

What If Consensus Lies? Selective-Complementary Reinforcement Learning at Test Time

Dong Yan^{1,2}, Jian Liang^{1,2*}, Yanbo Wang^{1,2}, Shuo Lu², Ran He^{1,2}, Tieniu Tan^{1,2,3}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²NLPR & MAIS, Institute of Automation, Chinese Academy of Sciences

³Nanjing University

yandong2025@ia.ac.cn, liangjian92@gmail.com

Abstract

Test-Time Reinforcement Learning (TTRL) enables Large Language Models (LLMs) to enhance reasoning capabilities on unlabeled test streams by deriving pseudo-rewards from majority voting consensus. However, existing TTRL methods rely exclusively on positive pseudo-labeling strategies. Such reliance becomes vulnerable under challenging scenarios where answer distributions are highly dispersed, resulting in weak consensus that inadvertently reinforces incorrect trajectories as supervision signals. In this paper, we propose SCRL (Selective-Complementary Reinforcement Learning), a robust test-time reinforcement learning framework that effectively mitigates label noise amplification. SCRL develops Selective Positive Pseudo-Labeling, which enforces strict consensus criteria to filter unreliable majorities. Complementarily, SCRL introduces Entropy-Gated Negative Pseudo-Labeling, the first negative supervision mechanism in TTRL, to reliably prune incorrect trajectories based on generation uncertainty. Extensive experiments on multiple reasoning benchmarks demonstrate that SCRL achieves substantial improvements over baselines, while maintaining robust generalization and training stability under constrained rollout budgets. Our code is available at <https://github.com/Jasper-Yan/SCRL>.

1 Introduction

Reinforcement Learning with Verifiable Rewards (RLVR) (Jaech et al., 2024; Shao et al., 2024; Yang et al., 2025a) has significantly advanced the reasoning capabilities of Large Language Models (LLMs), enabling state-of-the-art performance in verifiable domains such as mathematics and coding (Gao et al., 2024; Setlur et al., 2024; Wang et al., 2024). Guided by ground-truth labels or rule-based verification signals, RLVR allows policy optimization

*Corresponding Author

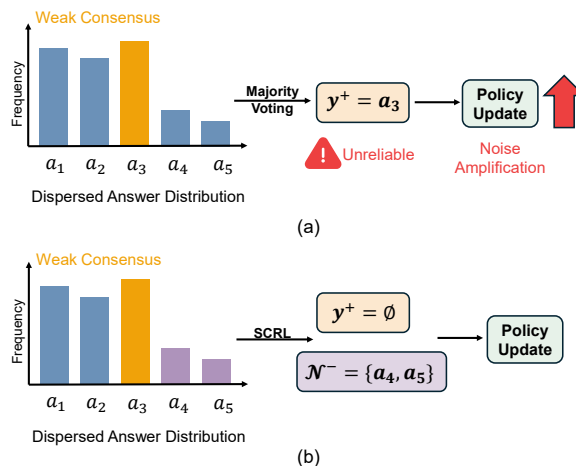


Figure 1: Comparison of pseudo-labeling strategies under weak consensus. (a) Majority voting assigns the positive label despite dispersed answer distribution. (b) SCRL abstains from positive labeling when consensus is insufficient and identifies negative labels.

to directly reinforce trajectories that lead to correct outcomes. However, the reliance on extensive manually-annotated data creates a fundamental limitation: as task complexity and diversity grow, acquiring high-quality supervision becomes increasingly difficult. To bridge this gap, Test-Time Reinforcement Learning (TTRL) has emerged as a critical paradigm for unsupervised reasoning (Zuo et al., 2025; Yang et al., 2025b; Jayalath et al., 2025; Yuan et al., 2025). TTRL allows models to self-improve on unlabeled test streams by generating diverse rollouts and leveraging majority voting consensus to derive pseudo-rewards for policy updates.

While TTRL offers a promising direction for unsupervised reasoning, existing methods (Zuo et al., 2025; Yu et al., 2025b; Wen et al., 2025; Wang et al., 2025a) that rely on majority voting and its variants, such as soft-weighted consensus and self-play mechanisms, face inherent limitations rooted in their exclusive focus on positive pseudo-labeling. These methods require substantial rollout budgets

to achieve reliable consensus; however, on challenging problems, the answer distribution remains highly dispersed even with extensive sampling. As shown in Figure 1(a), this dispersion weakens the consensus, which may result in incorrect trajectories being utilized as supervision signals (Stahlberg et al., 2022; Liu et al., 2025a). Consequently, the model prematurely converges toward spurious solutions through policy optimization (Shi and Jin, 2025; Huang et al., 2024). Note that when rollout budgets are constrained, this vulnerability is particularly pronounced, where insufficient sampling coverage increases consensus instability. In addition, while identifying a correct answer is difficult under high uncertainty, recognizing incorrect answers is comparatively reliable. Nevertheless, existing methods overlook the potential of negative labeling. As illustrated in Figure 1(b), when credible positive consensus is absent, a robust strategy is to employ negative labels to prune the search space, which allows the model to eliminate errors and update toward more promising regions without prematurely committing to any single answer.

To address these critical issues, we propose **SCRL** (Selective-Complementary Reinforcement Learning), a robust framework that effectively mitigates label noise amplification in unsupervised test-time reinforcement learning. SCRL develops *Selective Positive Pseudo-Labeling* which enforces strict consensus and margin criteria, ensuring that positive supervision is only provided when the answer distribution exhibits sharp concentration and clear separation from alternatives. This mechanism can prevent the amplification of unreliable majorities when the answer distribution is dispersed. Complementing this, SCRL introduces *Entropy-Gated Negative Pseudo-Labeling*, the first mechanism in test-time reinforcement learning that integrates negative supervision to identify and penalize incorrect trajectories. By isolating answers that exhibit both low frequency and high uncertainty, the model reliably prunes implausible solutions without eliminating potentially correct low-frequency answers. To calibrate the reinforcement magnitude based on consensus strength, we design *Dynamic Reward Shaping* that integrates credible positive signals with informative negative signals, enabling SCRL to maintain exploration capacity while systematically narrowing the search space and achieve robust unsupervised reinforcement learning.

Extensive experiments on multiple reasoning benchmarks consistently demonstrate that SCRL

significantly outperforms baseline methods, particularly on challenging problems and under constrained rollout budgets. Our contributions can be summarized as follows:

- We propose SCRL, a test-time reinforcement learning framework that mitigates label-noise amplification under weak consensus.
- SCRL incorporates strict consensus criteria to filter unreliable majorities, restricting positive supervision to concentrated distributions.
- SCRL introduces negative supervision in test-time reinforcement learning for the first time, which eliminates implausible trajectories without discarding potentially valid rare solutions.
- Extensive experiments consistently demonstrate that SCRL outperforms baselines particularly under constrained rollout budgets, while ablation studies and label-quality analyses validate the necessity of each component.

2 Related Work

2.1 RL for Reasoning

Reinforcement learning (RL) has emerged as a critical approach for enhancing the instruction-following and complex reasoning capabilities of LLMs. Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022) aligns base models with human preferences via annotated preference data and policy optimization methods (Schulman et al., 2017; Rafailov et al., 2023; Meng et al., 2024; Cui et al., 2025; Li et al., 2026; Lu et al., 2025a). To reduce reliance on human labels, Reinforcement Learning with Verifiable Rewards (RLVR) (Shao et al., 2024; Yu et al., 2025a; Feng et al., 2025; Lu et al., 2025b) replaces preference rewards with verifiable signals, enabling objective and automated evaluation which has proven especially effective in math and code domains (Yang et al., 2025a; Lambert et al., 2024; Wang et al., 2025b). Furthermore, Reinforcement Learning from Internal Feedback (RLIF) (Zhao et al., 2025b; Shafayat et al., 2025; Prasad et al., 2025; Liu et al., 2025b) derives intrinsic rewards from the model’s confidence, entropy or self-consistency across its reasoning paths (Tan et al., 2025; Zhang et al., 2025b,a; Zhao et al., 2025a; Yan et al., 2025). For example, Intuitor (Zhao et al., 2025b) utilizes the model’s confidence as a sparse intrinsic reward to reinforce high confidence reasoning paths, while

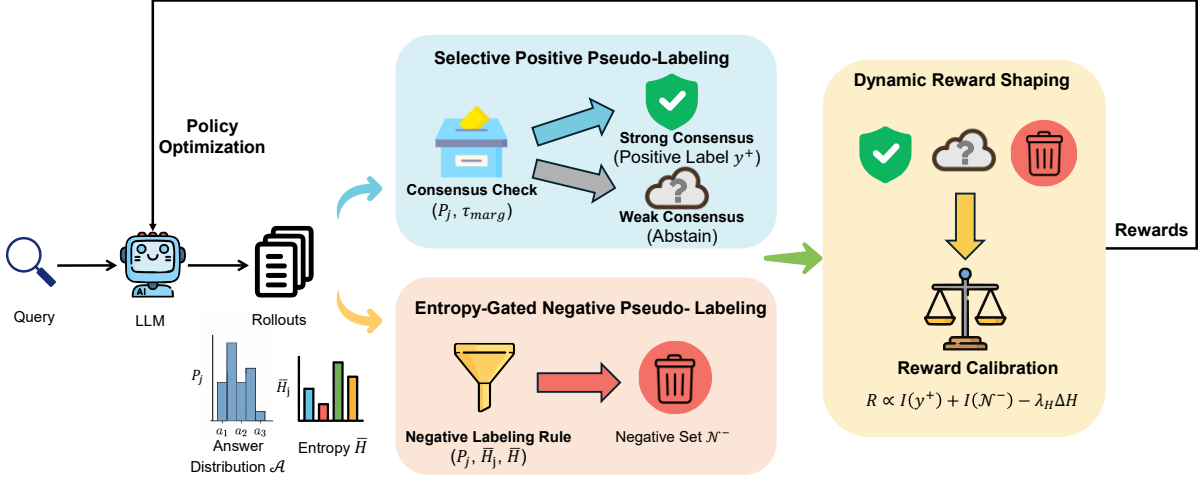


Figure 2: Overview of the SCRL framework. SCRL addresses test-time label noise through three components: selective positive pseudo-labeling enforces strict consensus thresholds to prevent reinforcing unreliable majorities; entropy-gated negative pseudo-labeling identifies negative labels by isolating answers that are both rare and exhibit high uncertainty, pruning the search space without eliminating valid candidates; dynamic reward shaping constructs distribution-aware rewards that scale with consensus strength and penalize uncertainty trajectories.

EMPO (Zhang et al., 2025b) incentivizes reasoning by minimizing the predictive entropy of LLM outputs in a latent semantic space. Our work belongs to the RLIF paradigm and uniquely leverages both positive and negative signals derived from the model’s output distribution to enable robust test-time reinforcement learning.

2.2 Unsupervised Reasoning at Test Time

Test-Time Reinforcement Learning (TTRL) has emerged as a crucial paradigm for adapting LLMs to unlabeled test streams, utilizing majority voting consensus as a verifiable pseudo-reward for online policy optimization (Zuo et al., 2025; Wei et al., 2025; Yu et al., 2025b; Liu et al., 2025a; Wang et al., 2025a; Prabhudesai et al., 2025; Wu et al., 2025; Yu et al., 2026; Zhou et al., 2025). Recent research has focused on robust unsupervised reward estimation: RESTRAIN (Yu et al., 2025b) employs soft-weighted pseudo-labels and penalizes low-confidence responses to enhance training stability, while Self-Harmony (Wang et al., 2025a) utilizes a self-play mechanism to verify positive labels. SPINE (Wu et al., 2025) stabilizes training by restricting updates to high-entropy forking tokens, whereas ETTRL (Liu et al., 2025a) enhances efficiency through entropy-fork tree rollouts to mitigate early-stage estimation bias. MM-UPT (Wei et al., 2025) extends this paradigm to the multi-modal domain, validating the approach for complex vision-language reasoning tasks. However, relying solely on voting-based methods for positive label

assignment can amplify noise when consensus is weak. Our work introduces selective positive labeling with strict consensus criteria and complements it with negative labeling to prune the search space without premature convergence.

3 Method

In this section, we propose **SCRL** (Selective-Complementary Reinforcement Learning), a robust framework for test-time reinforcement learning to mitigate label noise amplification in unsupervised settings. As illustrated in Figure 2, SCRL consists of three components: Selective Positive Pseudo-Labeling (Section 3.2), Entropy-Gated Negative Pseudo-Labeling (Section 3.3), and Dynamic Reward Shaping (Section 3.4).

3.1 Preliminaries

We adopt the Grouped Relative Policy Optimization (GRPO) (Shao et al., 2024) as our main RL algorithm. For a given query q , the policy samples a group of G responses $\mathcal{O} = \{o_1, \dots, o_G\}$ from the sampling policy π_{old} . Each response receives a reward R_i and GRPO constructs a group-normalized advantage \hat{A}_i shared across tokens:

$$\hat{A}_i = \frac{R_i - \mu}{\sigma}. \quad (1)$$

The parameters θ are updated by maximizing the

objective:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q \sim \mathcal{Q}, \mathcal{O} \sim \pi_{old}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \min \left(\rho_{i,t} \hat{A}_i, \text{clip}(\rho_{i,t}, 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) \right], \quad (2)$$

where $\rho_{i,t} = \pi_{\theta}(o_{i,t}|q, o_{i,<t}) / \pi_{old}(o_{i,t}|q, o_{i,<t})$ is the importance sampling ratio.

3.2 Selective Positive Pseudo-Labeling

Majority voting assigns positive pseudo-label by selecting the most frequent answer among rollouts (Zuo et al., 2025; Wang et al., 2025a; Liu et al., 2025a). This implicitly assumes that the most frequent answer is a reliable label for the unknown ground truth. However, on difficult queries or with limited rollout budgets, the answer distribution becomes dispersed: correct trajectories are sparse and erroneous solutions are diverse. In this situation, majority voting can produce false-positive supervision by promoting a wrong answer, which is further amplified under GRPO due to group normalization. Let $f = \frac{1}{G} \sum_{i=1}^G R_i$ be the fraction of rollouts labeled positive within a group. The normalized advantage for positive samples is:

$$\hat{A}^+ = \frac{1 - f}{\sqrt{f(1 - f)}} = \sqrt{\frac{1 - f}{f}}. \quad (3)$$

When consensus is weak, f is small and \hat{A}^+ becomes large, causing a small subset of positive pseudo-labeled trajectories to disproportionately influence policy updates. If the voted answer is incorrect, GRPO can rapidly reinforce this spurious signal, driving the policy toward premature convergence on an erroneous solution.

To mitigate the amplification of label noise under group normalization, we adopt a conservative principle: if we cannot credibly identify a correct answer, we abstain from providing positive supervision. Concretely, we convert majority voting into a selective pseudo-labeling rule with abstention. Given N responses for query q , let $\mathcal{A} = \{a_j\}_{j=1}^K$ be the answer distribution with counts n_j and proportions $p_j = n_j/N$. We denote $j^* = \arg \max_j p_j$ as the index of the most frequent answer and $p_{(2)}$ as the second-largest proportion. We declare a positive pseudo-label y^+ only when the answer distribution is sharply concentrated and well-separated. Formally, $y^+ = a_{j^*}$ if:

$$p_{j^*} \geq \tau_{\text{pos}} \wedge (p_{j^*} - p_{(2)}) > \tau_{\text{marg}}, \quad (4)$$

otherwise, $y^+ = \emptyset$. The threshold τ_{pos} prevents positive labeling when the top-ranked answer has insufficient support, while the margin threshold τ_{marg} enforces separation from the second-ranked answer, preventing unreliable majorities from being reinforced as supervision signals. When $y^+ = \emptyset$, we simply avoid positive reinforcement learning, shifting the learning focus to the negative pseudo-labeling described in Section 3.3.

3.3 Entropy-Gated Negative Pseudo-Labeling

When the answer distribution is dispersed, reinforcing any single trajectory with a positive label is unreliable. Nevertheless, the model’s responses still contain useful signal: while correct answers are difficult to identify with confidence, incorrect answers can be detected more reliably. By constructing high-confidence negative labels, we can prune the search space and encourage the model to update toward more plausible regions without forcing a premature collapse (Zhu et al., 2025).

Entropy-based uncertainty estimation Given responses $\{o_i\}_{i=1}^N$ and the answer distribution $\mathcal{A} = \{a_j\}_{j=1}^K$ with counts n_j and proportions p_j , we distinguish between low-frequency but valid answers and incorrect responses by computing an uncertainty measure from the policy. The Shannon entropy of the next-token distribution over the vocabulary \mathcal{V} at step t is:

$$h_{i,t} = - \sum_{v \in \mathcal{V}} \pi_{\text{old}}(v | o_{i,<t}) \log \pi_{\text{old}}(v | o_{i,<t}). \quad (5)$$

For response o_i with length T_i , we derive the trajectory-level uncertainty \bar{h}_i :

$$\bar{h}_i = \frac{1}{T_i} \sum_{t=1}^{T_i} h_{i,t}. \quad (6)$$

We then aggregate uncertainty at the answer level:

$$\bar{H}_j = \frac{1}{n_j} \sum_{i: a_i = a_j} \bar{h}_i, \quad \bar{H} = \frac{1}{N} \sum_{i=1}^N \bar{h}_i. \quad (7)$$

Intuitively, \bar{H}_j measures the uncertainty with trajectories leading to answer a_j .

Negative pseudo-labeling rule We identify answer a_j as a negative pseudo-label that simultaneously (i) has proportion below the low-support threshold τ_{neg} , and (ii) has generation uncertainty exceeding the query-level average \bar{H} :

$$\mathcal{N}^- = \{a_j \in \mathcal{A} \mid p_j < \tau_{\text{neg}} \wedge \bar{H}_j \geq \bar{H}\}. \quad (8)$$

The condition $\bar{H}_j \geq \bar{H}$ ensures that we only penalize rare answers with high generation uncertainty (Prabhudesai et al., 2025; Liang et al., 2025), thereby preserving potentially correct low-frequency trajectories until stronger consensus emerges in subsequent iterations.

3.4 Dynamic Reward Shaping

In test-time reinforcement learning, assigning fixed rewards to pseudo-labels may amplify noise and destabilize training. To address this, we introduce Dynamic Reward Shaping, which scales reinforcement magnitude based on consensus strength and incorporates uncertainty penalties to guide policy updates. For each response o_i with answer a_i and proportions p_i , the reward is defined as:

$$R_i = p(a_i)\mathbb{I}[a_i = y^+] + (p(a_i) - \tau_{\text{neg}})\mathbb{I}[a_i \in \mathcal{N}^-] - \lambda_H (\bar{H}(a_i) - \bar{H}). \quad (9)$$

The first two terms jointly calibrate the reinforcement magnitude based on the strength of the consensus. The final entropy term, weighted by the coefficient λ_H , gently biases learning toward lower-uncertainty responses. This reward implements a risk-averse strategy that reinforces answers only under credible consensus and eliminates trajectories that are simultaneously rare and uncertain, preventing premature convergence to unreliable solutions.

4 Experiments

4.1 Experimental Setup

Benchmarks and Baselines To evaluate the reasoning capabilities of SCRL, we conduct experiments on six challenging datasets: AIME24 (Li et al., 2024), AIME25 (Li et al., 2024), AMC (Li et al., 2024), MATH-500 (Hendrycks et al., 2021), Minerva (Team et al., 2024), and GPQA (Rein et al., 2024). We compare our method against: (1) **TTRL** (Zuo et al., 2025), which enables the model to self-evolve on unlabeled test data through majority voting. (2) **COMPASS** (Tang et al., 2025), which introduces a composite reward mechanism that jointly optimizes answer reliability and reasoning quality on unlabeled data. (3) **ETMR** (Liu et al., 2025a), which enhances exploration-exploitation balance through entropy-fork tree rollouts and entropy-based advantage reshaping. (4) **RESTRAIN** (Yu et al., 2025b), a self-penalizing framework that penalizes overcon-

fidant and low-consistency rollouts to derive useful learning signals without supervision.

Models To evaluate the generalizability of our method across varying architectures and scales, we utilize a diverse set of open-weight models, including Qwen2.5-3B (Yang et al., 2025a), Qwen2.5-Math-7B (Yang et al., 2025a), Qwen3-4B (Yang et al., 2025a), Llama-3.2-1B-Instruct and Llama-3.1-8B-Instruct (Grattafiori et al., 2024).

Evaluation Metric We report pass@1 (Guo et al., 2025; Chen, 2021) as the primary evaluation metric. For each question, we generate $N = 16$ responses utilizing a temperature of 0.6 and a top- p value of 0.95, with maximum generation length set to 3072 tokens. The pass@1 score is computed as: $\text{pass@1} = \frac{1}{k} \sum_{i=1}^k p_i$, where $p_i \in \{0, 1\}$ indicates the correctness of the i -th response.

Hyperparameter Configuration We optimize the policy model using the AdamW optimizer with a cosine learning rate schedule peaking at 5×10^{-7} . During the rollout phase, we generate 64 (or 32) candidate responses for label estimation, employing a temperature of 1.0 for Qwen2.5-Math and 0.6 for other models to ensure appropriate exploration. Subsequently, we downsample 32 (or 16) responses per prompt for the training update. The maximum token generation length is set to 3072. For the labeling thresholds, we set $\tau_{\text{pos}} = 0.375$, $\tau_{\text{marg}} = 0.125$, and $\tau_{\text{neg}} = 0.125$ across all experiments. The entropy penalty weight is set to $\lambda_H = 0.1$. To accommodate varying dataset sizes and difficulties, we set the number of training episodes to 10 for MATH-500 and Minerva, 30 for AMC, and 80 for AIME, unless otherwise specified. All experiments are conducted on $8 \times$ NVIDIA A100 80GB GPUs.

4.2 Main Results

Tables 1 and 2 present the main results across multiple reasoning benchmarks and models. Our results demonstrate that SCRL consistently achieves substantial improvements over baselines, with particularly pronounced gains on challenging problems and under constrained rollout budgets—precisely the scenarios where weak consensus poses the greatest risk of label noise amplification.

Performance on Vanilla Base Model Under the constrained budget setting, SCRL achieves an average improvement of 1.6% across all benchmarks on Qwen2.5-3B. The most significant improvement is observed on AIME25, where SCRL achieves

Method	AIME25	AMC	MATH-500	Minerva	GPQA	Avg
<i>32 candidate responses, 16 training samples</i>						
Qwen2.5-3B	1.9	23.2	50.0	21.4	21.6	23.6
+ TTRL (Zuo et al., 2025)	2.6	39.4	66.9	31.6	25.0	33.1
+ SCRL (Ours)	8.4	41.5	68.2	29.5	25.7	34.7
Δ	+5.8	+2.1	+1.3	-2.1	+0.7	+1.6
Qwen2.5-Math-7B	4.6	34.0	46.5	10.1	23.0	23.6
+ TTRL (Zuo et al., 2025)	16.8	65.7	85.7	14.5*	25.5	41.6
+ SCRL (Ours)	26.9	66.9	85.6	41.6	25.4	49.3
Δ	+10.1	+1.2	-0.1	+27.1	-0.1	+7.7
<i>64 candidate responses, 32 training samples</i>						
Qwen2.5-3B	1.9	23.2	50.0	21.4	21.6	23.6
+ TTRL (Zuo et al., 2025)	3.9	39.5	72.8	33.5	23.5	34.6
+ SCRL (Ours)	7.7	41.3	73.4	29.5	24.5	35.3
Δ	+3.8	+1.8	+0.6	-4.0	+1.0	+0.7
Qwen2.5-Math-7B	4.6	34.0	46.5	10.1	23.0	23.6
+ TTRL (Zuo et al., 2025)	19.0	68.1	83.4	10.8*	25.5	41.4
+ SCRL (Ours)	22.8	68.5	86.2	43.1	26.0	49.3
Δ	+3.8	+0.4	+2.8	+32.3	+0.5	+7.9

Table 1: Main results on various reasoning benchmarks. We report the pass@1 accuracy (%) across five datasets under two rollout budgets. Results denoted by * represent peak performance before significant degradation during training. Δ shows the performance gain over TTRL (Zuo et al., 2025).

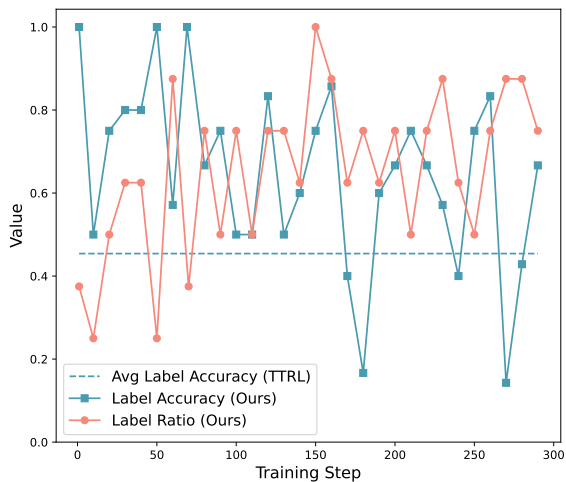
Method	AIME24	MATH / AMC*	Average
<i>Instruct Models</i>			
Llama-3.2-1B [†]	1.5	24.7	13.1
+ TTRL [†]	6.7	27.8	17.3
+ COMPASS [†]	3.5	28.7	16.1
+ SCRL	6.7	39.7	23.2
Δ	+3.2	+11.0	+7.1
Llama-3.1-8B [‡]	4.6	23.3	14.0
+ TTRL [‡]	10.0	32.3	21.2
+ ETMR [‡]	16.9	35.4	26.2
+ RESTRAIN [‡]	16.7	40.0	28.4
+ SCRL	21.9	36.1	29.0
Δ	+5.2	-3.9	+0.6

Table 2: Pass@1 accuracy (%) on Llama-3-Instruct models. Column denoted by * refers to MATH-500 for Llama-3.2-1B-Instruct and AMC for Llama-3.1-8B-Instruct. Results denoted by [†] and [‡] are reported from COMPASS (Tang et al., 2025) and RESTRAIN (Yu et al., 2025b), respectively. Δ shows the performance gain over the corresponding baseline.

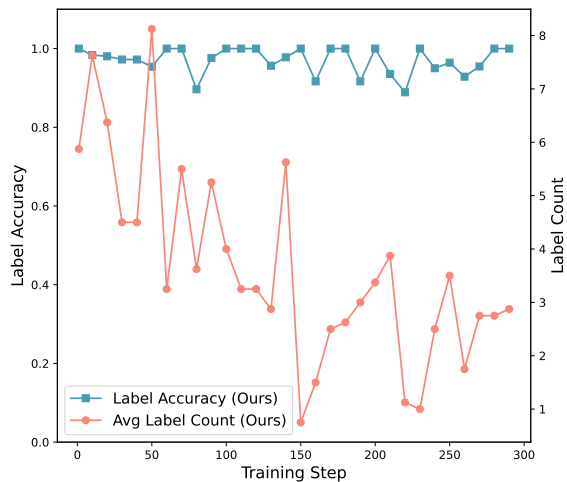
8.4% accuracy compared to TTRL’s 2.6%. This substantial improvement on the most challenging benchmark validates our core hypothesis that SCRL effectively prevents the reinforcement of unreliable majorities when answer distributions are highly dispersed. On AMC and MATH-500, SCRL achieves 41.5% and 68.2% respectively, compared to TTRL’s 39.4% and 66.9%, demonstrating con-

sistent gains. When the rollout budget is increased to 64 candidate responses with 32 training samples, SCRL maintains an average improvement of 0.7%, which aligns with theoretical expectations: with more rollouts, majority voting naturally produces more reliable consensus. Notably, on AIME25, SCRL maintains a substantial improvement of 3.8%, demonstrating that for difficult problems, SCRL can remain effective regardless of computational resources. In addition to mathematics, SCRL consistently improves performance on the GPQA dataset, suggesting that its robustness extends beyond mathematical reasoning to broader general reasoning tasks.

Performance on Math-Specialized Model On Qwen2.5-Math-7B, SCRL achieves an average improvement of 7.7% under the constrained budget. On AIME25, SCRL achieves 26.9% accuracy compared to TTRL’s 16.8%, an absolute improvement of 10.1%. On AMC and MATH-500, SCRL achieves 66.9% and 85.6% respectively, outperforming TTRL on AMC and achieving comparable performance on MATH-500. Notably, on Minerva, while TTRL experiences significant training instability and reaches only 14.5% before performance degradation, SCRL maintains robust training dynamics and achieves 41.6% accuracy. With



(a) Statistics of positive pseudo-label



(b) Statistics of negative pseudo-label

Figure 3: Statistics of positive and negative pseudo-label estimation on the AMC dataset using Qwen2.5-3B.

increased rollout budget of 64 candidate responses and 32 training samples, SCRL sustains its advantage with an average improvement of 7.9%. On AIME25, SCRL attains 22.8% compared to TTRL’s 19.0%. On AMC and MATH-500, SCRL demonstrates continued improvements. Minerva exhibits the most dramatic improvement at 32.3%, while TTRL still shows training instability even with doubled rollout budget. This indicates the necessity for sophisticated label quality assessment and credible reinforcement mechanisms.

Performance on Instruct Model Table 2 further demonstrates that SCRL generalizes effectively to instruction-tuned models with different architectures. We evaluate on Llama-3.2-1B-Instruct and Llama-3.1-8B-Instruct using 64 candidate responses and 32 training samples. On Llama-3.2-1B-Instruct, SCRL achieves an average accuracy of 23.2%, outperforming both TTRL and COM-PASS. On Llama-3.1-8B-Instruct, SCRL achieves the highest average accuracy of 29.0%, surpassing all competing baselines including TTRL, ETMR, and RESTRAIN. The consistency of results across different model families validates that our approach exhibits model-agnostic properties and demonstrates broad applicability across diverse architectures and model scales.

5 Analysis

5.1 Ablation Study

To validate the individual contributions of the proposed components in SCRL, we conduct an ablation study on the AIME25 and AMC datasets using Qwen2.5-3B. Table 3 presents the results. Removing the selective positive pseudo-labeling

Method	AIME25	AMC
SCRL	8.4	41.5
w/o Selective Labeling	3.9 \downarrow 4.5	41.8 \uparrow 0.3
w/o Negative Labeling	1.2 \downarrow 7.2	37.7 \downarrow 3.8
w/o Entropy Gate	3.6 \downarrow 4.8	39.3 \downarrow 2.2
w/o Dynamic Reward	5.1 \downarrow 3.3	39.8 \downarrow 1.7

Table 3: Ablation study of SCRL components on Qwen2.5-3B. The table reports pass@1 accuracy (%).

mechanism reduces the positive labeling method to majority voting and causes a substantial performance drop on the AIME25 dataset. This confirms that weak consensus on challenging problems leads to noise amplification when the most frequent answer is reinforced without credibility checks. On the relatively easier AMC dataset, removing the consensus thresholds achieves a slight performance gain, suggesting that when the model is reasonably accurate, stricter thresholds may overly constrain positive supervision. Furthermore, the substantial degradation when removing negative labeling confirms that negative labeling serves as a complementary signal, pruning the search space and maintaining training stability especially when positive signals are sparse to guide the model.

Removing the entropy gate for negative pseudo-labeling also causes a substantial degradation, indicating that frequency alone is insufficient and uncertainty-aware filtering is necessary to avoid penalizing rare but valid answers. Replacing dynamic reward shaping with hard rewards (+1,0,-1) consistently degrades performance, suggesting that distribution-aware reward magnitudes are important for stabilizing policy updates and calibrating learning to consensus strength.

Method	AMC		MATH-500		Minerva	
	Pass@1	Pass@16	Pass@1	Pass@16	Pass@1	Pass@16
Qwen2.5-3B	23.2	67.5	50.0	86.4	21.4	50.0
+ TTRL (Zuo et al., 2025)	27.3	54.2	58.5	84.8	23.4	48.9
+ SCRL (Ours)	27.5	60.2	65.1	87.0	26.3	49.3

Table 4: Generalization performance of Qwen2.5-3B trained on AIME25. Bold shows the best result.

5.2 Statistics of Label Estimation

To validate the reliability of the supervision signals, we track the dynamics of pseudo-label estimation during training on the AMC dataset using Qwen2.5-3B. As shown in Figure 3a, our method maintains higher positive label accuracy over TTRL, which confirms that the selective positive labeling mechanism effectively filters out unreliable majorities, particularly in early training steps when the initial policy is weak and answer distributions are highly dispersed. Figure 3b shows that entropy-gated negative labeling achieves near-perfect accuracy throughout training, which validates our hypothesis that under high uncertainty, incorrect answers can be identified more reliably than correct ones by examining both frequency and generation uncertainty. Meanwhile, the positive label ratio increases with a decrease in negative labels, indicating that the model progressively develops stronger consensus as its reasoning capability improves, and the search space is being effectively pruned through the elimination of implausible trajectories. The complementary dynamics of positive and negative labels demonstrate that SCRL successfully guides the policy toward higher-quality solution regions while maintaining stable supervision signals.

5.3 Generalization Capabilities

To evaluate whether SCRL develops transferable reasoning capabilities rather than overfitting to training-specific patterns, we train both SCRL and TTRL on AIME25 and evaluate their generalization performance on three out-of-distribution datasets. We use Qwen2.5-3B as the backbone and report both pass@1 and pass@16 results. As shown in Table 4, SCRL consistently improves pass@1 accuracy over TTRL across all three benchmarks. Notably, while TTRL’s pass@16 performance degrades substantially from the base model, SCRL maintains more stable pass@16 scores. This indicates that majority-voting mechanism narrows the solution space during training, whereas SCRL’s

conservative labeling strategy preserves exploration capacity while improving answer quality. These results confirm that selective positive labeling and entropy-gated negative pruning generalize robustly without overfitting to the training dataset.

5.4 Analysis of Labeling Thresholds

To investigate the sensitivity of SCRL to the labeling threshold hyperparameters, we conduct an analysis of τ_{pos} and τ_{neg} using Qwen2.5-Math-7B with 64 candidate responses and 32 training samples. Table 5 presents the results across different threshold configurations. Increasing τ_{pos} from 0.25 to 0.375 consistently improves performance on AIME25. This substantial gain on the most challenging benchmark validates our hypothesis that stricter consensus criteria are crucial for difficult problems where answer distributions are highly dispersed. The more conservative threshold effectively filters out unreliable majorities that would introduce label noise. The impact of τ_{neg} exhibits a complementary pattern. On AMC and MATH-500, increasing τ_{neg} from 0.0625 to 0.125 generally improves performance, indicating that keeping a moderately larger set of negatives can strengthen the contrastive learning signal and enhance exploration of alternative solution strategies. The performance gaps across different configurations become more pronounced on harder problems, reinforcing that SCRL’s conservative labeling strategy is particularly beneficial in scenarios where weak consensus poses the risk of label noise amplification. See Appendix C for analysis of τ_{marg} and λ_H .

τ_{pos}	τ_{neg}	AIME25	AMC	MATH
0.25	0.0625	19.2	64.5	85.5
0.25	0.125	19.0	66.6	85.4
0.375	0.0625	22.1	66.3	85.0
0.375	0.125	22.8	67.5	86.2

Table 5: Parameter analysis of τ_{pos} and τ_{neg} on Qwen2.5-Math-7B. The table reports pass@1 accuracy (%).

Type	AIME24	AMC	MATH-500	Minerva
Correct Trajectories	0.208	0.225	0.245	0.258
Incorrect Trajectories	0.288	0.279	0.272	0.316
Overall Average	0.286	0.266	0.259	0.303

Table 6: Average trajectory-level entropy for correct and incorrect answers across mathematical reasoning benchmarks. Results are computed using Qwen2.5-3B during the initial rollout phase.

5.5 Entropy Analysis

To validate the effectiveness of our entropy-gated negative labeling mechanism, we analyze the relationship between trajectory-level entropy and answer correctness. As shown in Table 6, we compute the average trajectory entropy \bar{h}_i for correct and incorrect trajectories across multiple benchmarks using the Qwen2.5-3B model. The results reveal a consistent and substantial gap between correct and incorrect trajectories across all benchmarks, which validates our hypothesis that generation uncertainty serves as a robust signal for distinguishing implausible solutions. This empirical evidence supports the entropy-gated mechanism, which leverages this uncertainty differential to construct high-confidence negative labels while preserving potentially correct low-frequency trajectories.

5.6 Performance on Large Reasoning Models

To evaluate SCRL’s applicability to long chain-of-thought reasoning models (Wei et al., 2023), we conduct experiments using Qwen3-4B with thinking mode enabled on the AIME25 dataset. As shown in Table 7, SCRL consistently outperforms TTRL across both maximum generation length settings, achieving +2.7% and +0.9% improvements, respectively. These results confirm that selective positive labeling and entropy-gated negative pruning remain effective on long-CoT reasoning models with substantially extended generation lengths.

Method	Max Generation Length	
	10,240	15,360
Qwen3-4B	38.9	51.9
+ TTRL	50.9	59.3
+ SCRL	53.6	60.2
Δ	+2.7	+0.9

Table 7: Pass@1 accuracy (%) on large reasoning models using Qwen3-4B under maximum generation length of 10k and 15k tokens. Δ shows the performance gain over TTRL (Zuo et al., 2025).

5.7 Failure Analysis

SCRL underperforms TTRL on the Minerva dataset with Qwen2.5-3B, contrasting with consistent improvements on other benchmarks. Our analysis suggests this discrepancy arises from the interaction between the smaller model’s limited domain-specific knowledge and the strictness of our selective labeling mechanism.

The Minerva dataset consists of physics-based mathematical problems requiring specialized domain knowledge beyond pure mathematical reasoning. The 3B model, lacking sufficient capacity to store such extensive domain knowledge, produces highly dispersed answer distributions. This dispersion is further exacerbated by the nature of physics-based mathematical problems in Minerva, which involve variations in numerical precision and equivalent formulations, fragmenting the correct solution across similar but non-identical answers. Under these conditions, SCRL’s selective positive labeling becomes over-conservative, frequently abstaining when strict thresholds are unmet, starving the model of positive supervision.

6 Conclusion

In this work, we propose SCRL, a novel framework for test-time reinforcement learning to mitigate label noise amplification in unsupervised settings. SCRL integrates Selective Positive Pseudo-Labeling, which enforces strict consensus criteria to prevent reinforcing unreliable majorities under weak consensus, with Entropy-Gated Negative Pseudo-Labeling, the first negative supervision mechanism in test-time reinforcement learning, which reliably prunes implausible solutions by isolating answers with both low frequency and high uncertainty. This dual mechanism ensures that the model reinforces high-confidence trajectories while reliably identifying and penalizing incorrect answers, without discarding rare but potentially correct trajectories. Empirical results confirm that SCRL significantly outperforms baselines, delivering robust generalization and training stability.

Limitations

While SCRL demonstrates substantial empirical improvements across multiple benchmarks, several limitations remain to be addressed in future work. First, SCRL lacks theoretical analysis for the effectiveness of selective positive pseudo-labeling and entropy-gated negative pseudo-labeling. Although we provide empirical validation through label accuracy statistics and ablation studies, we have not established theoretical guarantees on conditions where our mechanisms provably outperform standard majority voting. Second, the comparison with concurrent methods is constrained by reproducibility challenges. Most related works in the TTRL paradigm have not released code and critical hyperparameter settings cannot be precisely aligned, potentially limiting the scope of our baseline analysis.

Acknowledgements

This research is supported by the National Natural Science Foundation of China under Grants (62276256, U2441251). We thank Yongcan Yu at NLPR for early discussions and feedback to this project. Besides, we also extend our sincere thanks to the anonymous reviewers for their constructive suggestions.

References

- Mark Chen. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Yuchen Zhang, Jiacheng Chen, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, and 1 others. 2025. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*.
- Yunzhen Feng, Parag Jain, Anthony Hartshorn, Yaqi Duan, and Julia Kempe. 2025. Don't waste mistakes: Leveraging negative rl-groups via confidence reweighting. *arXiv preprint arXiv:2510.08696*.
- Jiaxuan Gao, Shusheng Xu, Wenjie Ye, Weilin Liu, Chuyi He, Wei Fu, Zhiyu Mei, Guangju Wang, and Yi Wu. 2024. On designing effective rl reward at training time for llm reasoning. *arXiv preprint arXiv:2410.15115*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- 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 *Proc. NeurIPS*.
- Siyuan Huang, Zhiyuan Ma, Jintao Du, Changhua Meng, Weiqiang Wang, and Zhouhan Lin. 2024. Mirror-consistency: Harnessing inconsistency in majority voting. In *Proc. EMNLP*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Dulhan Jayalath, Shashwat Goel, Thomas Foster, Parag Jain, Suchin Gururangan, Cheng Zhang, Anirudh Goyal, and Alan Schelten. 2025. Compute as teacher: Turning inference compute into reference-free supervision. *arXiv preprint arXiv:2509.14234*.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, and 1 others. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- 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):9.
- Zhiheng Li, Zongyang Ma, Yuntong Pan, Ziqi Zhang, Xiaolei Lv, Bo Li, Jun Gao, Jianing Zhang, Chunfeng Yuan, Bing Li, and Weiming Hu. 2026. Making mllms blind: Adversarial smuggling attacks in mllm content moderation. *Preprint*, arXiv:2604.06950.
- Jian Liang, Ran He, and Tieniu Tan. 2025. A comprehensive survey on test-time adaptation under distribution shifts. *International Journal of Computer Vision*, 133(1):31–64.
- Jia Liu, Changyi He, Yingqiao Lin, Mingmin Yang, Fei Yang Shen, and ShaoGuo Liu. 2025a. Ettl: Balancing exploration and exploitation in llm test-time reinforcement learning via entropy mechanism. *arXiv preprint arXiv:2508.11356*.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025b. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*.

- Shuo Lu, Yanyin Chen, Wei Feng, Jiahao Fan, Fengheng Li, Zheng Zhang, Jingjing Lv, Junjie Shen, Ching Law, and Jian Liang. 2025a. Uni-layout: Integrating human feedback in unified layout generation and evaluation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 7709–7718.
- Shuo Lu, Yinu Xu, Jianjie Cheng, Lingxiao He, Meng Wang, and Jian Liang. 2025b. Deepresearch-slice: Bridging the retrieval-utilization gap via explicit text slicing. *arXiv preprint arXiv:2601.03261*.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. In *Proc. NeurIPS*.
- 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. In *Proc. NeurIPS*.
- Mihir Prabhudesai, Lili Chen, Alex Ippoliti, Katerina Fragkiadaki, Hao Liu, and Deepak Pathak. 2025. Maximizing confidence alone improves reasoning. *arXiv preprint arXiv:2505.22660*.
- Archiki Prasad, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, Mohit Bansal, Sainbayar Sukhbaatar, Jason Weston, and Jane Yu. 2025. Self-consistency preference optimization. In *Proc. ICML*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. In *Proc. NeurIPS*.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First conference on language modeling*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Amrith Setlur, Saurabh Garg, Xinyang Geng, Naman Garg, Virginia Smith, and Aviral Kumar. 2024. RL on incorrect synthetic data scales the efficiency of LLM math reasoning by eight-fold. In *Proc. NeurIPS*.
- Sheikh Shafayat, Fahim Tajwar, Ruslan Salakhutdinov, Jeff Schneider, and Andrea Zanette. 2025. Can large reasoning models self-train? *arXiv preprint arXiv:2505.21444*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: 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*, pages 1279–1297.
- Wenlei Shi and Xing Jin. 2025. Heimdall: test-time scaling on the generative verification. *arXiv preprint arXiv:2504.10337*.
- Felix Stahlberg, Ilya Kulikov, and Shankar Kumar. 2022. Uncertainty determines the adequacy of the mode and the tractability of decoding in sequence-to-sequence models. In *Proc. ACL*.
- Chuyi Tan, Peiwen Yuan, Xinglin Wang, Yiwei Li, Shaoxiong Feng, Yueqi Zhang, Jiayi Shi, Ji Zhang, Boyuan Pan, Yao Hu, and 1 others. 2025. Diagnosing and mitigating system bias in self-rewarding RL. *arXiv preprint arXiv:2510.08977*.
- Chenwei Tang, Jingyu Xing, Lin Long, Xinyu Liu, Deng Xiong, Wei Ju, Shudong Huang, Jiancheng Lv, and Ziyue Qiao. 2025. Rewarding the journey, not just the destination: A composite path and answer self-scoring reward mechanism for test-time reinforcement learning. *arXiv preprint arXiv:2510.17923*.
- Qwen Team and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).
- Junqiao Wang, Zeng Zhang, Yangfan He, Zihao Zhang, Xinyuan Song, Yuyang Song, Tianyu Shi, Yuchen Li, Hengyuan Xu, Kunyu Wu, and 1 others. 2024. Enhancing code LLMs with reinforcement learning in code generation: A survey. *arXiv preprint arXiv:2412.20367*.
- Ru Wang, Wei Huang, Qi Cao, Yusuke Iwasawa, Yutaka Matsuo, and Jiaxian Guo. 2025a. Self-harmony: Learning to harmonize self-supervision and self-play in test-time reinforcement learning. *arXiv preprint arXiv:2511.01191*.
- Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, and 1 others. 2025b. Reinforcement learning for reasoning in large language models with one training example. In *Proc. NeurIPS*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models. *Preprint*, arXiv:2201.11903.
- Lai Wei, Yuting Li, Chen Wang, Yue Wang, Linghe Kong, Weiran Huang, and Lichao Sun. 2025. Unsupervised post-training for multi-modal LLM reasoning via grpo. In *Proc. NeurIPS*.
- Wen Wen, Tianwu Zhi, Kanglong Fan, Yang Li, Xinge Peng, Yabin Zhang, Yiting Liao, Junlin Li, and Li Zhang. 2025. Self-evolving vision-language models for image quality assessment via voting and ranking. *arXiv preprint arXiv:2509.25787*.

- Jianghao Wu, Yasmeeen George, Jin Ye, Yicheng Wu, Daniel F Schmidt, and Jianfei Cai. 2025. Spine: Token-selective test-time reinforcement learning with entropy-band regularization. *arXiv preprint arXiv:2511.17938*.
- Dong Yan, Gaochen Wu, and Bowen Zhou. 2025. Mission impossible: Feedback-guided dynamic interactive planning for improving reasoning on llms. *arXiv preprint arXiv:2510.05577*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Ziyi Yang, Weizhou Shen, Ruijun Chen, Chenliang Li, Fanqi Wan, Ming Yan, Xiaojun Quan, and Fei Huang. 2025b. Spell: Self-play reinforcement learning for evolving long-context language models. *arXiv preprint arXiv:2509.23863*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gao-hong Liu, Lingjun Liu, and 1 others. 2025a. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Yongcan Yu, Lingxiao He, Jian Liang, Kuangpu Guo, Meng Wang, Qianlong Xie, Xingxing Wang, and Ran He. 2026. Understanding and mitigating spurious signal amplification in test-time reinforcement learning for math reasoning. In *Findings of the Association for Computational Linguistics: ACL 2026*.
- Zhaoning Yu, Will Su, Leitian Tao, Haozhu Wang, Aashu Singh, Hanchao Yu, Jianyu Wang, Hongyang Gao, Weizhe Yuan, Jason Weston, and 1 others. 2025b. Restrain: From spurious votes to signals—self-driven rl with self-penalization. *arXiv preprint arXiv:2510.02172*.
- Wenzhen Yuan, Shengji Tang, Weihao Lin, Jiacheng Ruan, Ganqu Cui, Bo Zhang, Tao Chen, Ting Liu, Yuzhuo Fu, Peng Ye, and 1 others. 2025. Wisdom of the crowd: Reinforcement learning from coevolutionary collective feedback. *arXiv preprint arXiv:2508.12338*.
- Kongcheng Zhang, Qi Yao, Shunyu Liu, Yingjie Wang, Baisheng Lai, Jieping Ye, Mingli Song, and Dacheng Tao. 2025a. Consistent paths lead to truth: Self-rewarding reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.08745*.
- Qingyang Zhang, Haitao Wu, Changqing Zhang, Peilin Zhao, and Yatao Bian. 2025b. Right question is already half the answer: Fully unsupervised llm reasoning incentivization. *arXiv preprint arXiv:2504.05812*.
- Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. 2025a. Absolute zero: Reinforced self-play reasoning with zero data. *arXiv preprint arXiv:2505.03335*.
- Xuandong Zhao, Zhewei Kang, Aosong Feng, Sergey Levine, and Dawn Song. 2025b. Learning to reason without external rewards. *arXiv preprint arXiv:2505.19590*.
- Yujun Zhou, Zhenwen Liang, Haolin Liu, Wenhao Yu, Kishan Panaganti, Linfeng Song, Dian Yu, Xiangliang Zhang, Haitao Mi, and Dong Yu. 2025. Evolving language models without labels: Majority drives selection, novelty promotes variation. *arXiv preprint arXiv:2509.15194*.
- Xinyu Zhu, Mengzhou Xia, Zhepei Wei, Wei-Lin Chen, Danqi Chen, and Yu Meng. 2025. The surprising effectiveness of negative reinforcement in llm reasoning. In *Proc. NeurIPS*.
- Yuxin Zuo, Kaiyan Zhang, Li Sheng, Shang Qu, Ganqu Cui, Xuekai Zhu, Haozhan Li, Yuchen Zhang, Xinwei Long, Ermo Hua, and 1 others. 2025. Ttrl: Test-time reinforcement learning. In *Proc. NeurIPS*.

A Implementation Details

A.1 Prompt Design

Consistent with TTRL (Zuo et al., 2025), we adopt the standard chat templates corresponding to each model architecture. For **Qwen2.5-3B**, we employ the following prompt template:

```
system
You are a helpful assistant.
user
{question} Let's think step by step and output the final answer within \boxed{ }.
assistant
```

For **Qwen2.5-Math-7B**, we employ the following prompt template:

```
system
Please reason step by step, and put your final answer within \boxed{ }.
user
{question}
assistant
```

A.2 Benchmarks

Our evaluation suite comprises six challenging reasoning datasets that span different difficulty levels and problem-solving domains:

- **AIME24 and AIME25** (Li et al., 2024): The official problem sets from the 2024 and 2025 American Invitational Mathematics Examination, respectively.

Hyperparameter	Qwen2.5-3B	Qwen2.5-Math-7B
Max Prompt Length	512	512
Max Response Length	3072	3072
Train Batch Size	8	8
PPO Mini Batch Size	1	1
PPO Micro Batch Size per GPU	2	2
N GPUs per Node	4	8
Learning Rate	5×10^{-7}	5×10^{-7}
LR Warmup Steps Ratio	0.03	0.03
Warmup Style	Cosine	Cosine
Temperature	0.6	1.0
Optimizer	AdamW	AdamW

Table 8: Training hyperparameter configuration for Qwen2.5-3B and Qwen2.5-Math-7B based on the verl (Sheng et al., 2025) framework.




Figure 4: Training dynamics of SCRL and TTRL on Qwen2.5-3B across three mathematical benchmarks.




Figure 5: Training dynamics of SCRL and TTRL on Qwen2.5-Math-7B across three mathematical benchmarks.

- **AMC** (Li et al., 2024): A collection of problems from the American Mathematics Competitions, covering core high-school mathematical domains.
- **MATH-500** (Hendrycks et al., 2021): A 500-problem subset sampled from the MATH dataset, covering mathematical reasoning problems across algebra, geometry, number theory, and combinatorics.
- **Minerva** (Team et al., 2024): The mathematics subset from the Minerva quantitative reasoning benchmark.
- **GPQA** (Rein et al., 2024): A challenging dataset of 448 graduate-level multiple-choice questions written by domain experts in biology, physics, and chemistry.

A.3 Training Episode

For the MATH-500 dataset using Qwen2.5-3B with 32 candidate responses and 16 training samples, the TTRL method experiences training collapse due to excessive optimization steps when the training episode is set to 10. Consequently, we set the training episode to 4 for this specific configuration to ensure fair comparison. All other settings remain

consistent with the descriptions in the Section 4.1.

A.4 Training Hyperparameter

Our training hyperparameters are shown in Table 8.

B Additional Results

Figures 4 and 5 show the pass@1 accuracy (%) throughout training.

C Analysis of Consensus Margin and Entropy Penalty

We conduct an analysis of τ_{marg} and λ_H using Qwen2.5-Math-7B with 64 candidate responses and 32 training samples. Table 9 presents the results across different threshold configurations.

The entropy penalty weight λ_H controls the strength of uncertainty penalization in our dynamic reward shaping mechanism. Setting $\lambda_H = 0.1$ achieves the optimal performance across both benchmarks. Compared to $\lambda_H = 0$, introducing a moderate penalty provides a significant performance gain, confirming that discouraging high-uncertainty trajectories helps the policy avoid converging to unstable solutions. However, increasing the coefficient leads to a performance degradation, which suggests that moderate entropy regularization effectively balances the trade-off between encouraging confident reasoning and maintaining sufficient exploration capacity.

The margin threshold τ_{marg} enforces separation between the top-ranked answer and alternatives, preventing unreliable majorities from being reinforced. The results highlight a distinct trade-off based on problem difficulty. On AIME25, increasing τ_{marg} from 0 to 0.125 achieves substantial improvements, with further increase to 0.25 maintaining strong performance at 22.6%. This demonstrates that for challenging problems where answer distributions are highly dispersed, strict margin requirements are crucial to filter out false-positive consensus. Conversely, on the relatively easier AMC dataset, a lower margin of $\tau_{\text{marg}} = 0.0625$ achieves higher performance, which suggests that on problems where the model exhibits sufficient accuracy, overly strict margins may unnecessarily discard valid positive signals. Our configuration in main experiments demonstrates the most robust generalization across varying difficulty levels, effectively mitigating label noise on hard tasks while maintaining supervision frequency on easier ones.

τ_{marg}	λ_H	AIME25	AMC
0.125	0	19.0	64.4
0.125	0.2	19.0	63.7
0.125	0.5	18.5	66.5
0	0.1	18.8	66.5
0.0625	0.1	18.8	68.8
0.25	0.1	22.6	65.6
0.125	0.1	22.8	68.5

Table 9: Parameter analysis of τ_{marg} and λ_H on Qwen2.5-Math-7B. The table reports pass@1 accuracy (%). Bold shows the configuration in main experiments.

D Algorithmic Details of SCRL

Algorithm 1 SCRL’s Reward Construction

Require: Unlabeled queries \mathcal{Q} ; rollout policy π_{old} ; number of rollouts N ; training samples G ; thresholds $\tau_{\text{pos}}, \tau_{\text{marg}}, \tau_{\text{neg}}$; entropy weight λ_H .

Ensure: Per-query training samples $\{(o_i, R_i)\}_{i=1}^G$

- 1: **for** each query $q \in \mathcal{Q}$ **do**
- 2: Sample responses $\{o_i\}_{i=1}^N \sim \pi_{\text{old}}(\cdot | q)$
- 3: $\{a_i\}_{i=1}^N \leftarrow \{\text{ExtractAnswer}(o_i)\}_{i=1}^N$
- 4: Build answer distribution $\mathcal{A} = \{a_j\}_{j=1}^K$ with counts n_j and proportions $p_j = n_j/N$
- 5: $j^* \leftarrow \arg \max_j p_j$; $p_{(2)} \leftarrow \max_{j \neq j^*} p_j$
- 6: $y^+ \leftarrow \emptyset$
- 7: **if** $p_{j^*} \geq \tau_{\text{pos}} \wedge (p_{j^*} - p_{(2)}) > \tau_{\text{marg}}$ **then**
- 8: $y^+ \leftarrow a_{j^*}$
- 9: **end if**
- 10: **for** $i = 1$ to N **do**
- 11: **for** $t = 1$ to $|o_i|$ **do**
- 12: $h_{i,t} \leftarrow - \sum_{v \in \mathcal{V}} \pi_{\text{old}}(v | q, o_{i < t}) \cdot \log \pi_{\text{old}}(v | q, o_{i < t})$
- 13: **end for**
- 14: $\bar{h}_i \leftarrow \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} h_{i,t}$
- 15: **end for**
- 16: $\bar{H} \leftarrow \frac{1}{N} \sum_{i=1}^N \bar{h}_i$
- 17: **for** each $a_j \in \mathcal{A}$ **do**
- 18: $\bar{H}_j \leftarrow \frac{1}{n_j} \sum_{i: a_i = a_j} \bar{h}_i$
- 19: **end for**
- 20: $\mathcal{N}^- \leftarrow \{a_j \in \mathcal{A} \mid p_j < \tau_{\text{neg}} \wedge \bar{H}_j \geq \bar{H}\}$
- 21: Sample training set $\{o_i\}_{i=1}^G$ from $\{o_i\}_{i=1}^N$
- 22: **for** $i = 1$ to G **do**
- 23: $p(a_i) \leftarrow \frac{1}{N} \sum_{k=1}^N \mathbb{I}[a_k = a_i]$
- 24: $\bar{H}(a_i) \leftarrow \bar{H}_j$ where $a_j = a_i$
- 25: $R_i \leftarrow p(a_i) \mathbb{I}[a_i = y^+] + (p(a_i) - \tau_{\text{neg}}) \cdot \mathbb{I}[a_i \in \mathcal{N}^-] - \lambda_H (\bar{H}(a_i) - \bar{H})$
- 26: **end for**
- 27: **end for**
