

Less Noise, More Voice: Reinforcement Learning for Reasoning via Instruction Purification

Yiju Guo^{♫,†}, Tianyi Hu[♫], Zexu Sun[♫], Yankai Lin^{♫,✉}

[♫] Gaoling School of Artificial Intelligence, Renmin University of China

[♫] Department of Computer Science, Aarhus University

[♫] Baidu Inc.

✉{yijuguo, yankailin}@ruc.edu.cn

 **GitHub:** RUCBM/LENS

Abstract

Reinforcement Learning with Verifiable Rewards (RLVR) has advanced LLM reasoning, but remains constrained by inefficient exploration under limited rollout budgets, leading to low sampling success and unstable training in complex tasks. We find that many exploration failures arise not from problem difficulty, but from a small number of prompt tokens that introduce interference. Building on this insight, we propose the **Less Noise Sampling Framework (LENS)**, which first low-success prompts by identifying and removing interference tokens. then transfers successful rollouts from the purification process to supervise policy optimization on the original noisy prompts, enabling the model to learn to ignore interference in the real-world, noisy prompting settings. Experimental results show that LENS significantly outperforms GRPO, delivering higher performance and faster convergence, with a 3.88% average gain and over $1.6\times$ speedup on math reasoning, and a 1.83% gain on scientific and general reasoning. Our work highlights the critical role of pruning interference tokens in improving rollout efficiency, offering a new perspective for RLVR research.

1 Introduction

Reinforcement Learning with Verifiable Rewards (RLVR), such as GRPO (Shao et al., 2024), has significantly advanced the reasoning capabilities of large language models (LLMs). However, RLVR fundamentally relies on sampling correct rollouts to generate informative learning signals (Yu et al., 2025; Zheng et al., 2025a). In complex reasoning tasks, reward sparsity arises from long-horizon decision making with delayed and binary feedback. When combined with a high-dimensional action space, correct rollouts become exceedingly rare

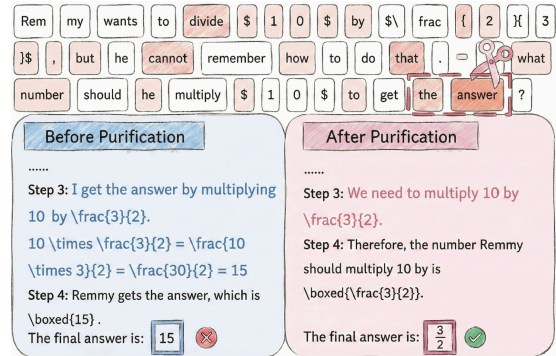


Figure 1: An example of *interference token purification*: Removing a few interference tokens corrects the reasoning rollout and turns it into a successful one.

(Figure 2a), leading to a lack of positive samples and causing training to collapse or become highly inefficient (Hare, 2019).

To mitigate these issues, recent works have primarily followed two directions: (1) scaling exploration by increasing rollout (Xu et al., 2025; Yang et al., 2025b; Zhan et al., 2025; Xiong et al., 2025b), and (2) filtering zero-variance prompts (Yu et al., 2025; Zheng et al., 2025a). However, the former has substantially higher computational cost without improving efficiency, while the latter sacrifices exploration on challenging samples, limiting the model’s ability to solve complex problems. As a result, neither approach truly addresses the core issue of inefficient exploration on challenging samples.

To address this, we investigate why the model fails to explore successful rollouts. Through a fine-grained token-level analysis, surprisingly, we find that *many failures arise not from problem difficulty, but from a few ($< 5\%$) tokens that introduce excessive interference*, as shown in Figure 1 and Figure 2b. We define tokens that are likely to cause failures as *Interference Tokens*. Simply pruning these tokens improves rollout accuracy on previously failed DeepMath (He et al., 2025) samples by over 20% across all the model families (Figure 2c).

✉ Corresponding author: Yankai Lin.

† Work done during an internship at Baidu.




Figure 2: (a) **Zero-Reward Prompt Analysis:** Comparison of the zero-reward prompt ratio across different models and rollout sizes (n). LENS significantly reduces the proportion of zero-reward samples compared to GRPO, enhancing training efficiency. (b) **Distribution of token-level Interference Scores** (log scale): Only a few tokens exhibit high interference. (c) **Rollout Accuracy Improvement:** Removing these interference tokens leads to an improvement in rollout success rates (Average@8).

Building on these insights, we introduce **Less Noise Sampling Framework (LENS)**, an online selective rollout framework that improves high-quality evolution by extracting informative learning signals from low-success prompts. In the first stage, LENS identifies and removes interference tokens within low-success prompts via interference score (Figure 2b), producing prompts that yield a higher proportion of successful rollouts. In the second stage, LENS performs a transfer from the purification process to the original noisy setting: successful rollouts generated under denoised prompts are used as high-reward supervision to calibrate policy optimization on the original prompts. Unlike standard filtering, this mechanism encourages the model to learn to ignore interference tokens, rather than merely fitting solutions under cleaner conditions, ultimately enhancing the robustness of LLM reasoning through self-exploration.

Experimental results show that LENS significantly outperforms GRPO, achieving a *Pareto improvement* in performance–efficiency trade-offs, with an average performance gain of 3.88% and over $1.6\times$ faster convergence across seven math reasoning benchmarks, with further gains on four out-of-domain general reasoning benchmarks. Furthermore, LENS exhibits superior performance over both scaling exploration and prompt filtering baselines while using substantially fewer computational resources. These results further provide empirical support for our hypothesis: low-success, challenging prompts contain valuable training signals, highlighting the critical role of pruning interference tokens in improving rollout efficiency, offering a new perspective for RLVR research.

2 LENS: Less Noise Sampling Framework

We introduce **LENS (LEss Noise Sampling Framework)**, a plug-and-play rollout framework designed to facilitate effective policy exploration. LENS consists of two key components: (1) *Interference Token Identification and Purification*: identifying and purifying interference tokens in low-success prompts through interference score; (2) *Calibrated Rollout Policy Optimization (CRPO)*, an efficient post-training algorithm that transfers learning signals from the denoised prompts produced by Component (1) to the original noisy prompts, equipping the model with the ability to ignore interference and perform robust reasoning under noisy inputs.

2.1 Interference Token Identification and Purification

In this section, we first describe how interference tokens are identified and then explain how prompt purification is performed.

Interference Token Identification. We start from the observation that the reference model π_{ref} provides a stable reference distribution learned from the training data. In contrast, large token-level deviations of the learned policy π_{θ} from this reference often signal over-optimization or spurious behavior driven by noise, which can destabilize exploration (Rafailov et al., 2024).

Motivated by this intuition, we identify interference tokens by measuring the token-level deviation between the current policy and the reference model. Specifically, for a token prefix s and the next generated token a , we define *Interference Score* as:

$$S_I(s, a) \triangleq \left| \log \pi_{\theta}(a | s) - \log \pi_{\text{ref}}(a | s) \right| \quad (1)$$

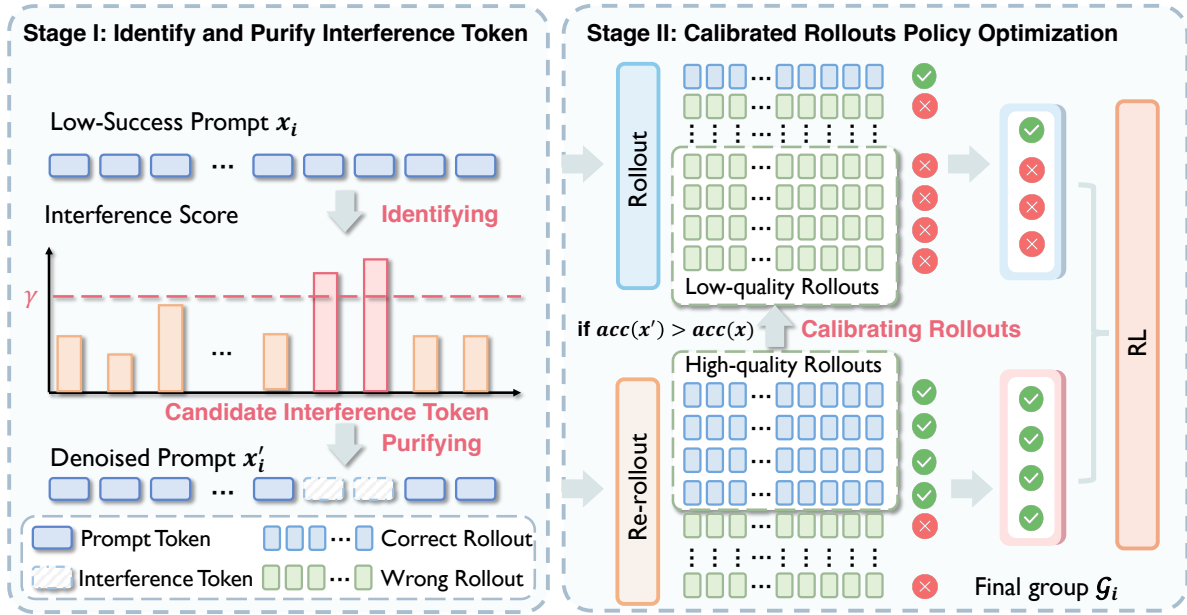


Figure 3: **Method Overview.** In the first stage, LENS identifies and purifies interference tokens within low-success prompts via *Interference Score* (Defined in Section 2.1), thereby generating a higher proportion of successful rollouts. In the second stage, LENS uses successful rollouts from the denoised prompts as high-reward supervision to calibrate policy optimization on the original prompt, correcting gradient updates distorted by interference.

Tokens with large interference scores contribute disproportionately to the KL divergence from the reference distribution and are therefore treated as interference tokens. Such deviations are commonly induced by reward over-optimization or noisy and misleading signals (Rafailov et al., 2024; Gao et al., 2023), and can hinder effective exploration and generalization in the high-dimensional token action space (Engstrom et al., 2020; Dai et al., 2025).

Interference Purification. Based on the observations in Section 1 and the *Interference Score*, we propose a token-wise inspection and pruning mechanism. By removing only a small number of tokens, this mechanism effectively eliminates interference while preserving the original semantics of the prompt with minimal impact.

Specifically, for the i -th prompt x_i , which contains $|x_i|$ tokens after tokenization, we compute token-level *Interference Scores* for all the tokens in the prompt and rank the tokens in descending order. Then, we introduce a deletion ratio γ and select the top $k = \lceil \gamma \cdot |x_i| \rceil$ tokens to form an interference token set I_i . We define the denoised prompt as $x'_i = x_i \setminus I_i$, which denotes removing all tokens in the interference set I_i . We set γ to a small value (e.g., 1%–5%) to preserve the original semantics, with further discussion in Section 4.3.

2.2 Calibrated Rollout Policy Optimization

While interference tokens can be identified, directly removing them is not always beneficial, as only

about 20% of prompts exhibit improvement in roll-out accuracy after removal (Figure 2c).

Therefore, we adopt **Calibrated Rollout Policy Optimization (CRPO)**, which applies interference token purification (Section 2.1) to obtain rollouts with a higher proportion of successful samples. When the original prompt exhibits low sampling success (i.e., success rate below τ ; see Appendix C for sensitivity analysis), CRPO treats rollouts generated from the denoised prompt x'_i as a source of transferable supervision and applies this signal to guide policy optimization on the original prompt x_i , enabling interference-aware calibration. Such calibration equips the model to recognize interference under noisy prompts and thereby prevents training collapse. The complete algorithmic procedure is provided in Algorithm 1.

Sample Reweighting. To mitigate the impact of interference-induced failures and properly incorporate the successful rollouts of the denoised prompt, we adopt a sample reweighting strategy to calibrate the training signal.

Let $Y_i = Y_i^+ \cup Y_i^-$ denote the set of rollouts sampled from the original prompt x_i , where Y_i^+ and Y_i^- are the sets of successful and failed rollouts, respectively. We define the initial sampling success rate as $\bar{a}_i = |Y_i^+|/|Y_i|$.

In addition, we collect a set of successful rollouts P_i by sampling from the denoised prompt x'_i . To ensure that pruning provides a genuine improve-

Algorithm 1 CRPO: Calibrated Rollout Policy Optimization

```

1: Input: Policy  $\pi_\theta$ , reference policy  $\pi_{\text{ref}}$ , dataset  $\mathcal{D}$ , group  $m$ , accuracy threshold  $\tau$ , pruning count  $k$ .
2: for iteration  $t = 1, \dots, N$  do
3:   Sample batch  $\mathcal{B} = \{x_i\}_{i=1}^B \sim \mathcal{D}$ .
4:   for each prompt  $x_i \in \mathcal{B}$  do
5:     Sample  $m$  rollouts  $Y_i = \{y_{i,1}, \dots, y_{i,m}\} \sim \pi_\theta(\cdot | x_i)$ .
6:     Partition  $Y_i$  into success set  $Y_i^+$  and failure set  $Y_i^-$ .
7:     Compute initial success rate  $\bar{a}_i = |Y_i^+|/|Y_i|$ .
8:     Initialize training group  $\mathcal{G}_i \leftarrow Y_i$  and prompt mapping  $x^{\text{roll}}(y) \leftarrow x_i, \forall y \in \mathcal{G}_i$ .
9:     if  $\bar{a}_i < \tau$  then
10:      Interference Identification: ▷ Interference Token Identification and Purification.
11:      Compute  $S_I(t) = |\log \pi_\theta(t | x_i) - \log \pi_{\text{ref}}(t | x_i)|$  for tokens in  $x_i$  (Eq. 1).
12:      Obtain  $x'_i$  by deleting top- $k$  tokens in  $x_i$  with highest  $S_I(t)$ .
13:      Sample  $m$  rollouts  $Y'_i \sim \pi_\theta(\cdot | x'_i)$  and compute success rate  $\text{acc}(x'_i) = |Y'_i^+|/m$ .
14:      if  $\text{acc}(x'_i) > \bar{a}_i$  then
15:        Let  $P_i$  be successful rollouts in  $Y'_i$ .
16:        Select  $R_i \subseteq Y_i^-$  with  $|R_i| = \min(|Y_i^-|, |P_i|)$ .
17:        Update  $\mathcal{G}_i \leftarrow (Y_i \setminus R_i) \cup P_i$ .
18:        Update  $x^{\text{roll}}(y) \leftarrow x'_i$  for all  $y \in P_i$ .
19:      end if
20:    end if
21:    Policy Calibration: ▷ Calibrated Rollout Policy Optimization.
22:    Compute importance ratios  $\rho(y; \theta)$  (Eq. 5) and weights  $\tilde{w}(y)$  based on  $\bar{a}_i$  (Eq. 4).
23:    Compute calibrated advantages  $\hat{A}(y)$  (Eq. 6) and update  $\pi_\theta$  using  $\mathcal{L}(\theta)$  (Eq. 7).
24:  end for
25: end for

```

ment, we activate rollout replacement only when the denoised prompt achieves higher empirical accuracy:

$$g_i = \mathbb{I}[\text{acc}(x'_i) > \text{acc}(x_i)]. \quad (2)$$

When $g_i = 1$, we replace a subset of failed rollouts in Y_i^- with successful ones from P_i . Specifically, we randomly sample a subset $R_i \subseteq Y_i^-$ with $|R_i| = \min(|Y_i^-|, |P_i|)$. Let \mathcal{G}_i be the reconstructed rollout set:

$$\mathcal{G}_i = \begin{cases} Y_i^+ \cup (Y_i^- \setminus R_i) \cup P_i, & \text{if } g_i = 1, \\ Y_i, & \text{otherwise.} \end{cases} \quad (3)$$

We define the unnormalized weight $\tilde{w}(y)$ for each rollout $y \in \mathcal{G}_i$ by applying \bar{a}_i as a scaling factor:

$$\tilde{w}(y) = \begin{cases} \bar{a}_i, & y \in Y_i^+, \\ 1 - \bar{a}_i, & y \in P_i \cup (Y_i^- \setminus R_i). \end{cases} \quad (4)$$

Objective Function. To enable transfer of learning signals from the denoised prompt x'_i back to the original prompt x_i , we formulate a unified objective that corrects for distribution mismatch and stabilizes policy optimization. Let $x^{\text{roll}}(y) \in \{x_i, x'_i\}$

denote the prompt variant used to sample rollout y under the rollout policy π_{old} . We apply importance correction by defining the ratio:

$$\rho(y; \theta) = \frac{\pi_\theta(y | x_i)}{\tilde{w}(y)\pi_{\text{old}}(y | x^{\text{roll}}(y))}. \quad (5)$$

We compute group-relative advantages by normalizing rewards over the reconstructed rollout set \mathcal{G}_i :

$$\hat{A}(y) = \frac{r(y) - \text{mean}_{y' \in \mathcal{G}_i}[r(y')]}{\text{std}_{y' \in \mathcal{G}_i}[r(y')]} \quad (6)$$

Finally, we optimize a PPO-style clipped surrogate objective with KL regularization:

$$\begin{aligned} \mathcal{L}(\theta) = & - \sum_{y \in \mathcal{G}_i} \min \left(\rho(y; \theta) \hat{A}(y), \right. \\ & \left. \text{CLIP}(\rho(y; \theta), 1 - \epsilon, 1 + \epsilon) \hat{A}(y) \right) \\ & + \beta \mathbb{D}_{\text{KL}}(\pi_\theta(\cdot | x_i) \| \pi_{\text{ref}}(\cdot | x_i)). \end{aligned} \quad (7)$$

Overall, CRPO provides a selective and stable calibration mechanism that improves policy optimization by exploiting high-quality rollouts revealed through interference purification.

3 Experiments

3.1 Experiment Settings

Base Models. We conduct experiments across three distinct model families, Llama-3.2 (Meta, 2024), Qwen-2.5 (Team et al., 2024) and Qwen-3 (Yang et al., 2025a), covering various parameter scales. Specifically, we utilize five models: Llama-3.2-3B-Instruct (Meta, 2024), Qwen2.5-3B, Qwen2.5-7B (Team et al., 2024), Qwen3-4B-Base and Qwen3-8B-Base (Yang et al., 2025a).

Baseline. We evaluate LENS against vanilla GRPO (Shao et al., 2024) and two representative strategies for handling zero-variance prompts: increased sampling and zero-variance filtering. For increased sampling, we implement GRPO_{extended}, which doubles the rollout budget per instruction. For zero-variance filtering, we compare with DAPO (Yu et al., 2025) and GRESO (Zheng et al., 2025a), maintain training stability by discarding or skipping zero-variance prompts via post-rollout filtering and pre-rollout prediction, respectively. We also report their extended variants, DAPO_{extended} and GRESO_{extended}, which run twice the number of training epochs. Notably, LENS operates under a strictly lower computational budget, without increasing rollout counts or training epochs.

Training and Evaluation Datasets. For the RL training phase, we employ Openr1-Math-46k (Yan et al., 2025), a large-scale, high-quality dataset designed for mathematical reasoning. In the evaluation, we assess model performance across a diverse set of seven benchmarks: MATH500 (Hendrycks et al., 2021), AMC23 (AI-MO, 2024), AIME24, AIME25 (Li et al., 2024), GaokaoEN-2023 (Zhang et al., 2023), Minerva (Lewkowycz et al., 2022), and OlympiadBench (He et al., 2024). This selection covers a broad range of difficulty levels, enabling comprehensive evaluation. To further assess out-of-domain generalization, we additionally evaluate on scientific reasoning (SuperGPQA (Du et al., 2025), GPQA-Diamond (Rein et al., 2024)) and general complex reasoning benchmarks (BBEH (Kazemi et al., 2025), MMLU-Pro (Wang et al., 2024)).

Training Settings. We leverage GRPO as the basis for LENS, setting the KL coefficient $\beta = 0.001$ and the imitation coefficient $\gamma = 0.001$. Both preliminary and validation experiments are conducted using Llama-3.2, Qwen-2.5, and Qwen-3 models.

We configure the maximum response length to 4096 tokens, the learning rate to 1×10^{-6} , both the rollout and update batch sizes to 128, the number of rollouts to 8, top- p to 1, and the temperature to 1. Detailed training hyperparameters for GRPO are provided in Appendix A.

Evaluation Settings. For evaluation, we set both the temperature and top- p to 1.0, with a maximum generation length of 4,096 tokens. We primarily report the Pass@1 accuracy. To ensure robust evaluation on high-difficulty benchmarks such as AMC23, AIME24, and AIME25, we present the results averaged over 16 generation samples.

3.2 Main Results and Analysis

Table 1 and Table 2 summarize the performance on Qwen3-4B-Base, Qwen3-8B-Base and Llama-3.2-3B-Instruct, while results on Qwen2.5-3B and Qwen2.5-7B are included in Appendix 4. We draw two conclusions: (1) **Superiority over rollout-intensive baselines.** LENS consistently outperforms GRPO (Shao et al., 2024) and GRPO_{extended} given the same training corpus. This indicates that simply increasing the rollout budget is insufficient for generating informative samples when exploration is affected by interference tokens. In contrast, LENS improves sample efficiency by enabling the model to focus on critical information, thereby effectively enhancing its reasoning performance. (2) **Advantage over filtering strategies.** LENS also surpasses the pre-rollout filter GRESO (Zheng et al., 2025a) and the post-rollout filter DAPO (Yu et al., 2025), including in settings with fewer rollouts. These results suggest that aggressively discarding zero-variance prompts can limit capability expansion, particularly on challenging benchmarks, which explains why LENS yields larger improvements on such datasets (e.g. AMC23 & AIME24). (3) **Effectiveness on general reasoning tasks.** LENS consistently outperforms GRPO on both Qwen3 and Llama3.2 families across SuperGPQA, GPQA-Diamond, BBEH, and MMLU-Pro, with an average gain of +1.83%. This demonstrates that interference token purification improves general reasoning capability rather than overfitting to math-specific patterns.

Figure 4 illustrates the learning curves for Qwen3-4B-Base and Qwen3-8B-Base across MATH-500 and OlympiadBench. LENS consistently achieves higher progressive and final accuracy compared to GRPO across all configurations.

Model	Pass@1				Average@16			Avg.
	MATH	Minerva	Olympiad	GAO	AMC23	AIME24	AIME25	
Llama3.2-3B-Instruct	40.02	16.54	15.26	35.58	10.16	1.88	0.00	17.06
+ GRPO	51.60	<u>21.69</u>	20.00	44.68	22.81	6.25	<u>0.83</u>	23.98
+ DAPO	53.00	<u>21.69</u>	19.26	47.01	26.09	9.79	0.42	25.32
+ GRESO	51.80	<u>21.69</u>	19.70	45.45	26.09	6.46	1.04	24.60
+ GRPO _{extended}	51.20	20.96	19.41	44.68	26.76	6.25	1.04	24.33
+ DAPO _{extended}	52.20	20.22	<u>20.15</u>	<u>48.31</u>	<u>27.34</u>	<u>10.21</u>	0.42	<u>25.55</u>
+ GRESO _{extended}	<u>53.20</u>	18.38	18.81	46.75	27.19	7.29	<u>0.83</u>	24.64
+ LENS _{GRPO}	55.80	22.43	21.04	48.83	29.84	10.62	<u>0.83</u>	27.03
Qwen3-4B-Base	72.20	27.21	38.96	65.19	26.72	5.62	6.70	34.66
+ GRPO	79.40	34.56	45.19	<u>70.91</u>	53.28	12.71	<u>14.79</u>	44.41
+ DAPO	80.20	35.29	42.96	68.57	52.97	13.12	11.04	43.50
+ GRESO	<u>80.80</u>	36.76	<u>45.33</u>	69.61	55.78	16.46	12.50	45.32
+ GRPO _{extended}	78.20	<u>39.34</u>	44.30	70.39	54.53	12.92	11.88	44.51
+ DAPO _{extended}	<u>80.80</u>	<u>39.34</u>	<u>45.33</u>	68.57	53.44	16.04	11.04	44.94
+ GRESO _{extended}	80.60	37.50	44.74	69.06	<u>58.75</u>	<u>18.33</u>	12.71	<u>45.96</u>
+ LENS _{GRPO}	83.20	41.54	50.37	72.99	60.94	18.54	16.46	49.15
Qwen3-8B-Base	77.00	34.29	41.19	64.68	38.59	5.62	5.83	38.17
+ GRPO	84.00	<u>45.22</u>	48.15	73.25	58.91	18.33	18.75	49.52
+ DAPO	85.00	44.12	52.00	73.51	62.03	21.46	17.92	50.86
+ GRESO	84.20	41.18	49.93	71.69	64.69	20.83	15.42	49.71
+ GRPO _{extended}	<u>85.20</u>	43.38	52.15	76.01	<u>66.25</u>	22.08	19.17	<u>52.03</u>
+ DAPO _{extended}	85.00	44.49	<u>52.44</u>	74.55	65.47	22.50	18.75	51.88
+ GRESO _{extended}	84.60	43.38	50.81	<u>75.32</u>	60.62	21.46	<u>19.38</u>	50.80
+ LENS _{GRPO}	86.00	48.16	53.33	<u>75.32</u>	66.56	23.96	20.21	53.36

Table 1: **Detailed evaluation results on seven math reasoning benchmarks.** The best and second best results are in **bold** and underlined. (*) Even in the unfavorable setting where GRPO_{extended}, DAPO_{extended} and GRESO_{extended} are trained for 2× more rollouts, LENS still outperforms them on the majority of benchmarks.

Model	SuperGPQA	GPQA-D	BBEH	MMLU-Pro	Avg.
Llama3.2-3B-Instruct	18.49	25.25	5.53	34.69	20.99
+ GRPO	<u>18.77</u>	<u>29.29</u>	<u>6.99</u>	<u>36.84</u>	<u>22.97</u>
+ LENS _{GRPO}	20.05	30.17	8.34	38.21	24.19
Qwen3-4B-Base	25.99	26.30	8.05	51.60	27.99
+ GRPO	<u>30.17</u>	<u>40.40</u>	<u>10.88</u>	<u>59.35</u>	<u>35.20</u>
+ LENS _{GRPO}	31.66	43.94	11.95	60.65	37.05
Qwen3-8B-Base	30.40	33.30	10.55	58.00	33.06
+ GRPO	<u>35.26</u>	<u>48.99</u>	<u>13.08</u>	<u>62.41</u>	<u>39.94</u>
+ LENS _{GRPO}	35.59	55.56	13.36	64.88	42.35

Table 2: **Out-of-domain generalization results.** Models are trained only on mathematical data and evaluated on scientific reasoning (SuperGPQA, GPQA-D) and general complex reasoning (BBEH, MMLU-Pro). GPQA-D denotes GPQA-Diamond. The best and second best results are in **bold** and underlined.

Notably, on the more challenging OlympiadBench, LENS exhibits more stable and continuous improvement, in contrast to the fluctuations observed with

GRPO. We attribute this advantage to two primary factors: (1) LENS enhances the model’s exploration capacity by effectively removing potential inter-




Figure 4: **Learning curves of LENS and GRPO across model scales and task difficulties.** We compare Qwen3-4B/8B-Base backbones on MATH-500 (Medium) and OlympiadBench (High). LENS converges faster and achieves comparable or higher final accuracy than GRPO under the same training step, indicating more efficient optimization.

ference tokens, leading to improved sample quality, and facilitating more effective exploration of capability boundaries; (2) By contrasting correct responses from interference-free rollouts with erroneous responses in the original sampling, LENS helps the model focus on key information, thereby boosting its reasoning abilities.

4 Further Analysis

In this section, we conduct training dynamics analysis (Section 4.1), efficiency analysis (Section 4.2), threshold sensitivity analysis (Section 4.3), pruning strategies analysis (Appendix E) and computational overhead (Appendix D).

4.1 Training Dynamics Analysis

We analyze the training dynamics of LENS in two complementary views: (1) the sampling accuracy distribution across training stages (Figure 5), and (2) the evolution of training reward, policy entropy, and response length (Figure 6). Together, they characterize both the *quality* of rollouts and the *trajectory* of policy learning under LENS.

Higher-quality rollouts throughout training.

We first examine the sampling accuracy distribution across three training stages: early (Steps 1–100), middle (Steps 101–200), and late (Steps 201–300). As shown in Figure 5, LENS substantially reduces the proportion of zero-reward prompts (*Failure* category) compared to GRPO and GRPO_{extended} across all stages. Correspondingly, more prompts are shifted into the *Mid* and *High* categories, which contain more informative learning signals.

More stable and decisive policy learning. Beyond rollout quality, LENS also exhibits healthier learning trajectories. (1) **Stable accuracy progression.** In Figures 6(a, d), LENS achieves

steadier accuracy improvements than GRPO and GRPO_{extended} across model scales. (2) **Emergence of confident long-form reasoning.** LENS yields more rapid gains in both accuracy and response length while maintaining moderate entropy (Cheng et al., 2025) (Figure 6), with the earlier emergence of long-form reasoning (“aha moment” (Guo et al., 2025a)) reflecting more decisive reasoning.

4.2 Efficiency Analysis

We evaluate the learning efficiency of LENS relative to GRPO on MATH-500 and OlympiadBench. We quantify efficiency by the relative speedup in terms of gradient steps required to reach the peak average accuracy attained by the GRPO baseline.

Results. Figure 7 illustrates the comparative learning curves of LENS and GRPO. Notably, LENS achieves the same peak performance as GRPO while requiring $1.67\times$ fewer gradient steps on MATH-500 and attaining a $1.64\times$ acceleration on OlympiadBench. We attribute this substantial gain in training efficiency to the enhanced quality of the rollouts generated during exploration. By effectively mitigating noise from interference tokens, LENS provides more stable and informative rollouts, which significantly accelerates convergence compared to the vanilla GRPO.

4.3 Threshold Sensitivity Analysis

We conduct a sensitivity analysis on the pruning threshold for interference tokens using Qwen2.5-3B and Qwen2.5-7B on MATH-500. Figure 8 shows validation accuracy curves for various thresholds ranging from 1% to 5%, alongside the vanilla GRPO baseline. Detailed sensitivity analysis of the success rate threshold τ in Appendix C.

Results. We observe two key findings: (1) LENS consistently matches or outperforms the GRPO




Figure 5: **Sampling accuracy distribution across three training phases.** We compare the sampling distributions of GRPO, GRPO_{extended} and LENS across the early, middle and late training stages.

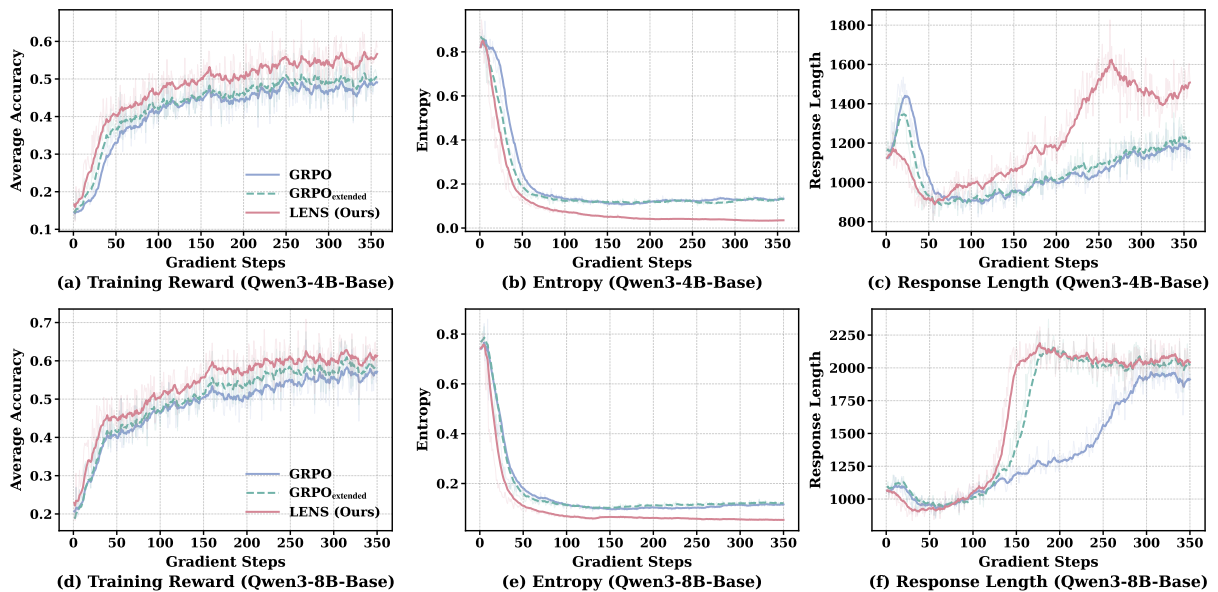


Figure 6: **Training dynamics across different model scales.** Each row reports average accuracy, entropy, and response length during training for Qwen3-4B-Base (top) and Qwen3-8B-Base (bottom). Compared with GRPO and GRPO_{extended}, LENS exhibits more consistent and stable trends.

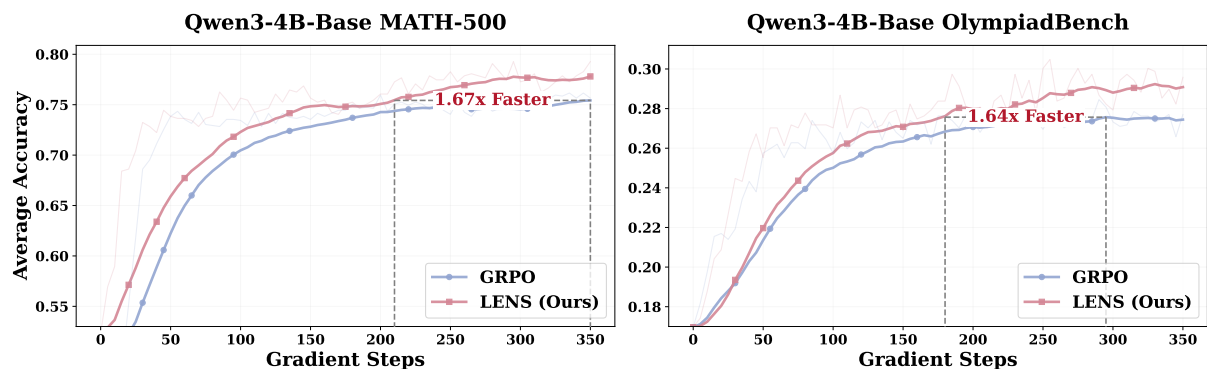


Figure 7: **Training efficiency comparison of LENS and GRPO on MATH-500 and OlympiadBench.** The gray dashed lines indicate the number of training steps required for both methods to reach the highest average accuracy of the baseline during the entire training period. LENS demonstrates superior sample efficiency and faster convergence.

throughout the entire training process across all evaluated thresholds (1%–5%). (2) As highlighted in the insets, the optimal threshold exhibits a clear

correlation with model capacity: the Qwen2.5-3B model yields superior performance with a higher threshold, whereas the more capable Qwen2.5-7B

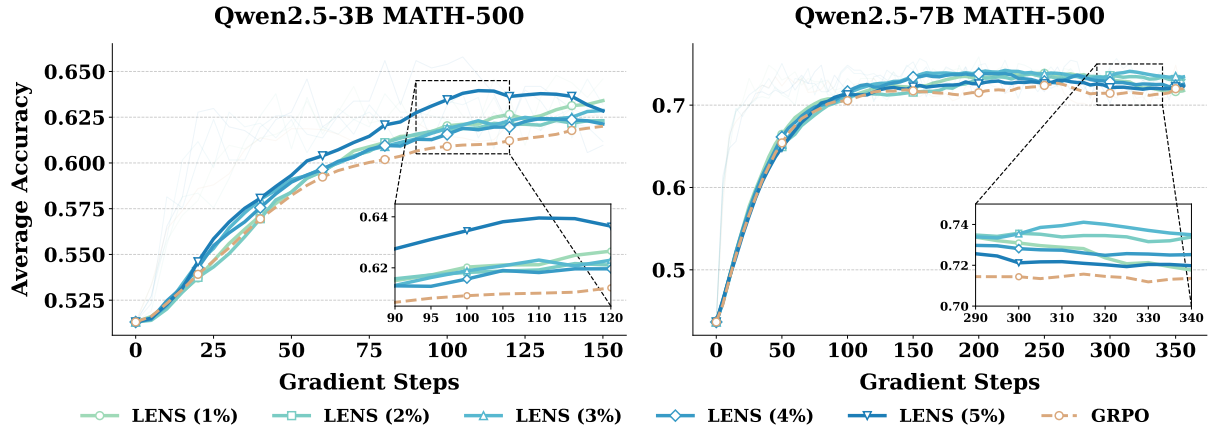


Figure 8: **Performance convergence of LENS on MATH-500 with Qwen2.5-3B/7B.** LENS matches or exceeds GRPO throughout training. Insets highlight that the optimal threshold depends on model capacity: the weaker 3B model requires a higher threshold, while the stronger 7B model achieves optimal results with a lower threshold.

model achieves optimal performance with a lower threshold. This pattern suggests that models with limited capacity are more strongly affected by interference tokens.

5 Related Work

Credit Assignment. Credit assignment (Kazemnejad et al., 2024; Bentegeac et al., 2025; Chai et al., 2024) is crucial for improving Reinforcement Learning with Verifiable Rewards (RLVR) (Guo et al., 2025a; Liu et al., 2025; Yue et al., 2025; Yu et al., 2025; Shao et al., 2024; Wen et al., 2025), helping the model identify key decision paths that contribute most to the final decision. Credit assignment methods use forward signals, such as attention (Chan et al., 2024; Li et al., 2025), log probabilities (Bentegeac et al., 2025), and entropy (Tan et al., 2025), to evaluate the importance of each token based on the model’s forward pass. However, existing methods primarily focus on output tokens during generation, while overlooking how distracting information within instruction tokens can mislead model behavior (Guo et al., 2025b). To address this gap, we propose extending the log-probability signal to the instruction level, analyzing and removing low-information instruction tokens to enhance the model’s focus during inference. This approach effectively improves the model’s sampling success rate and sample diversity.

GRPO Signal Collapse. GRPO suffers from signal collapse that identical rewards in a group lead to vanishing gradients, effectively halting further learning. Current studies can be categorized into 4 directions: (1) **Prompt Filtering**, removing extremely difficult prompts to maintain a productive

training set; (Yu et al., 2025; Xiong et al., 2025a; Zheng et al., 2025b); (2) **Scaling Exploration**, dynamically allocating a greater number of rollouts to harder samples, thereby increasing reward variance within groups (Xu et al., 2025; Yang et al., 2025b; Zhan et al., 2025; Xiong et al., 2025b); and (3) **Reward Function Design**, revising the advantage computation to prevent vanishing gradients (Le et al., 2025). These approaches focus on optimizing rollout strategies or reward functions while overlooking the interference tokens in the instruction that degrade sample quality. To address this gap, we propose a sampling framework to dynamically identify and purify these tokens, mitigating signal collapse and improving training efficiency.

6 Conclusion

In this paper, we first reveal a novel observation: in many cases, RLVR is only a few interference tokens away from discovering correct rollouts. Building on this insight, we propose LENS to identify interference tokens and to transfer successful rollouts generated from denoised prompts to calibrate policy optimization on the original noisy prompts. Extensive experiments demonstrate that LENS consistently outperforms GRPO in both performance and efficiency, with a 3.88% average gain and over $1.6\times$ speedup. LENS also exhibits better performance over both scaling exploration and prompt filtering baselines, and generalizes to scientific and general reasoning tasks with a 1.83% average gain. Overall, our findings offer a fundamentally new perspective on improving sampling efficiency in RLVR and open up promising directions for future research.

Limitations

Our work has the following limitations: (1) **Model Scale Constraints:** Due to limited computational resources, our experiments were conducted on models with up to 8B parameters. Evaluating the performance and scalability of our method on larger-scale models (e.g., 32B or 70B) remains an avenue for future research. (2) **Binary Reward Limitation:** The effectiveness of our approach has been validated primarily in tasks with binary rewards. Its applicability to more complex environments, such as those with multi-dimensional scoring, requires further investigation. (3) **Algorithmic Integration:** While we demonstrated the efficacy of our method within the GRPO framework, we have not yet explored its integration with other GRPO-based variants. Specifically, our method could be combined with algorithms that optimize rollout frequency or reward functions, potentially enhancing exploration capabilities and training stability.

Acknowledgments

This work was supported by The National Natural Science Foundation of China (No. 62376273 and No. U2436209). The research is also funded by the Pioneer Centre for AI and the Novo Nordisk Foundation (NNF24OC0099109). We also thank the Department of Computer Science at Aarhus University for providing Tianyi Hu’s travel support.

We sincerely thank all the anonymous reviewers and (S)ACs for their constructive comments, Wentong Chen, Yupeng Huo and Bokai Ji for their help with the implementation, and Tingchen Fu for valuable feedback on the writing.

References

- AI-MO. 2024. Amc 2023. <https://huggingface.co/datasets/AI-MO/aimo-validation-amc>.
- Raphael Bentegeac, Bastien Le Guellec, Gregory Kuchcinski, Philippe Amouyel, and Aghiles Hamroun. 2025. Token probabilities to mitigate large language models overconfidence in answering medical questions: Quantitative study. *Journal of medical Internet research*, 27:e64348.
- Yekun Chai, Haoran Sun, Huang Fang, Shuohuan Wang, Yu Sun, and Hua Wu. 2024. Ma-rlhf: Reinforcement learning from human feedback with macro actions. *arXiv preprint arXiv:2410.02743*.
- Alex J Chan, Hao Sun, Samuel Holt, and Mihaela Van Der Schaar. 2024. Dense reward for free in reinforcement learning from human feedback. *arXiv preprint arXiv:2402.00782*.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu Wei. 2025. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*.
- Juntao Dai, Taiye Chen, Yaodong Yang, Qian Zheng, and Gang Pan. 2025. Mitigating reward over-optimization in rlhf via behavior-supported regularization. *arXiv preprint arXiv:2503.18130*.
- Xinrun Du, Yifan Yao, Kaijing Ma, Bingli Wang, Tianyu Zheng, King Zhu, Minghao Liu, Yiming Liang, Xiaolong Jin, Zhenlin Wei, and 1 others. 2025. Supergpqa: Scaling llm evaluation across 285 graduate disciplines. *arXiv preprint arXiv:2502.14739*.
- Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. 2020. Implementation matters in deep policy gradients: A case study on ppo and trpo. *arXiv preprint arXiv:2005.12729*.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pages 10835–10866. PMLR.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025a. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Yiju Guo, Wenkai Yang, Zexu Sun, Ning Ding, Zhiyuan Liu, and Yankai Lin. 2025b. Learning to focus: Causal attention distillation via gradient-guided token pruning. *arXiv preprint arXiv:2506.07851*.
- Joshua Hare. 2019. Dealing with sparse rewards in reinforcement learning. *arXiv preprint arXiv:1910.09281*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, and 1 others. 2025. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint arXiv:2504.11456*.
- 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. *arXiv preprint arXiv:2103.03874*.

- Mehran Kazemi, Bahare Fatemi, Hritik Bansal, John Palowitch, Chrysovalantis Anastasiou, Sanket Vaibhav Mehta, Lalit K Jain, Virginia Aglietti, Disha Jindal, Yuanzhu Peter Chen, and 1 others. 2025. Big-bench extra hard. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26473–26501.
- Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy, Aaron Courville, and Nicolas Le Roux. 2024. Vineppo: Accurate credit assignment in rl for llm mathematical reasoning. In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24*.
- Thanh-Long V Le, Myeongho Jeon, Kim Vu, Viet Lai, and Eunho Yang. 2025. No prompt left behind: Exploiting zero-variance prompts in llm reinforcement learning via entropy-guided advantage shaping. *arXiv preprint arXiv:2509.21880*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and 1 others. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857.
- 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.
- Yang Li, Zhichen Dong, Yuhan Sun, Weixun Wang, Shaopan Xiong, Yijia Luo, Jiashun Liu, Han Lu, Jiamang Wang, Wenbo Su, and 1 others. 2025. Attention illuminates llm reasoning: The preplan-and-anchor rhythm enables fine-grained policy optimization. *arXiv preprint arXiv:2510.13554*.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*.
- Llama Meta. 2024. 3.2: Revolutionizing edge ai and vision with open, customizable models, 2024. URL: <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices>, 6.
- Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. 2024. From r to q
*: Your language model is secretly a q -function. *arXiv preprint arXiv:2404.12358*.
- 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*.
- 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*.
- Hongze Tan, Jianfei Pan, Jinghao Lin, Tao Chen, Zhihang Zheng, Zhihao Tang, and Haihua Yang. 2025. Gtpo and grpo-s: Token and sequence-level reward shaping with policy entropy. *arXiv preprint arXiv:2508.04349*.
- Qwen Team and 1 others. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2(3).
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Xumeng Wen, Zihan Liu, Shun Zheng, Shengyu Ye, Zhirong Wu, Yang Wang, Zhijian Xu, Xiao Liang, Junjie Li, Ziming Miao, and 1 others. 2025. Reinforcement learning with verifiable rewards implicitly incentivizes correct reasoning in base llms. *arXiv preprint arXiv:2506.14245*.
- Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong Zhang, Caiming Xiong, and 1 others. 2025a. A minimalist approach to llm reasoning: from rejection sampling to reinforce. *arXiv preprint arXiv:2504.11343*.
- Wei Xiong, Chenlu Ye, Baohao Liao, Hanze Dong, Xinxing Xu, Christof Monz, Jiang Bian, Nan Jiang, and Tong Zhang. 2025b. Reinforce-ada: An adaptive sampling framework for reinforce-style llm training. *arXiv preprint arXiv:2510.04996*.
- Yixuan Even Xu, Yash Savani, Fei Fang, and J Zico Kolter. 2025. Not all rollouts are useful: Down-sampling rollouts in llm reinforcement learning. *arXiv preprint arXiv:2504.13818*.
- Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang. 2025. [Learning to reason under off-policy guidance](#). *Preprint*, arXiv:2504.14945.
- 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*.
- Zhicheng Yang, Zhijiang Guo, Yinya Huang, Yongxin Wang, Dongchun Xie, Yiwei Wang, Xiaodan Liang, and Jing Tang. 2025b. Depth-breadth synergy in rlvr: Unlocking llm reasoning gains with adaptive exploration. *arXiv preprint arXiv:2508.13755*.

Qiyong Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.

Yu Yue, Yufeng Yuan, Qiyong Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiaze Chen, Chengyi Wang, Tiantian Fan, Zhengyin Du, and 1 others. 2025. Vapo: Efficient and reliable reinforcement learning for advanced reasoning tasks. *arXiv preprint arXiv:2504.05118*.

Runzhe Zhan, Yafu Li, Zhi Wang, Xiaoye Qu, Dongrui Liu, Jing Shao, Derek F Wong, and Yu Cheng. 2025. Exgrpo: Learning to reason from experience. *arXiv preprint arXiv:2510.02245*.

Xiaotian Zhang, Chunyang Li, Yi Zong, Zhengyu Ying, Liang He, and Xipeng Qiu. 2023. Evaluating the performance of large language models on gaokao benchmark. *arXiv preprint arXiv:2305.12474*.

Haizhong Zheng, Yang Zhou, Brian R Bartoldson, Bhavya Kailkhura, Fan Lai, Jiawei Zhao, and Beidi Chen. 2025a. Act only when it pays: Efficient reinforcement learning for llm reasoning via selective rollouts. *arXiv preprint arXiv:2506.02177*.

Haizhong Zheng, Yang Zhou, Brian R Bartoldson, Bhavya Kailkhura, Fan Lai, Jiawei Zhao, and Beidi Chen. 2025b. Act only when it pays: Efficient reinforcement learning for llm reasoning via selective rollouts. *arXiv preprint arXiv:2506.02177*.

A Detailed Training Settings

The complete training hyper-parameters in GRPO and LENS are put in Table 3.

Hyper-parameter	Value
Train Batch Size	128
Micro Batch Size	128
Rollout n	8
Maximum Prompt Length	512
Maximum Response Length	4096
Temperature	1.0
Top p	1.0
LR	1×10^{-6}
KL Coefficient	0.005

Table 3: Basic training hyper-parameters of both GRPO and LENS.

B Performance on Various Models

In addition to Qwen3-4B-Base, Qwen3-8B-Base, and Llama3.2-3B-Instruct, we further evaluate LENS on Qwen2.5-3B and Qwen2.5-7B. The results are summarized in Table 4. Overall, LENS consistently achieves the strongest performance across all math reasoning benchmarks.

C Success Rate Sensitivity Analysis

To investigate the impact of the success rate threshold τ on the stability of LENS, we conducted experiments with $\tau \in \{0.125, 0.25, 0.375, 0.5\}$.

Results. The main results are presented in Table 5. On one hand, $\tau = 0.125$ concentrates the training signal on the most difficult samples near the model’s capability frontier, which particularly benefits high-difficulty benchmarks such as Minerva and AMC23. On the other hand, $\tau = 0.5$ provides a broader coverage of both medium- and high-difficulty samples, leading to the best aggregate performance and enhanced stability. Accordingly, we select $\tau = 0.5$ as the default setting for all experiments.

D Computational Overhead

We report detailed runtime and memory analyses of LENS compared with standard GRPO under identical hardware configurations, utilizing the verl framework on $8 \times$ NVIDIA A800 GPUs. LENS incurs a higher wall-clock cost per update to priori-

Model	Pass@1				Average@16			Avg.
	MATH	Minerva	Olympiad	GAO	AMC23	AIME24	AIME25	
Qwen2.5-3B	56.00	26.10	25.33	44.16	12.97	1.04	0.21	23.69
+ GRPO	65.00	27.21	27.26	<u>56.62</u>	31.09	4.48	1.88	30.51
+ DAPO	<u>67.40</u>	27.21	29.48	<u>57.40</u>	34.06	4.58	1.04	31.60
+ GRESO	65.40	28.31	32.00	58.96	<u>37.19</u>	<u>6.46</u>	<u>2.50</u>	<u>32.97</u>
+ GRPO _{extended}	66.20	27.21	26.81	56.10	27.81	3.96	<u>2.29</u>	30.05
+ LENS (Ours)	68.20	<u>27.94</u>	<u>29.78</u>	58.96	38.91	7.08	3.12	33.43
Qwen2.5-7B	62.60	27.21	29.48	51.43	19.69	3.12	0.42	27.71
+ GRPO	76.60	35.29	<u>39.85</u>	65.45	51.09	11.46	5.83	40.79
+ DAPO	77.00	36.40	39.41	67.01	53.91	11.04	8.96	41.96
+ GRESO	77.80	37.50	40.89	66.49	48.28	14.17	6.67	41.69
+ GRPO _{extended}	78.00	<u>36.76</u>	38.96	69.61	51.41	15.62	<u>5.62</u>	<u>42.28</u>
+ LENS	<u>77.40</u>	37.87	40.89	<u>69.39</u>	56.41	<u>14.58</u>	10.62	43.88

Table 4: **Detailed evaluation results on seven math reasoning benchmarks.** The best and second best results are in **bold** and underlined. (*) Even in the unfavorable setting where DAPO and GRESO are trained for $2\times$ more epochs and GRPO ($n = 16$) uses more rollouts, LENS still outperforms them on the majority of benchmarks.

Value	Pass@1				Average@16			Avg.
	MATH	Minerva	Olympiad	GAO	AMC23	AIME24	AIME25	
$\tau = 0.125$	67.00	29.78	29.04	56.88	41.41	6.67	2.50	<u>33.33</u>
$\tau = 0.250$	67.40	<u>28.31</u>	29.78	<u>57.92</u>	36.88	7.50	2.50	32.90
$\tau = 0.375$	<u>68.00</u>	26.47	29.78	56.36	37.19	6.04	<u>2.92</u>	32.39
$\tau = 0.500$	68.20	27.94	29.78	58.96	<u>38.91</u>	<u>7.08</u>	3.12	33.43

Table 5: **Validation result under different τ .**

tize signal quality, the specific average step times are detailed in Table 6.

Results. LENS incurs a computational overhead ranging from $1.27\times$ to $1.62\times$ compared to standard GRPO ($G = 8$). Notably, this remains significantly more efficient than the brute-force approach of doubling sample size ($G = 16$). Crucially, this additional computation translates directly into enhanced reasoning capabilities, evidenced by the 2–3% absolute accuracy gains reported in Table 1. These results demonstrate that LENS offers a superior trade-off between efficiency and effectiveness, successfully converting marginal temporal costs into substantial reasoning capabilities.

E Pruning Strategies Analysis

To demonstrate the effectiveness of our pruning strategy, we compare LENS with three representative alternatives: (1) **Resampling**, which uses an additional n rollouts to replace unsuccessful rollouts with successful ones; (2) **Random Pruning**,

which prunes the same fraction of tokens uniformly at random per instance; and (3) **Gradient-based Pruning**, which prunes tokens with the smallest gradient norm. For fairness, all methods share the same training setup and pruning ratio, and we evaluate them on seven benchmarks.

Results. We present the results in Table 7. Our observations are as follows: LENS consistently outperforms the three baselines, achieving the best accuracy on seven benchmarks, which indicates more consistent and stable improvements than competing methods.

F Semantic Preservation of Interference Purification

A natural concern is whether removing interference tokens might distort prompt semantics. We provide two pieces of evidence showing that LENS preserves semantic integrity.

Functional Equivalence. In RLVR, rewards are computed against the unique ground-truth answer.

Model	Algorithm	Group Size (G)	Avg. Step Time (s)	Relative Cost
Qwen2.5-3B	GRPO	8	214	1.0 \times
	GRPO _{extended}	16	384	1.79 \times
	LENS (Ours)	8	333	1.56 \times
Qwen2.5-7B	GRPO	8	344	1.0 \times
	GRPO _{extended}	16	523	1.52 \times
	LENS (Ours)	8	463	1.35 \times
Qwen3-4B-Base	GRPO	8	331	1.0 \times
	GRPO _{extended}	16	524	1.58 \times
	LENS (Ours)	8	475	1.44 \times
Qwen3-8B-Base	GRPO	8	477	1.0 \times
	GRPO _{extended}	16	802	1.79 \times
	LENS (Ours)	8	726	1.62 \times
Llama-3.2-3B-Instruct	GRPO	8	229	1.0 \times
	GRPO _{extended}	16	338	1.48 \times
	LENS (Ours)	8	291	1.27 \times

Table 6: **Computational efficiency comparison on $8 \times \text{A800}$.** G denotes the group size.

Method	Pass@1				Average@16			Avg.
	MATH	Minerva	Olympiad	GAO	AMC23	AIME24	AIME25	
GRPO	65.00	27.21	27.26	56.62	31.09	4.48	1.88	30.51
Resampling	66.40	26.47	29.48	57.92	37.50	6.67	2.50	32.42
Random Pruning	65.80	26.10	29.19	57.66	37.50	5.83	2.92	32.14
Gradient-based Pruning	65.60	26.84	29.78	58.18	37.81	6.67	1.88	32.39
LENS	68.20	27.94	29.78	58.96	38.91	7.08	3.12	33.43

Table 7: **Validation result under different pruning strategies.**

Metric	Qwen3-4B	Qwen3-8B	Llama3.1-8B
F1	0.9875	0.9871	0.9877
Precision	0.9880	0.9875	0.9883
Recall	0.9871	0.9866	0.9871

Table 8: BERTScore between original and pruned prompts exceeds 0.98, showing semantic preservation.

If pruning distorted critical semantics, the generated reasoning would diverge from the correct solution and yield zero reward. The substantial post-pruning accuracy gains (Figure 3) therefore serve as functional proof that pruned prompts remain semantically equivalent to the originals.

Quantitative Verification. We compute BERTScore between original and pruned prompts across three backbones (Table 8). All F1 scores exceed 0.98, confirming strong semantic equivalence. Qualitatively, removed tokens are predominantly formatting noise (e.g., `\`, `$$`), template redundancy (e.g., `Let, Given`), and irrelevant entities

rather than semantically load-bearing content.

G Checklist

Potential Risks Our work does not involve any identifiable ethical or legal risks.

Artifacts We check that the data does not contain any information that names or uniquely identifies individual people or offensive content. All models and datasets used in this work comply with their respective open-source or research licenses. We ensure that all artifacts are used strictly within the permitted scope of their terms. The Code we released will be under a permissive open-source license, enabling reproducibility and reuse. Documentation for all artifacts will be updated and made available in the project’s GitHub repository upon release.

AI Assistants We used AI assistants (ChatGPT) solely for textual and grammatical refinement, without influencing the core content or experiments.