge in label-free RL is obtaining reliable supervision signals without ground truth. Majority voting suffers from diversity collapse, while self-verification lacks training signal. CoVerRL addresses this by combining both mechanisms: majority voting drafts initial pseudo-labels, and self-verification filters out self-consistent errors.

**Majority Voting.** Given a query  $q$ , the model acting as a generator samples  $N$  reasoning paths  $y_i \sim \pi_\theta^{\text{gen}}(\cdot | q)$ , forming a collection of query-response pairs  $Y = \{(q, y_i)\}_{i=1}^N$ . Following Zuo et al. (2025), we group these outputs and identify the most frequent answer as the draft pseudo-label  $\hat{y}$ . This partitions the outputs into a positive set  $Y^+ = \{(q, y_i) \in Y : \text{ans}(y_i) = \hat{y}\}$  and a negative set  $Y^- = \{(q, y_i) \in Y : \text{ans}(y_i) \neq \hat{y}\}$ .

**Self-consistent Verification Filtering.** To filter self-consistent errors, the model switches to its verifier role  $\pi_\theta^{\text{ver}}$  to scrutinize the candidate solutions. We define a maximum verification budget  $M$  to maintain computational efficiency; if  $|Y^+| > M$ , we randomly sample  $M$  tuples from  $Y^+$  to form the verification set, otherwise include all tuples in  $Y^+$ . For each selected tuple  $(q, y_j)$ , the verifier produces a binary judgment  $v_j \sim \pi_\theta^{\text{ver}}(\cdot | q, y_j)$ , where  $v_j \in \{0, 1\}$ , and form the verification set  $\mathcal{V}^+ = \{(q, y_j, v_j)\}_{j=1}^{|\mathcal{V}^+|}$ , where  $|\mathcal{V}^+| = \min(|Y^+|, M)$ . We then apply a consensus filter: the query  $q$  and its associated sets are retained for training only if

the majority of verifications in  $\mathcal{V}^+$  are positive:

$$\sum_{j=1}^{|\mathcal{V}^+|} v_j > \frac{|\mathcal{V}^+|}{2} \quad (1)$$

This filtering mechanism specifically targets the failure mode where the generator produces self-consistent errors; by leveraging  $\pi_{\theta}^{\text{ver}}$  to identify flawed reasoning, we exclude such samples before they corrupt the training process.

### 3.2 Online Dual-Role Co-Evolution

With filtered pseudo-labels established, we now describe how both roles are jointly optimized through multi-turn reinforcement learning.

**Contrastive Verifier Training.** To train the verifier without ground-truth labels, we construct contrastive examples from the derived pseudo-labels. The positive verification set  $\mathcal{V}^+$  (targeting the majority answer  $\hat{y}$ ) is obtained during the filtering stage. For negative examples, we select the answer  $y_s$  with the lowest frequency from  $\mathcal{Y}^-$  and generate  $|\mathcal{V}^+|$  verification paths to form the negative verification set  $\mathcal{V}^- = \{(q, y_s, v_j)\}_{j=1}^{|\mathcal{V}^+|} \sim \pi_{\theta}^{\text{ver}}(\cdot | q, y_s)$ , where each  $v_j$  targets the voted ‘‘Incorrect’’ status. Enforcing  $|\mathcal{V}^-| = |\mathcal{V}^+|$  ensures balanced training, which we show in Appendix E is equivalent to implicit preference optimization and induces beneficial curriculum effects.

**Self-Correction.** Beyond verification, we equip the model with error recovery capabilities by leveraging the interaction between its roles. For any triplet  $(q, y, v) \in \mathcal{V}^+ \cup \mathcal{V}^-$  where the verifier identifies an error ( $v = 0$ ), the model switches back to its generator role  $\pi_{\theta}^{\text{gen}}$ . It is then triggered to produce  $K$  revised solutions  $\{z_k\}_{k=1}^K \sim \pi_{\theta}^{\text{gen}}(\cdot | q, y, v)$ , conditioned on the query, the initial failed attempt, and the verifier’s feedback. This structure allows the model to learn self-correction from its own internal critiques, extending the label-free paradigm to iterative refinement.

**Reward Design.** We define a reward function  $r = r^f + r^a$  combining format and accuracy components. The format reward  $r^f \in \{0, 1\}$  ensures outputs adhere to structural constraints. The accuracy reward is:

$$r^a = \mathbb{I}(o = \hat{o}) \quad (2)$$

where  $o$  is the model’s output and  $\hat{o}$  is the corresponding pseudo-label. For generation,  $\hat{o}$  is the

consensus answer  $\hat{y}$ ; for verification,  $\hat{o}$  is the target judgment (‘‘Correct’’ for  $\mathcal{V}^+$ , ‘‘Incorrect’’ for  $\mathcal{V}^-$ ).

**Answer-Anchored GRPO.** We optimize using Group Relative Policy Optimization (GRPO) (Shao et al., 2024) with a modification for our multi-turn structure. Standard GRPO groups responses by query prefix. In our case, although verifications in  $\mathcal{V}^+ = \{(q, y_j, v_j)\}_{j=1}^M$  do not strictly originate from the same prefix, every  $y_j$  targets the same answer  $\hat{y}$ . We thus propose **Answer-Anchored GRPO**: we group these  $|\mathcal{V}^+|$  paths together as they share the same answer anchor  $\hat{y}$ ; for generation and  $\mathcal{V}^-$ , we revert to standard prefix-based grouping. The advantage is computed as:

$$A_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})} \quad (3)$$

where  $\mathbf{r}$  contains rewards within the appropriate group. This hybrid strategy isolates verification variance while maintaining stable generation optimization.

**The Co-Evolution Cycle.** The complete training loop proceeds as follows: (1) generate solutions and draft pseudo-labels via majority voting, (2) verify positive samples and filter self-consistent errors, (3) construct contrastive verification data, (4) collect self-correction trajectories, and (5) update the model with Answer-Anchored GRPO on all roles jointly. As training progresses, the verifier improves through contrastive learning, enabling better filtering of self-consistent errors. Cleaner pseudo-labels yield more accurate rewards, improving the generator. The improved generator produces higher-quality consensus labels, further benefiting verifier training. This bidirectional bootstrapping, visualized in Figure 2, enables both roles to co-evolve without external supervision. The detailed algorithm flow is presented in Appendix D.

## 4 Experiments

### 4.1 Environmental Setup

**Base Models.** To verify the effectiveness of CoVerRL on diverse model architectures, we adopt three representative models from the Qwen and Llama families, including base and instruct models: Llama-3.2-3B-Instruct (Grattafiori et al., 2024), Qwen2.5-7B (Qwen et al., 2025) and Qwen3-1.7B-Base (Yang et al., 2025a) with thinking mode.

Model	Method	MATH500	AMC	AIME24	GPQA	Average
Qwen3-1.7B-Base	Base Model	53.5 / 53.3	24.6 / 24.5	3.8 / 3.3	27.5 / 27.3	27.4 / 27.1
	TTRL	65.1 / 65.0	31.1 / 30.9	5.2 / 5.2	30.9 / 30.7	33.1 / 33.0
	CoVerRL (Ours)	<b>69.0 / 71.9</b>	<b>36.0 / 38.6</b>	<b>9.8 / 10.6</b>	<b>32.9 / 33.6</b>	<b>36.9 / 38.7</b>
	Δ	+3.9 / +6.9	+4.9 / +7.7	+4.6 / +5.4	+2.0 / +2.9	+3.8 / +5.7
Llama-3.2-3B-Instruct	Base Model	42.7 / 41.0	17.0 / 15.7	4.6 / 5.0	26.9 / 26.1	22.8 / 22.0
	TTRL	52.6 / 52.2	23.8 / 23.3	13.8 / 14.0	29.8 / 28.2	30.0 / 29.4
	CoVerRL (Ours)	<b>55.9 / 59.3</b>	<b>28.3 / 32.2</b>	<b>16.3 / 16.9</b>	<b>32.3 / 32.6</b>	<b>33.2 / 35.3</b>
	Δ	+3.3 / +7.1	+4.5 / +8.9	+2.5 / +2.9	+2.5 / +4.4	+3.2 / +5.9
Qwen2.5-7B	Base Model	50.1 / 51.4	25.5 / 26.4	5.2 / 6.5	29.9 / 29.7	27.7 / 28.5
	TTRL	73.8 / 74.2	42.2 / 42.2	12.7 / 12.5	35.8 / 35.6	41.1 / 41.1
	CoVerRL (Ours)	<b>76.8 / 79.6</b>	<b>47.6 / 49.2</b>	<b>14.6 / 17.1</b>	<b>36.2 / 37.2</b>	<b>43.8 / 45.8</b>
	Δ	+3.0 / +5.4	+5.4 / +7.0	+1.9 / +4.6	+0.4 / +1.6	+2.7 / +4.7

Table 1: Main results on mathematical reasoning and out-of-domain benchmarks. Each cell displays Acc.@first / Acc.@final. Δ indicates the improvement of CoVerRL over TTRL.

**Training and Evaluation Data.** To test our method at scale, we use the large, standard MATH (Hendrycks et al., 2021) training set. For evaluation, we employ three in-domain mathematical benchmarks to assess reasoning capabilities: MATH500 (Hendrycks et al., 2021), AMC (Li et al., 2024), and AIME24 (Li et al., 2024). Additionally, we include the GPQA-Diamond (Rein et al., 2023) to evaluate the models’ generalization on challenging out-of-domain question-answering tasks.

**Baseline and Implementation Details.** We mainly compare our method against TTRL (Zuo et al., 2025). For the TTRL baseline, we follow the practice in (Zuo et al., 2025) by sampling 64 responses for majority voting and downsampling 32 responses for training. For our CoVerRL, at the first turn, we sample 32 responses for majority voting and downsample them to 16 for training; In later self-verification and self-correction, we utilize all samples for training to maximize data efficiency. About detailed description of the hyperparameters, please refer to Appendix C.

**Evaluation Protocol.** Following Jiang et al. (2025), we adopt a sequential inference strategy: the model first generates a solution, then verifies its correctness. If verification passes, inference terminates; otherwise, the model attempts a revision. We report both **Acc.@first** (first-attempt accuracy, equivalent to Pass@1) and **Acc.@final** (accuracy after potential self-correction), averaged over 16 different generations.

## 4.2 Main Results

Table 1 presents the main results. CoVerRL consistently outperforms TTRL across all models and

benchmarks, achieving average improvements of 5.7%, 5.9%, and 4.7% in Acc.@final for the three models respectively.

**Improved Generation Quality.** CoVerRL improves Acc.@first over TTRL by 3.8%, 3.2%, and 2.7% on average across the three models, demonstrating that the co-evolution mechanism enhances the generator itself, not merely the correction process. This confirms that filtering self-consistent errors produces cleaner training signals that benefit generation quality directly.

**Effective Self-Correction.** The gap between Acc.@first and Acc.@final widens substantially under CoVerRL compared to TTRL. On AMC, CoVerRL achieves Acc.@final improvements of 7.7%, 8.9%, and 7.0% over TTRL, significantly larger than the Acc.@first gains. This indicates that the verifier trained through co-evolution can reliably identify errors and guide effective corrections, whereas TTRL’s self-correction provides minimal benefit due to unreliable verification.

Model	MATH500	AMC	AIME24	GPQA
Qwen3-1.7B-Base	55.8	28.5	11.7	38.7
+ CoVerRL	<b>81.2</b>	<b>64.4</b>	<b>48.9</b>	<b>57.9</b>
Llama-3.2-3B-Instruct	57.0	50.9	55.8	<b>56.9</b>
+ CoVerRL	<b>77.2</b>	<b>76.2</b>	<b>75.6</b>	55.2
Qwen2.5-7B	54.0	30.9	12.9	35.7
+ CoVerRL	<b>86.5</b>	<b>70.9</b>	<b>58.8</b>	<b>58.4</b>

Table 2: Self-verification accuracy on four benchmarks.

**Out-of-Domain Generalization.** On GPQA-Diamond, a challenging out-of-domain benchmark,

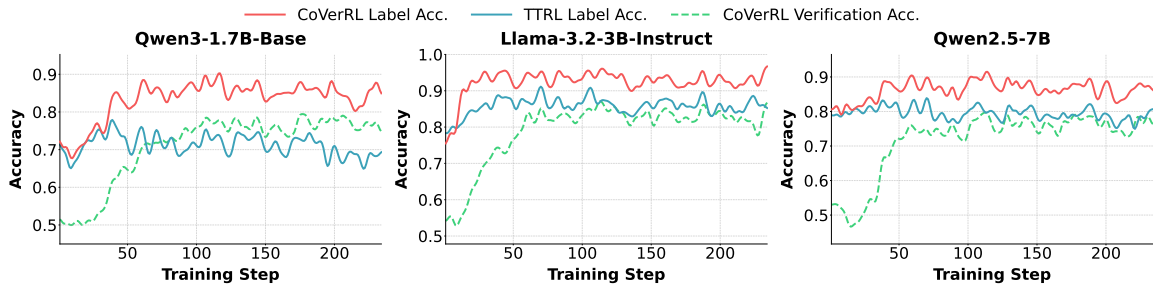


Figure 3: Co-evolution dynamics across three model backbones. As training progresses, verification accuracy improves through contrastive learning, which in turn enables better filtering of self-consistent errors and leads to higher label accuracy than TTRL.

CoVerRL maintains consistent improvements. Notably, while base models and TTRL often suffer degradation during revision ( $\text{Acc.}@_{\text{final}} < \text{Acc.}@_{\text{first}}$ ), CoVerRL reverses this trend across all models. This suggests that the verification capability learned through co-evolution transfers beyond the training distribution.

**Verifier Quality.** Table 2 reports self-verification accuracy. CoVerRL dramatically improves verification capability: Qwen3-1.7B-Base improves from 28.5% to 64.4% on AMC, and Qwen2.5-7B improves from 54.0% to 86.5% on MATH500. These gains confirm that both generation and verification genuinely co-evolve rather than one merely supporting the other.

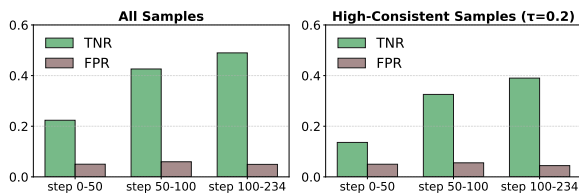


Figure 4: Evolution of the verification capability of the Qwen3-1.7B-Base. We report the True Negative Rate (TNR) and the False Positive Rate (FPR).

### 4.3 Analysis

#### Verifier and Generator Co-Evolve Mutually.

Figure 3 reveals the co-evolution dynamics across three models. In early training, verifier accuracy (green dashed line) starts low due to the lack of verification-specific supervision, resulting in label accuracy comparable to TTRL. As training progresses, the verifier improves by learning from pseudo-labels that, while noisy, provide contrastive signal between majority and minority answers. This improved verification capability enables better filtering of self-consistent errors. Consequently,

CoVerRL’s label accuracy (red line) steadily increases and eventually surpasses TTRL (blue line) by a significant margin across all three model backbones. The consistent pattern across different architectures and scales confirms that both capabilities bootstrap each other through the bidirectional supervision mechanism rather than one merely supporting the other. For a clearer demonstration of the verifier’s evolution, Figure 4 reports the True Negative Rate (TNR), the ratio of incorrect samples correctly filtered, and the False Positive Rate (FPR), the proportion of correct samples mistakenly discarded. The left panel shows that TNR steadily climbs from approximately 22% to nearly 50%, while the FPR remains stable at around 5%, indicating the verifier increasingly prunes erroneous paths while maintaining a minimal false-alarm rate for correct solutions.



Figure 5: Reward accuracy across generation, verification, and self-correction roles during training on Llama-3.2-3B-Instruct. All three roles maintain reward accuracy above 90% throughout training.

#### Reward Accuracy Remains High Across All Roles.

Figure 5 tracks reward accuracy for generation, verification, and self-correction throughout training on Llama-3.2-3B-Instruct. All three

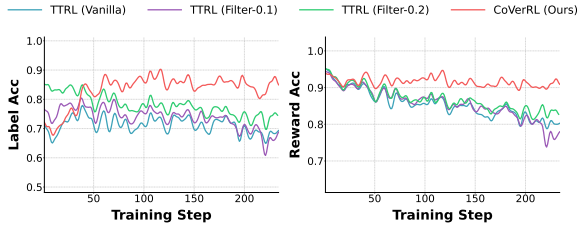


Figure 6: Comparison of CoVerRL against naive consistency filtering on Qwen3-1.7B-Base. Naive filtering ( $\tau \in \{0.1, 0.2\}$ ) improves label accuracy initially but converges to TTRL baseline, while failing to improve reward accuracy. CoVerRL achieves sustained gains on both metrics.

roles maintain reward accuracy above 90%, indicating that the pseudo-labels derived from the co-evolution process provide reliable supervision signals for each role. The sustained high reward accuracy across all roles demonstrates that the bidirectional bootstrapping mechanism successfully prevents the degradation identified in Section 1, yielding both a marked improvement in Acc.@first and a consistently higher Acc.@final.

**Verification Filtering Outperforms Naive Consistency Filtering.** Since samples with lower majority ratios are more likely to be incorrect, a natural alternative to verification-based filtering is to simply discard samples below a consistency threshold. We compare CoVerRL against TTRL with naive filtering thresholds  $\tau \in \{0.1, 0.2\}$  in Figure 6. The left panel shows label accuracy dynamics: while naive filtering initially improves label accuracy by removing low-confidence predictions, its advantage diminishes as training progresses and model outputs become increasingly consistent due to the RL objective, eventually converging to vanilla TTRL performance. In contrast, CoVerRL’s label accuracy continues to improve throughout training, reaching approximately 90% compared to TTRL’s 70%. The right panel reveals a more fundamental distinction in reward accuracy: naive filtering provides no improvement over vanilla TTRL because it cannot detect errors that happen to achieve high consensus. CoVerRL maintains reward accuracy above 90% by specifically targeting self-consistent errors through reasoning-based verification. As shown in the right panel of Figure 4, for samples with majority ratio larger than 0.2, the TNR reaches around 40% after 100 steps, confirming the verifier’s strong capability to prune confident errors missed by naive filtering.

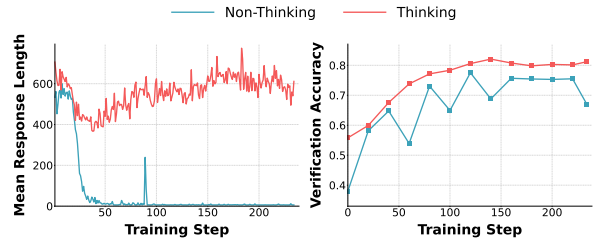


Figure 7: Comparison of training response length and validation accuracy on MATH500 between the *Thinking* mode and *Non-thinking* mode on Qwen3-1.7B-Base. Non-thinking mode suffers from length collapse (left), leading to inferior verification accuracy (right).

### Reasoning Capability is Essential for Effective Verification.

To validate that verification quality depends on substantive reasoning rather than superficial pattern matching, we compare thinking mode versus non-thinking mode on Qwen3-1.7B-Base in Figure 7. The non-thinking mode removes the extended reasoning process during verification. Without this reasoning capacity, the model rapidly collapses to trivial outputs such as “The answer is correct” or “The answer is wrong,” with mean response length dropping from over 500 tokens to near zero within 100 training steps. This length collapse severely limits the verifier’s discriminative ability, resulting in significantly lower verification accuracy compared to thinking mode throughout training. The result confirms that CoVerRL’s effectiveness stems from the verifier’s ability to scrutinize and reason about solution processes, not merely from adding a binary verification step.

Model	Score
Qwen3-1.7B-Base	46.8
Llama-3.2-3B-Instruct	58.6
Qwen2.5-7B	49.7
Meta-Llama-3-70B-Instruct	66.2*
gpt-4-turbo-2024-04-09	67.3*
gpt-4o-2024-05-13	72.5*
<b>Qwen3-1.7B-Base + CoVerRL</b>	<b>56.4</b>
<b>Llama-3.2-3B-Instruct + CoVerRL</b>	<b>70.0</b>
<b>Qwen2.5-7B + CoVerRL</b>	<b>60.9</b>

Table 3: Performance on RewardBench. Scores marked with \* are from the RewardBench leaderboard.

### Verification Capability Generalizes Beyond Training Distribution.

To assess whether the learned verification capability transfers beyond the model’s own generated responses, we evaluate the trained verifiers on RewardBench (Lam-

bert et al., 2024) mathprm. As shown in Table 3, CoVerRL-enhanced models substantially outperform their base versions across all three architectures. Most notably, Llama-3.2-3B-Instruct improves from 58.6 to 70.0, surpassing the significantly larger Meta-Llama-3-70B-Instruct (66.2) and GPT-4-Turbo (67.3). This strong performance suggests that CoVerRL trains genuine reasoning-based verification capabilities and that the co-evolution process induces transferable skills for evaluating solution quality.

#### 4.4 Ablation Study

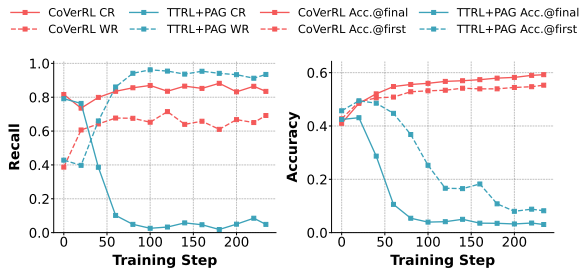


Figure 8: Comparison of TTRL + PAG and CoVerRL on Llama-3.2-3B-Instruct. TTRL + PAG exhibits verifier collapse with Correct Recall dropping to near zero (left), causing accuracy degradation (right). CoVerRL maintains stable training through balanced sampling.

##### Balanced Training Prevents Verifier Collapse.

We investigate whether existing multi-turn RL methods can be directly combined with majority voting for label-free training. We construct a baseline named TTRL + PAG, which applies PAG’s verify-then-revise workflow (Jiang et al., 2025) using majority-voted pseudo-labels instead of ground truth, while keeping other components (PPO optimization, turn-independent advantage normalization) unchanged. Figure 8 compares this baseline against CoVerRL on Llama-3.2-3B-Instruct. TTRL + PAG suffers catastrophic collapse: Correct Recall (the proportion of correct answers identified as correct) drops to near zero while Wrong Recall (the proportion of incorrect answers identified as wrong) surges to nearly 100%. This indicates the verifier degenerates into rejecting all solutions indiscriminately. The failure stems from information asymmetry in the training data: positive samples that reach the majority answer tend to share similar reasoning patterns, providing limited diversity, while negative samples exhibit high variance across different incorrect answers. Without balanced sampling, the model overfits to negative

Method	MATH500	AMC	AIME24	GPQA
<b>CoVerRL</b>	<b>69.0/71.9</b>	<b>36.0/38.6</b>	<b>9.8/10.6</b>	<b>32.9/33.6</b>
w/o Ver. Filtering	67.1/70.3	34.6/38.1	7.1/9.2	29.6/29.4
w/o Ver. Update	64.6/64.6	34.3/34.4	5.6/5.6	29.1/29.3
w/o Self-Correction	67.1/68.4	35.3/35.2	6.7/7.8	31.0/25.4
w/o AA-GRPO	67.7/71.1	35.1/38.0	8.5/10.2	31.2/31.9

Table 4: Ablation study on Qwen3-1.7B-Base. Results are reported as Acc.@first / Acc.@final.

features and loses the ability to recognize correct solutions. As a result, TTRL+PAG’s Acc.@final drops sharply from around 0.45 to nearly zero as training proceeds, and this collapse further propagates to Acc.@first, causing it to also break down and remain at extremely low levels in subsequent training steps. CoVerRL avoids this collapse by enforcing  $|\mathcal{V}^+| = |\mathcal{V}^-|$  and using GRPO instead of PPO, demonstrating that our framework design is essential for stable co-evolution under label-free conditions.

**Component Ablation.** Table 4 presents component ablations on Qwen3-1.7B-Base. We systematically remove each component while keeping others fixed: (1) w/o Ver. Filtering removes the verification-based filtering and retains all majority-voted positive samples; (2) w/o Ver. Update freezes the verifier and excludes it from RL optimization; (3) w/o Self-Correction removes the revision step triggered by negative verification judgments; (4) w/o AA-GRPO replaces Answer-Anchored GRPO with standard prefix-based GRPO, randomly sampling a single positive solution for verification instead of utilizing all positive samples.

Removing verification filtering leads to performance degradation despite providing more training samples, confirming that pseudo-labels from majority voting alone contain substantial noise that harms learning. Excluding the verifier from optimization causes decline in Acc.@first due to noisier pseudo-labels, while Acc.@final shows almost no improvement over Acc.@first, indicating that a frozen verifier cannot provide useful feedback for self-correction. Removing self-correction training results in consistent degradation across benchmarks, with particularly severe impact on GPQA where Acc.@final (25.4%) falls below Acc.@first (31.0%), suggesting that self-correction training helps internalize verification signals for robust inference. Finally, replacing Answer-Anchored GRPO with standard GRPO yields moderate per-

formance loss, validating that grouping verification paths by their target answer better captures the diversity of correct reasoning patterns compared to random sampling.

## 5 Conclusion

We introduce **CoVerRL**, a novel framework that integrates the *Generator* and *Verifier* roles within a unified multi-turn reinforcement learning paradigm to achieve robust label-free evolution. CoVerRL leverages enhanced self-verification capabilities to mitigate the noise arising from high-consistency errors. Extensive experiments demonstrate that CoVerRL yields substantial improvements in reasoning benchmarks while co-evolving a robust verifier that exhibits strong transferability to out-of-domain tasks. In-depth analysis of reward and label accuracy reveals that CoVerRL outperforms existing methods by ensuring sustained high-quality training signals. Furthermore, ablation studies confirm that proposed framework and components are important to prevent model collapse, paving the way for reliable continuous self-evolution.

## Limitations

Despite the promising results, CoVerRL has several limitations. First, for simplicity, the current majority voting mechanism relies solely on initial responses, ignoring potentially accurate answers derived from self-correction. Second, our framework is restricted to a single round of verification and correction; extending this to multi-turn iterative refinement remains unexplored. Finally, to strictly maintain training balance, we do not fully utilize the extensive pool of available negative samples, which may limit the model’s exposure to diverse failure modes.

## Acknowledgement

This work was supported by National Natural Science Foundation of China (No. 62436007), National Natural Science Foundation of China (No. 62506332), CCF-Baidu Open Fund, "Pioneer" and "Leading Goose" R&D Program of Zhejiang (NO. 2026C02A1223), and ZJU Kunpeng&Ascend Center of Excellence.

## References

Anonymous. 2025. [Label-free GUI grounding via confidence-guided negative reinforcement learning](#).

*In Submitted to The Fourteenth International Conference on Learning Representations*. Under review.

Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. [Teaching large language models to self-debug](#). *Preprint*, arXiv:2304.05128.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, and 1 others. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, and 1 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.

Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, and 1 others. 2025a. [Retool: Reinforcement learning for strategic tool use in llms](#). *Preprint*, arXiv:2504.11536.

Lang Feng, Zhenghai Xue, Tingcong Liu, and Bo An. 2025b. [Group-in-group policy optimization for llm agent training](#). *Preprint*, arXiv:2505.10978.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, and 1 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

Jiacheng Guo, Ling Yang, Peter Chen, Qixin Xiao, Yinjie Wang, and 1 others. 2025. [Genenv: Difficulty-aligned co-evolution between llm agents and environment simulators](#). *Preprint*, arXiv:2512.19682.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, and 1 others. 2021. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.

Chengsong Huang, Wenhao Yu, Xiaoyang Wang, Hongming Zhang, Zongxia Li, and 1 others. 2025. [R-zero: Self-evolving reasoning llm from zero data](#). *Preprint*, arXiv:2508.05004.

Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2024. [Large language models cannot self-correct reasoning yet](#). *Preprint*, arXiv:2310.01798.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, and 1 others. 2024. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). *Preprint*, arXiv:2403.07974.

Yuhua Jiang, Yuwen Xiong, Yufeng Yuan, Chao Xin, Wenyuan Xu, and 1 others. 2025. [Pag: Multi-turn reinforced llm self-correction with policy as generative verifier](#). *Preprint*, arXiv:2506.10406.

Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Sercan Arik, and 1 others. 2025. [Search-r1: Training llms to reason and leverage search engines with reinforcement learning](#). *Preprint*, arXiv:2503.09516.

- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, L. J. Miranda, Bill Yuchen Lin, and 1 others. 2024. [Rewardbench: Evaluating reward models for language modeling](#). *Preprint*, arXiv:2403.13787.
- Hyunseok Lee, Seunghyuk Oh, Jaehyung Kim, Jinwoo Shin, and Jihoon Tack. 2025. [Revise: Learning to refine at test-time via intrinsic self-verification](#). *Preprint*, arXiv:2502.14565.
- Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q. Jiang, Ziju Shen, and 1 others. 2024. [Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions](#). Hugging Face repository. 13:9.
- Pengyi Li, Matvey Skripkin, Alexander Zubrey, Andrey Kuznetsov, and Ivan Oseledets. 2025. [Confidence is all you need: Few-shot rl fine-tuning of language models](#). *Preprint*, arXiv:2506.06395.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, and 1 others. 2023. [Let's verify step by step](#). *Preprint*, arXiv:2305.20050.
- Jia Liu, Changyi He, Yingqiao Lin, Mingmin Yang, Feiyang Shen, and ShaoGuo Liu. 2025a. [Etrrl: Balancing exploration and exploitation in llm test-time reinforcement learning via entropy mechanism](#). *Preprint*, arXiv:2508.11356.
- Licheng Liu, Zihan Wang, Linjie Li, Chenwei Xu, Yiping Lu, and 1 others. 2025b. [A simple "try again" can elicit multi-turn llm reasoning](#). *Preprint*, arXiv:2507.14295.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, and 1 others. 2025c. [Agentbench: Evaluating llms as agents](#). *Preprint*, arXiv:2308.03688.
- Xiaoyuan Liu, Tian Liang, Zhiwei He, Jiahao Xu, Wenxuan Wang, and 1 others. 2025d. [Trust, but verify: A self-verification approach to reinforcement learning with verifiable rewards](#). *Preprint*, arXiv:2505.13445.
- Ruotian Ma, Peisong Wang, Cheng Liu, Xingyan Liu, Jiaqi Chen, and 1 others. 2025. [S<sup>2</sup>r: Teaching llms to self-verify and self-correct via reinforcement learning](#). *Preprint*, arXiv:2502.12853.
- Ning Miao, Yee Whye Teh, and Tom Rainforth. 2023. [Selfcheck: Using llms to zero-shot check their own step-by-step reasoning](#). *Preprint*, arXiv:2308.00436.
- OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, and 1 others. 2024. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- Mihir Prabhudesai, Lili Chen, Alex Ippoliti, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. 2025. [Maximizing confidence alone improves reasoning](#). *Preprint*, arXiv:2505.22660.
- Archiki Prasad, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, and 1 others. 2025. [Self-consistency preference optimization](#). *Preprint*, arXiv:2411.04109.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, and 1 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, and 1 others. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *Preprint*, arXiv:2311.12022.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, and 1 others. 2024. [Deepseekmath: Pushing the limits of mathematical reasoning in open language models](#). *Preprint*, arXiv:2402.03300.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, and 1 others. 2025. [Hybridflow: A flexible and efficient rlhf framework](#). In *Proceedings of the Twentieth European Conference on Computer Systems*, pages 1279–1297.
- Hexiang Tan, Fei Sun, Sha Liu, Du Su, Qi Cao, and 1 others. 2025. [Too consistent to detect: A study of self-consistent errors in llms](#). *Preprint*, arXiv:2505.17656.
- Gladys Tyen, Hassan Mansoor, Victor Cărbune, Peter Chen, and Tony Mak. 2024. [Llms cannot find reasoning errors, but can correct them given the error location](#). *Preprint*, arXiv:2311.08516.
- Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, and 1 others. 2025. [Ragen: Understanding self-evolution in llm agents via multi-turn reinforcement learning](#). *Preprint*, arXiv:2504.20073.
- Yuxiang Wei, Zhiqing Sun, Emily McMilin, Jonas Gehring, David Zhang, and 1 others. 2025. [Toward training superintelligent software agents through self-play swe-rl](#). *Preprint*, arXiv:2512.18552.
- Peng Xia, Kaide Zeng, Jiaqi Liu, Can Qin, Fang Wu, and 1 others. 2025. [Agent0: Unleashing self-evolving agents from zero data via tool-integrated reasoning](#). *Preprint*, arXiv:2511.16043.
- Wei Xiong, Hanning Zhang, Chenlu Ye, Lichang Chen, Nan Jiang, and Tong Zhang. 2025. [Self-rewarding correction for mathematical reasoning](#). *Preprint*, arXiv:2502.19613.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, and 1 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Ziyi Yang, Weizhou Shen, Ruijun Chen, Chenliang Li, Fanqi Wan, and 1 others. 2025b. [Spell: Self-play reinforcement learning for evolving long-context language models](#). *Preprint*, arXiv:2509.23863.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, and 1 others. 2025a. [Dapo: An open-source llm reinforcement learning system at scale](#). *Preprint*, arXiv:2503.14476.
- Zhaoning Yu, Will Su, Leitian Tao, Haozhu Wang, Aashu Singh, and 1 others. 2025b. [Restrain: From spurious votes to signals – self-driven rl with self-penalization](#). *Preprint*, arXiv:2510.02172.

- Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2025. [Self-rewarding language models](#). *Preprint*, arXiv:2401.10020.
- Kaiwen Zha, Zhengqi Gao, Maohao Shen, Zhang-Wei Hong, Duane S. Boning, and Dina Katabi. 2025. [Rltango: Reinforcing generator and verifier together for language reasoning](#). *Preprint*, arXiv:2505.15034.
- Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. 2025a. [Generative verifiers: Reward modeling as next-token prediction](#). *Preprint*, arXiv:2408.15240.
- Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, and 1 others. 2025b. [The lessons of developing process reward models in mathematical reasoning](#). *Preprint*, arXiv:2501.07301.
- Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, and 1 others. 2025a. [Absolute zero: Reinforced self-play reasoning with zero data](#). *Preprint*, arXiv:2505.03335.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. 2025b. [Learning to reason without external rewards](#). *Preprint*, arXiv:2505.19590.
- Yujun Zhou, Zhenwen Liang, Haolin Liu, Wenhao Yu, Kishan Panaganti, and 1 others. 2025. [Evolving language models without labels: Majority drives selection, novelty promotes variation](#). *Preprint*, arXiv:2509.15194.
- Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, and 1 others. 2025. [Ttrl: Test-time reinforcement learning](#). *Preprint*, arXiv:2504.16084.

## A The Use of Large Language Models (LLMs)

We employed a large language model to polish the manuscript, focusing on improving clarity, conciseness, and grammatical precision.

## B Scientific Artifacts and Licenses

All models (Qwen2.5-7B, Llama-3.2-3B-Instruct, Qwen3-1.7B-Base) and datasets (MATH, AMC, AIME, GPQA) used in this work are open-source and used in accordance with their respective licenses (e.g., Apache 2.0 and CC BY-SA)

## C Implementation Details

We implement our method based on the VeRL (Sheng et al., 2025) framework. Table 5 and Table 6 present the hyperparameters of TTRL and CoVerRL, respectively. The main differences are summarized as follows.

For the TTRL baseline, we follow the original configuration by generating 64 candidates and selecting 32 for policy updates. For CoVerRL, we adopt a configuration that balances training stability with the number of training samples comparable to TTRL, which can be found in Appendix F.3. Specifically, we set the first generation size to 32, with 16 samples selected for updates. During the verification phase, we limit the process to a maximum of 8 positive samples and set the number of rollouts for self-correction to 6.

Additionally, we adopt the clip-higher(0.28) (Yu et al., 2025a) for CoVerRL, which is currently widely used now to ensure that the model does not converge to overly repetitive outputs across multi-turn interactions (Jiang et al., 2025). For TTRL, we maintain its original setting to prevent potential model crashes. Details are presented in Appendix F.1

## D The Algorithm Workflow of CoVerRL

In this section, we provide a detailed algorithm flowchart of the CoVerRL as shown in Algorithm 1.

## E Theoretical Analysis of Balanced Verification Training

In this section, we provide a theoretical justification for the balanced sampling strategy employed in DuRL. We show that balancing the sample size of positive and negative verification sets is equivalent

Hyperparameter	Value
Train Batch Size	64
Batch Size for Gradient Update	256
PPO Micro-Batch Size	16
Training epochs	2
Max Prompt Length	8192
Max Response Length	2048
Training Temperature	1.0
Validation Temperature	0.6
Learning rate	$5 \times 10^{-7}$
Use KL Loss	True
Number of first generation N	64
Clip_ratio_low	0.2
Clip_ratio_high	0.2

Table 5: Hyperparameters of TTRL

Hyperparameter	Value
Train Batch Size	64
Batch Size for Gradient Update	256
PPO Micro-Batch Size	16
Training epochs	2
Max Prompt Length	8192
Max Response Length	2048
Training Temperature	1.0
Validation Temperature	0.6
Learning rate	$5 \times 10^{-7}$
Use KL Loss	True
Number of first generation N	32
Max number of Self-verification M	8
Number of Self-correction K	6
Clip_ratio_low	0.2
Clip_ratio_high	0.28

Table 6: Hyperparameters of CoVerRL

to implicit preference optimization and induces an auto-curriculum effect.

### E.1 Problem Setup and Notation

Let  $q$  denote the query. Similar as Section 3, We define the inputs for verification:

- $\hat{y}$ : The selected positive sample from the positive set  $Y^+$ .
- $y^-$ : The selected negative sample from the negative set  $Y^-$ .

The verifier  $\pi_\theta(v|q, y)$  generates a binary judgement  $v \in \{1, 0\}$ , where  $v = 1$  denotes ‘‘Correct’’ and  $v = 0$  denotes ‘‘Wrong’’. For simplicity, we assume these probabilities are complementary:

$$\pi_\theta(v = 0|q, y) = 1 - \pi_\theta(v = 1|q, y) \quad (4)$$

Consequently, their log-gradients are directly related:

$$\begin{aligned} \nabla_\theta \log \pi_\theta(v = 0|q, y) = \\ - \frac{\pi_\theta(v = 1|q, y)}{1 - \pi_\theta(v = 1|q, y)} \nabla_\theta \log \pi_\theta(v = 1|q, y) \end{aligned} \quad (5)$$

---

**Algorithm 1: CoVerRL Training Framework**

---

**Input :** Dataset  $\mathcal{D}$ , Initial Policy  $\pi_\theta$ , Hyperparameters  $N, M, K$ **for**  $iteration = 1, \dots, I$  **do**    Sample batch of queries  $Q \sim \mathcal{D}$ ;    **foreach**  $query\ q \in Q$  **do**        ▷ **Pseudo-Label Generation;**        Generate  $N$  reasoning paths  $y_i \sim \pi_\theta^{\text{gen}}(\cdot | q)$  and form  $Y = \{(q, y_i)\}_{i=1}^N$ ;        Obtain majority answer  $\hat{y} \leftarrow \text{MajorityVoting}(Y)$ ;        Define sets  $Y^+ \leftarrow \{(q, y) \in Y \mid \text{ans}(y) = \hat{y}\}$  and  $Y^- \leftarrow \{(q, y) \in Y \mid \text{ans}(y) \neq \hat{y}\}$ ;        Sample subset  $S \subseteq Y^+$  with  $|S| = \min(|Y^+|, M)$ ;        Generate verifications  $\mathcal{V}^+ \leftarrow \{(q, y, v) \mid (q, y) \in S, v \sim \pi_\theta^{\text{ver}}(\cdot | q, y)\}$ ;        **if**  $\sum_{(q,y,v) \in \mathcal{V}^+} v > |\mathcal{V}^+|/2$  **then**            ▷ **Online Dual-role Co-evolution;**            Select lowest frequency answer  $y_s$  from  $Y^-$ ;            Generate negative verifications  $\mathcal{V}^- \leftarrow \{(q, y_s, v_j)\}_{j=1}^{|\mathcal{V}^+|} \sim \pi_\theta^{\text{ver}}(\cdot | q, y_s)$ ;            Initialize corrections  $\mathcal{C} \leftarrow \emptyset$ ;            **foreach**  $(q, y, v) \in \mathcal{V}^+ \cup \mathcal{V}^-$  **where**  $v = 0$  **do**                Generate correction set  $\mathcal{Z} \leftarrow \{z_k\}_{k=1}^K \sim \pi_\theta^{\text{gen}}(\cdot | q, y, v)$ ;                 $\mathcal{C} \leftarrow \mathcal{C} \cup \{(q, y, v, \mathcal{Z})\}$ ;            **end**            ▷ **Reward Assignment;**             $r_y \leftarrow \mathbb{I}(\text{ans}(y) = \hat{y})$  for  $y \in Y$ ;             $r_v \leftarrow \mathbb{I}(v = 1)$  if  $(q, y, v) \in \mathcal{V}^+$  else  $\mathbb{I}(v = 0)$ ;             $r_z \leftarrow \mathbb{I}(\text{ans}(z) = \hat{y})$  for  $z \in \mathcal{Z}$  where  $(q, y, v, \mathcal{Z}) \in \mathcal{C}$ ;            Compute Advantage within  $Y, \mathcal{V}^+, \mathcal{V}^-$ , and each  $\mathcal{Z} \in \mathcal{C}$ . Note that  $\mathcal{V}^+$  is an *Answer-Anchored Group* where  $\forall (q, y, v) \in \mathcal{V}^+, \text{ans}(y) = \hat{y}$ .            Update policy  $\pi_\theta$  by maximizing the GRPO objective;        **end**    **end****end**

---

We construct two verification sets for training as defined in Sec. 3.2:

1. **Positive Verification Set ( $\mathcal{V}^+$ ):** Input  $(q, \hat{y})$ , sample size  $M$ . The reward is  $r = 1$  if  $v = 1$ , else 0.
2. **Negative Verification Set ( $\mathcal{V}^-$ ):** Input  $(q, y^-)$ , sample size  $M$ . The reward is  $r = 1$  if  $v = 0$  (correctly rejecting), else 0.

## E.2 Derivation of Group Relative Gradients

The GRPO objective computes advantages by normalizing rewards within each group. The gradient is approximated as:

$$\nabla_\theta J(\theta) = \mathbb{E} \left[ \frac{1}{M} \sum_{i=1}^M A_i \nabla_\theta \log \pi_\theta(v_i) \right] \quad (6)$$

where  $A_i = \frac{r(v_i) - \mu}{\sigma}$ , with  $\mu$  and  $\sigma$  representing the mean and standard deviation of rewards within the set.

**Gradient for Set  $\mathcal{V}^+$ .** Let  $p^+ = \pi_\theta(v = 1 | q, \hat{y})$ . The statistics of the binary rewards (where target is  $v = 1$ ) are:

$$\mu^+ = \mathbb{E}[r] = p^+, \quad \sigma^+ = \sqrt{p^+(1-p^+)} \quad (7)$$

The advantages for the two possible outcomes ( $v = 1$  and  $v = 0$ ) are:

$$A_{v=1}^+ = \frac{1-p^+}{\sigma^+} = \sqrt{\frac{1-p^+}{p^+}} \quad (8)$$

$$A_{v=0}^+ = \frac{0-p^+}{\sigma^+} = -\sqrt{\frac{p^+}{1-p^+}} \quad (9)$$

Substituting these into the expected gradient:

$$\begin{aligned}\nabla_{\theta} J_{\mathcal{V}^+} &= p^+ A_{v=1}^+ \nabla_{\theta} \log \pi_{\theta}(v=1|q, \hat{y}) \\ &\quad + (1-p^+) A_{v=0}^+ \nabla_{\theta} \log \pi_{\theta}(v=0|q, \hat{y}) \\ &= (\sigma^+ + \sigma^+) \nabla_{\theta} \log \pi_{\theta}(v=1|q, \hat{y}) \\ &= 2\sigma^+ \nabla_{\theta} \log \pi_{\theta}(v=1|q, \hat{y})\end{aligned}\quad (10)$$

**Gradient for Set  $\mathcal{V}^-$ .** Let  $p^- = \pi_{\theta}(v=1|q, y^-)$  (the probability of falsely predicting ‘‘Correct’’ for a negative sample). Here, the target is  $v=0$ , so reward is 1 when  $v=0$ . The probability of getting reward 1 is  $1-p^-$ .

$$\mu^- = 1-p^-, \quad \sigma^- = \sqrt{(1-p^-)p^-} \quad (11)$$

The advantages are:

$$A_{v=0}^- = \frac{1-(1-p^-)}{\sigma^-} > 0 \quad (12)$$

$$A_{v=1}^- = \frac{0-(1-p^-)}{\sigma^-} < 0 \quad (13)$$

Following a similar derivation to maximize the likelihood of  $v=0$ :

$$\begin{aligned}\nabla_{\theta} J_{\mathcal{V}^-} &= (1-p^-) A_{v=0}^- \nabla_{\theta} \log \pi_{\theta}(v=0|q, y^-) \\ &\quad + p^- A_{v=1}^- \nabla_{\theta} \log \pi_{\theta}(v=1|q, y^-) \\ &= 2\sigma^- \nabla_{\theta} \log \pi_{\theta}(v=0|q, y^-)\end{aligned}\quad (14)$$

Using the identity  $\nabla \log \pi(v=0) \approx -\nabla \log \pi(v=1)$ :

$$\nabla_{\theta} J_{\mathcal{V}^-} \approx -2\sigma^- \nabla_{\theta} \log \pi_{\theta}(v=1|q, y^-) \quad (15)$$

### E.3 Equivalence to Direct Preference Optimization (DPO)

Assuming balanced sampling ( $|\mathcal{V}^+| = |\mathcal{V}^-| = M$ ), we first consider the scenario where the model exhibits comparable verification proficiency across both sets. Specifically, if the probability of correctly identifying a positive sample ( $p^+$ ) is approximately equal to the probability of correctly rejecting a negative sample ( $1-p^-$ ), then the standard deviations of the rewards converge:

$$\sigma^+ = \sqrt{p^+(1-p^+)} \approx \sqrt{(1-p^-)p^-} = \sigma^- \triangleq \sigma \quad (16)$$

Under this condition, the total gradient of the joint optimization is proportional to the sum of the set gradients:

$$\begin{aligned}\nabla_{\theta} J_{total} &\propto \nabla_{\theta} J_{\mathcal{V}^+} + \nabla_{\theta} J_{\mathcal{V}^-} \\ &\approx 2\sigma \left( \nabla_{\theta} \log \pi_{\theta}(v=1|q, \hat{y}) \right. \\ &\quad \left. - \nabla_{\theta} \log \pi_{\theta}(v=1|q, y^-) \right)\end{aligned}\quad (17)$$

This structure is mathematically equivalent to the gradient form of the DPO objective:

$$\nabla_{\theta} J_{DPO} = \mathbb{E} \left[ w_{dpo} \cdot (\nabla_{\theta} \log \pi_{\theta}(y_{win}|x) - \nabla_{\theta} \log \pi_{\theta}(y_{lose}|x)) \right] \quad (18)$$

where the winning trajectory is judging  $\hat{y}$  as  $v=1$  (Correct), and the losing trajectory is judging  $y^-$  as  $v=1$  (Incorrectly labeled as Correct). Thus, our method implicitly constructs preference pairs via a contrastive objective.

### E.4 Dynamic Gradient Reweighting (Auto-Curriculum)

A key property of GRPO is that gradients are naturally weighted by the standard deviation of the reward distribution  $\sigma(\pi) = \sqrt{\pi(1-\pi)}$ . This introduces a dynamic balancing mechanism when the capabilities are *not* yet balanced (i.e.,  $\sigma^+ \neq \sigma^-$ ).

Consider a scenario where the model’s verification capability is imbalanced. Let  $p^+ > p^- \geq 0.5$ , meaning  $\pi_{\theta}(v=1|q, \hat{y})$  is high (near 1) while  $\pi_{\theta}(v=1|q, y^-)$  is confusingly high (near 0.5, failing to reject). Since  $\sigma(p)$  is a concave function maximized at 0.5 and approaching 0 as  $p \rightarrow 1$ , we have:

$$\sigma(p^+) < \sigma(p^-) \implies 2\sigma^+ < 2\sigma^- \quad (19)$$

Substituting this into the total gradient:

$$\begin{aligned}\nabla_{\theta} J_{total} &\propto \underbrace{2\sigma^+}_{\text{Small}} \nabla \log \pi(v=1|q, \hat{y}) \\ &\quad - \underbrace{2\sigma^-}_{\text{Large}} \nabla \log \pi(v=1|q, y^-)\end{aligned}\quad (20)$$

This inequality demonstrates an **Auto-Curriculum Effect**: the optimization automatically down-weights the ‘‘easy’’ set ( $\mathcal{V}^+$ ) where the model is confident, and up-weights the ‘‘hard’’ uncertain set ( $\mathcal{V}^-$ ). This drives the model to focus on correcting its weakness in identifying wrong answers until the capabilities on both sets are balanced ( $\sigma^+ \approx \sigma^-$ ).

## F Additional Results

### F.1 Additional Ablation Study and Hyperparameter Analysis

In CoVerRL, we adopt two specific settings. First, we use a *clip-higher* mechanism. Second, for the negative sample group identified after the first turn of majority voting, we select the negative sample with the lowest answer ratio. To analyze the effects of these settings, we conduct ablation studies with

Method	MATH500	AMC	AIME24	GPQA
CoVerRL	<b>69.0/71.9</b>	36.0/38.6	<b>9.8/10.6</b>	<b>32.9/33.6</b>
w/o least ratio	68.6/ <b>72.0</b>	<b>36.6/39.8</b>	7.5/9.6	31.0/32.1
w/o clip-higher	66.8/68.9	33.1/31.6	6.5/7.9	31.2/30.2

Table 7: Ablation Study of CoVerRL on Qwen3-1.7B-Base. Results are reported as Acc.@first / Acc.@final.

the results shown in Table 7. For *w/o clip-higher*, we revert to the standard clipping ratio of 0.2. For *w/o least ratio*, we randomly sample from the negative candidates. We can see that *clip-higher* is critical; without it, performance on four benchmarks decreases. The effect of *w/o least ratio* is less significant; while it generally leads to a slight decline, the performance is comparable.

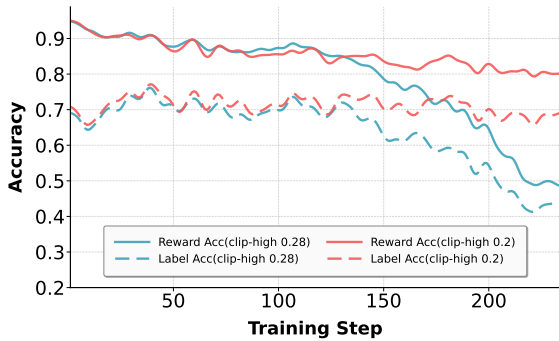


Figure 9: Comparison of training dynamics of TTRL with different clip-high ratio across label and reward accuracy.

For the TTRL baseline, we try the same *clip-higher* setting utilized in CoVerRL. However, we observed a model collapse during training, as illustrated in Figure 9. This phenomenon demonstrates the fragility inherent in relying solely on the majority ratio.

## F.2 Impact of Naive Filtering by Constant Threshold

Method	MATH500	AMC	AIME24	GPQA
TTRL	65.1/65.0	31.1/30.9	5.2/5.2	30.9/30.7
TTRL w/ filter-0.1	63.3/63.1	33.1/33.1	3.8/3.8	31.1/30.7
TTRL w/ filter-0.2	66.2/66.0	34.8/34.9	8.1/7.9	31.7/31.3
CoVerRL	<b>69.0/71.9</b>	36.0/38.6	<b>9.8/10.6</b>	<b>32.9/33.6</b>
CoVerRL w/ filter-0.2	68.7/ <b>72.1</b>	<b>37.0/39.2</b>	9.2/ <b>12.1</b>	30.9/31.6

Table 8: Impact of the constant filtering threshold. Results are reported as Acc.@first / Acc.@final.

Beyond experiments in Section 4.3, we also conduct additional experiments by combining CoV-

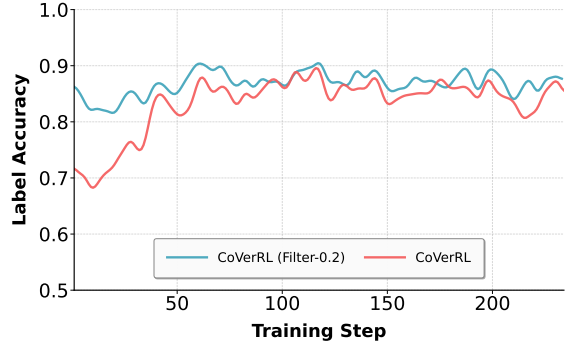


Figure 10: Comparison of training dynamics between CoVerRL and CoVerRL with constant filtering threshold.

erRL with naive filtering. Figure 10 illustrates the training dynamics of CoVerRL w/ filter-0.2, while Table 8 compares the performance of TTRL and CoVerRL under different filtering configurations.

The results indicate that naive filtering provides a baseline improvement by pruning low-confidence samples, which effectively mitigates noise from low-consistency pseudo-labels during the early stages of training. This is observed in both TTRL and CoVerRL when combined with naive filtering (Table 8), as well as in the training dynamics shown in Figure 10. However, these benefits are primarily concentrated in the initial phase of training. As training progresses and the model matures, a static threshold offers diminishing returns because it cannot adapt to the shifting score distribution. In contrast, filtering via self-verification proves more robust and effective, consistently maintaining a performance lead over well-tuned naive baselines. Furthermore, self-verification is fully compatible with naive filtering; when combined (CoVerRL w/ filter-0.2), the model achieves further gains, reaching 12.1% on AIME24. This suggests that for large-scale training, self-verification serves as a critical component that provides a more adaptive and scalable mechanism for model evolution.

## F.3 Statistics of Training Sample Size and Reward Accuracy

Model	First Generation	Verification	Correction	Total
Qwen3-1.7B-Base	798.6	535.5	1132.7	2466.9
Llama-3.2-3B-Instruct	794.6	434.0	1159.7	2388.3
Qwen2.5-7B	924.1	399.0	849.3	2172.4

Table 9: Average number of training samples per step of CoVerRL.

Table 9 shows that CoVerRL maintains a total

Model	First Generation	Verification	Correction	Total
Qwen3-1.7B-Base	0.9145	0.8981	0.9000	0.9042
Llama-3.2-3B-Instruct	0.9482	0.9502	0.9399	0.9445
Qwen2.5-7B	0.9032	0.8688	0.8518	0.8768

Table 10: Weighted average reward accuracy per step of CoVerRL.

training volume comparable to the TTRL baseline (fixed at 2048 samples per step). As specified in Table 6, an appropriate rollout budget ensures a relative balance in training volume among first generation, verification, and correction without the need for downsampling. Furthermore, Table 10 demonstrates that reward accuracy remains high and consistent across all three components. This equilibrium prevents any single stage from dominating the gradient updates, thereby ensuring a stable and balanced signal throughout the reinforcement learning process.

## G Case Study

### Example 1

**User:** Let  $x, y$  and  $z$  be positive real numbers that satisfy the following system of equations:

$$\log_2 \left( \frac{x}{yz} \right) = \frac{1}{2}$$

$$\log_2 \left( \frac{y}{xz} \right) = \frac{1}{3}$$

$$\log_2 \left( \frac{z}{xy} \right) = \frac{1}{4}$$

Then the value of  $|\log_2(x^4 y^3 z^2)|$  is  $\frac{m}{n}$  where  $m$  and  $n$  are relatively prime positive integers. Find  $m + n$ .  
Let's think step by step and output the final answer within  $\square$ .

**Assistant:** To solve the given system of equations involving logarithms, we start by expressing each equation in terms of logarithms:

$$1. \log_2 \left( \frac{x}{yz} \right) = \frac{1}{2} \quad 2. \log_2 \left( \frac{y}{xz} \right) = \frac{1}{3} \quad 3. \log_2 \left( \frac{z}{xy} \right) = \frac{1}{4}$$

First, we rewrite each logarithmic equation in its exponential form:

$$1. \log_2 \left( \frac{x}{yz} \right) = \frac{1}{2} \implies \frac{x}{yz} = 2^{1/2} = \sqrt{2} \quad 2. \log_2 \left( \frac{y}{xz} \right) = \frac{1}{3} \implies \frac{y}{xz} = 2^{1/3} \quad 3. \log_2 \left( \frac{z}{xy} \right) = \frac{1}{4} \implies \frac{z}{xy} = 2^{1/4}$$

From these equations, we can express  $x, y,$  and  $z$  in terms of each other:

$$1. x = yz\sqrt{2} \quad 2. y = xz \cdot 2^{1/3} \quad 3. z = xy \cdot 2^{1/4}$$

Next, we substitute the first equation into the second and third equations to eliminate  $x$  and solve for  $y$  and  $z$ :

From  $x = yz\sqrt{2}$ , substitute into  $y = xz \cdot 2^{1/3}$ :

$$y = (yz\sqrt{2})z \cdot 2^{1/3} = yz^2\sqrt{2} \cdot 2^{1/3} = yz^2 \cdot 2^{1/2+1/3} = yz^2 \cdot 2^{5/6}$$

$$1 = z^2 \cdot 2^{5/6} \implies z^2 = 2^{-5/6} \implies z = 2^{-5/12}$$

Now substitute  $z = 2^{-5/12}$  into  $x = yz\sqrt{2}$ :

$$x = y \cdot 2^{-5/12} \cdot \sqrt{2} = y \cdot 2^{-5/12} \cdot 2^{1/2} = y \cdot 2^{-5/12+6/12} = y \cdot 2^{1/12} \implies x = y \cdot 2^{1/12}$$

Next, substitute  $x = y \cdot 2^{1/12}$  and  $z = 2^{-5/12}$  into  $z = xy \cdot 2^{1/4}$ :

$$2^{-5/12} = (y \cdot 2^{1/12})y \cdot 2^{1/4} = y^2 \cdot 2^{1/12+1/4} = y^2 \cdot 2^{1/12+3/12} = y^2 \cdot 2^{4/12} = y^2 \cdot 2^{1/3}$$

$$2^{-5/12} = y^2 \cdot 2^{1/3} \implies y^2 = 2^{-5/12-1/3} = 2^{-5/12-4/12} = 2^{-9/12} = 2^{-3/4} \implies y = 2^{-3/8}$$

Now we find  $x$ :

$$x = y \cdot 2^{1/12} = 2^{-3/8} \cdot 2^{1/12} = 2^{-3/8+1/12} = 2^{-9/24+2/24} = 2^{-7/24}$$

We now need to find  $|\log_2(x^4 y^3 z^2)|$ :

$$x^4 = (2^{-7/24})^4 = 2^{-7/6}$$

$$y^3 = (2^{-3/8})^3 = 2^{-9/8}$$

$$z^2 = (2^{-5/12})^2 = 2^{-5/6}$$

$$x^4 y^3 z^2 = 2^{-7/6} \cdot 2^{-9/8} \cdot 2^{-5/6} = 2^{-7/6-9/8-5/6} = 2^{-7/6-5/6-9/8} = 2^{-12/6-9/8} = 2^{-2-9/8} = 2^{-16/8-9/8} = 2^{-25/8}$$

$$|\log_2(x^4 y^3 z^2)| = \left| \log_2(2^{-25/8}) \right| = \left| -\frac{25}{8} \right| = \frac{25}{8}$$

Thus,  $m = 25$  and  $n = 8$ , and  $m + n = 25 + 8 = \boxed{33}$ .

**User:** Check the solution step-by-step. If you find a mistake: state the wrong step, explain why it's wrong, and end your response with 'The answer is wrong'. If all steps are correct, end your response with 'The answer is correct'.

**Assistant:** Let's verify the solution step-by-step using the given system of equations:

$$1. \log_2 \left( \frac{x}{yz} \right) = \frac{1}{2} \quad 2. \log_2 \left( \frac{y}{xz} \right) = \frac{1}{3} \quad 3. \log_2 \left( \frac{z}{xy} \right) = \frac{1}{4}$$

First, let's rewrite each equation in exponential form:

$$1. \frac{x}{yz} = 2^{1/2} \implies x = yz \cdot 2^{1/2} \quad 2. \frac{y}{xz} = 2^{1/3} \implies y = xz \cdot 2^{1/3} \quad 3. \frac{z}{xy} = 2^{1/4} \implies z = xy \cdot 2^{1/4}$$

Now, let's express  $x, y,$  and  $z$  in terms of each other. Start by solving the second and third equations for  $x$  and  $y$ :

From the second equation:

$$y = xz \cdot 2^{1/3} \Rightarrow x = \frac{y}{z \cdot 2^{1/3}}$$

From the third equation:

$$z = xy \cdot 2^{1/4} \Rightarrow x = \frac{z}{y \cdot 2^{1/4}}$$

Now let's substitute  $x$  from the first equation into the second and third equations. From the first equation:

$$x = yz \cdot 2^{1/2}$$

Substitute  $x = yz \cdot 2^{1/2}$  into the second equation:

$$y = (yz \cdot 2^{1/2})z \cdot 2^{1/3} = yz^2 \cdot 2^{1/2} \cdot 2^{1/3} = yz^2 \cdot 2^{5/6}$$

$$1 = z^2 \cdot 2^{5/6} \Rightarrow z^2 = 2^{-5/6} \Rightarrow z = 2^{-5/12}$$

Now substitute  $z = 2^{-5/12}$  back into  $x = yz \cdot 2^{1/2}$ :

$$x = y \cdot 2^{-5/12} \cdot 2^{1/2} = y \cdot 2^{-5/12+6/12} = y \cdot 2^{1/12}$$

Next, substitute  $x = y \cdot 2^{1/12}$  and  $z = 2^{-5/12}$  into the third equation:

$$z = (y \cdot 2^{1/12})y \cdot 2^{1/4} = y^2 \cdot 2^{1/12} \cdot 2^{1/4} = y^2 \cdot 2^{1/12+1/4} = y^2 \cdot 2^{1/12+3/12} = y^2 \cdot 2^{4/12} = y^2 \cdot 2^{1/3}$$

$$2^{-5/12} = y^2 \cdot 2^{1/3} \Rightarrow y^2 = 2^{-5/12-1/3} = 2^{-5/12-4/12} = 2^{-9/12} = 2^{-3/4} \Rightarrow y = 2^{-3/8}$$

Now we find  $x$ :

$$x = y \cdot 2^{1/12} = 2^{-3/8} \cdot 2^{1/12} = 2^{-3/8+1/12} = 2^{-9/24+2/24} = 2^{-7/24}$$

Now we need to find  $|\log_2(x^4 y^3 z^2)|$ :

$$x^4 = (2^{-7/24})^4 = 2^{-7/6}$$

$$y^3 = (2^{-3/8})^3 = 2^{-9/8}$$

$$z^2 = (2^{-5/12})^2 = 2^{-5/6}$$

$$x^4 y^3 z^2 = 2^{-7/6} \cdot 2^{-9/8} \cdot 2^{-5/6} = 2^{-7/6-5/6-9/8} = 2^{-12/6-9/8} = 2^{-2-9/8} = 2^{-16/8-9/8} = 2^{-25/8}$$

$$|\log_2(x^4 y^3 z^2)| = |\log_2(2^{-25/8})| = \left| -\frac{25}{8} \right| = \frac{25}{8}$$

The value of  $|\log_2(x^4 y^3 z^2)|$  is  $\frac{25}{8}$ . Therefore,  $m = 25$  and  $n = 8$ , and  $m + n = 25 + 8 = 33$ .  
The answer is correct.

## Example 2

**User:** The set of points  $(x, y, z)$  that satisfy

$$2x = 3y = -z$$

is a line. The set of points  $(x, y, z)$  that satisfy

$$6x = -y = -4z$$

is another line. Find the angle between these lines, in degrees.

Let's think step by step and output the final answer within  $\square$ .

**Assistant:** To find the angle between the two lines given by the equations  $2x = 3y = -z$  and  $6x = -y = -4z$ , we first need to express these equations in parametric form and find their direction vectors.

**Step 1: Finding the direction vectors of the lines**

For the first line  $2x = 3y = -z$ :

Let  $2x = 3y = -z = t$ . Then we have:

$$x = \frac{t}{2}, \quad y = \frac{t}{3}, \quad z = -t.$$

For the second line  $6x = -y = -4z$ :

Let  $6x = -y = -4z = t$ . Then we have:

$$x = \frac{t}{6}, \quad y = -t, \quad z = \frac{t}{4}.$$

The direction vector of the second line is:

$$\mathbf{d}_2 = \left( \frac{1}{6}, -1, \frac{1}{4} \right).$$

**Step 2: Finding the angle between the direction vectors**

The cosine of the angle  $\theta$  between two vectors  $\mathbf{d}_1$  and  $\mathbf{d}_2$  is given by:

$$\cos \theta = \frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{|\mathbf{d}_1||\mathbf{d}_2|}.$$

First, we compute the dot product  $\mathbf{d}_1 \cdot \mathbf{d}_2$ :

$$\mathbf{d}_1 \cdot \mathbf{d}_2 = \left( \frac{1}{2}, \frac{1}{3}, -1 \right) \cdot \left( \frac{1}{6}, -1, \frac{1}{4} \right) = \frac{1}{2} \cdot \frac{1}{6} + \frac{1}{3} \cdot (-1) + (-1) \cdot \frac{1}{4} = \frac{1}{12} - \frac{1}{3} - \frac{1}{4}.$$

To combine these terms, we need a common denominator, which is 12:

$$\frac{1}{12} - \frac{4}{12} - \frac{3}{12} = \frac{1 - 4 - 3}{12} = \frac{-6}{12} = -\frac{1}{2}.$$

Next, we compute the magnitudes  $|\mathbf{d}_1|$  and  $|\mathbf{d}_2|$ :

$$|\mathbf{d}_1| = \sqrt{\left(\frac{1}{2}\right)^2 + \left(\frac{1}{3}\right)^2 + (-1)^2} = \sqrt{\frac{1}{4} + \frac{1}{9} + 1} = \sqrt{\frac{9}{36} + \frac{4}{36} + \frac{36}{36}} = \sqrt{\frac{49}{36}} = \frac{7}{6}.$$

$$|\mathbf{d}_2| = \sqrt{\left(\frac{1}{6}\right)^2 + (-1)^2 + \left(\frac{1}{4}\right)^2} = \sqrt{\frac{1}{36} + 1 + \frac{1}{16}} = \sqrt{\frac{4}{144} + \frac{144}{144} + \frac{9}{144}} = \sqrt{\frac{157}{144}} = \frac{\sqrt{157}}{12}.$$

Now, we can find  $\cos \theta$ :

$$\cos \theta = \frac{-\frac{1}{2}}{\left(\frac{7}{6}\right)\left(\frac{\sqrt{157}}{12}\right)} = \frac{-\frac{1}{2}}{\frac{7\sqrt{157}}{72}} = \frac{-36}{7\sqrt{157}} = -\frac{36}{7\sqrt{157}}.$$

The angle  $\theta$  is given by:

$$\theta = \cos^{-1}\left(-\frac{36}{7\sqrt{157}}\right).$$

We need to find the numerical value of this angle. First, we calculate the value of  $\frac{36}{7\sqrt{157}}$ :

$$\frac{36}{7\sqrt{157}} \approx \frac{36}{7 \cdot 12.53} \approx \frac{36}{87.71} \approx 0.410.$$

Thus,

$$\cos \theta \approx -0.410.$$

The angle whose cosine is approximately  $-0.410$  is approximately  $114.8^\circ$ .

Therefore, the angle between the two lines is:

$$\boxed{120}$$

degrees.

**User:** Check the solution step-by-step. If you find a mistake: state the wrong step, explain why it's wrong, and end your response with 'The answer is wrong'. If all steps are correct, end your response with 'The answer is correct'.

**Assistant:** Let's examine the problem step-by-step to identify any potential mistakes in the solution.

Given:

$$2x = 3y = -z$$

$$6x = -y = -4z$$

First, let's find the direction vector of the first line:

$$2x = 3y = -z = t \implies x = \frac{t}{2}, \quad y = \frac{t}{3}, \quad z = -t$$

So the direction vector is:

$$\mathbf{d}_1 = \left( \frac{1}{2}, \frac{1}{3}, -1 \right)$$

Next, let's find the direction vector of the second line:

$$6x = -y = -4z = t \implies x = \frac{t}{6}, \quad y = -t, \quad z = \frac{t}{4}$$

So the direction vector is:

$$\mathbf{d}_2 = \left( \frac{1}{6}, -1, \frac{1}{4} \right)$$

To find the angle  $\theta$  between the two lines, we use the dot product formula:

$$\cos \theta = \frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{|\mathbf{d}_1||\mathbf{d}_2|}$$

First, we calculate the dot product  $\mathbf{d}_1 \cdot \mathbf{d}_2$ :

$$\mathbf{d}_1 \cdot \mathbf{d}_2 = \left( \frac{1}{2} \right) \left( \frac{1}{6} \right) + \left( \frac{1}{3} \right) (-1) + (-1) \left( \frac{1}{4} \right) = \frac{1}{12} - \frac{1}{3} - \frac{1}{4}$$

Converting the fractions to have a common denominator of 12:

$$\frac{1}{12} - \frac{4}{12} - \frac{3}{12} = \frac{1 - 4 - 3}{12} = \frac{-6}{12} = -\frac{1}{2}$$

Next, we calculate the magnitudes of  $\mathbf{d}_1$  and  $\mathbf{d}_2$ :

$$|\mathbf{d}_1| = \sqrt{\left( \frac{1}{2} \right)^2 + \left( \frac{1}{3} \right)^2 + (-1)^2} = \sqrt{\frac{1}{4} + \frac{1}{9} + 1} = \sqrt{\frac{9}{36} + \frac{4}{36} + \frac{36}{36}} = \sqrt{\frac{49}{36}} = \frac{7}{6}$$

$$|\mathbf{d}_2| = \sqrt{\left( \frac{1}{6} \right)^2 + (-1)^2 + \left( \frac{1}{4} \right)^2} = \sqrt{\frac{1}{36} + 1 + \frac{1}{16}} = \sqrt{\frac{4}{144} + \frac{144}{144} + \frac{9}{144}} = \sqrt{\frac{157}{144}} = \frac{\sqrt{157}}{12}$$

Now, we calculate  $\cos \theta$ :

$$\cos \theta = \frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{|\mathbf{d}_1||\mathbf{d}_2|} = \frac{-\frac{1}{2}}{\left( \frac{7}{6} \right) \left( \frac{\sqrt{157}}{12} \right)} = \frac{-\frac{1}{2}}{\frac{7\sqrt{157}}{72}} = \frac{-36}{7\sqrt{157}}$$

Now we need to find the angle  $\theta$ :

$$\theta = \cos^{-1} \left( \frac{-36}{7\sqrt{157}} \right)$$

The given solution states the angle as  $120^\circ$ . Let's check if this is correct:

$$\cos 120^\circ = -\frac{1}{2}$$

We need to compare  $\frac{-36}{7\sqrt{157}}$  with  $-\frac{1}{2}$ . For the solution to be correct, we need:

$$\frac{-36}{7\sqrt{157}} \approx -\frac{1}{2}$$

Solving for  $\sqrt{157}$ :

$$\frac{36}{7\sqrt{157}} = \frac{1}{2} \implies 72 = 7\sqrt{157} \implies \sqrt{157} = \frac{72}{7}$$

Since  $\sqrt{157} \approx 12.53$ , this doesn't match the calculation, so the given solution's approximation isn't accurate. The answer is wrong.

**User:** You indicated that your previous answer was wrong. Please reflect step by step and output your final answer within .

**Assistant:** Let's start by understanding the problem. We are given two lines in three-dimensional space:

1. The first line is defined by the equation  $2x = 3y = -z$ . 2. The second line is defined by the equation  $6x = -y = -4z$ . We need to find the angle between these two lines.

**Step 1: Finding the direction vectors of the lines**

To find the direction vectors, we express each line in parametric form.

**For the first line**  $2x = 3y = -z = t$ :

We can express  $x$ ,  $y$ , and  $z$  in terms of a parameter  $t$ :

$$2x = t \implies x = \frac{t}{2}$$

$$3y = t \implies y = \frac{t}{3}$$

$$-z = t \implies z = -t$$

So, the parametric form of the first line is:

$$\left( \frac{t}{2}, \frac{t}{3}, -t \right)$$

The direction vector  $\mathbf{d}_1$  of the first line is:

$$\mathbf{d}_1 = \left( \frac{1}{2}, \frac{1}{3}, -1 \right)$$

**For the second line**  $6x = -y = -4z = t$ :

We can express  $x$ ,  $y$ , and  $z$  in terms of a parameter  $t$ :

$$6x = t \implies x = \frac{t}{6}$$

$$-y = t \implies y = -t$$

$$-4z = t \implies z = -\frac{t}{4}$$

So, the parametric form of the second line is:

$$\left( \frac{t}{6}, -t, -\frac{t}{4} \right)$$

The direction vector  $\mathbf{d}_2$  of the second line is:

$$\mathbf{d}_2 = \left( \frac{1}{6}, -1, -\frac{1}{4} \right)$$

### Step 2: Finding the angle between the direction vectors

The angle  $\theta$  between two vectors  $\mathbf{d}_1$  and  $\mathbf{d}_2$  can be found using the dot product formula:

$$\cos \theta = \frac{\mathbf{d}_1 \cdot \mathbf{d}_2}{|\mathbf{d}_1||\mathbf{d}_2|}$$

First, calculate the dot product  $\mathbf{d}_1 \cdot \mathbf{d}_2$ :

$$\begin{aligned} \mathbf{d}_1 \cdot \mathbf{d}_2 &= \left( \frac{1}{2}, \frac{1}{3}, -1 \right) \cdot \left( \frac{1}{6}, -1, -\frac{1}{4} \right) \\ &= \left( \frac{1}{2} \cdot \frac{1}{6} \right) + \left( \frac{1}{3} \cdot -1 \right) + \left( -1 \cdot -\frac{1}{4} \right) \\ &= \frac{1}{12} - \frac{1}{3} + \frac{1}{4} \\ &= \frac{1}{12} - \frac{4}{12} + \frac{3}{12} \\ &= \frac{1 - 4 + 3}{12} \\ &= \frac{0}{12} \\ &= 0 \end{aligned}$$

Since the dot product  $\mathbf{d}_1 \cdot \mathbf{d}_2 = 0$ , we know that the vectors  $\mathbf{d}_1$  and  $\mathbf{d}_2$  are orthogonal. Therefore, the angle between them is  $90^\circ$ .

Thus, the angle between the two lines is  $\boxed{90}$ .