

STABILIZING KNOWLEDGE, PROMOTING REASONING: DUAL-TOKEN CONSTRAINTS FOR RLVR

Anonymous authors

Paper under double-blind review

ABSTRACT

Reinforcement Learning with Verifiable Rewards (RLVR) has become an effective post-training method for improving the reasoning abilities of Large Language Models (LLMs), mainly by shaping higher-order behaviors such as reflection and planning. However, previous RLVR algorithms often apply uniform training signals to all tokens, without considering the different roles of low-entropy **knowledge-related tokens** and high-entropy **reasoning-related tokens**. Some recent methods try to separate these token types by gradient masking or asynchronous updates, but these approaches may break semantic dependencies in the model output and hinder effective learning. In this work, we propose **Archer**, an entropy-aware RLVR approach with **dual-token constraints** and synchronous updates. Specifically, our method applies weaker KL regularization and higher clipping thresholds to reasoning tokens to encourage exploration, while using stronger constraints on knowledge tokens to maintain factual knowledge. Experimental results on several mathematical reasoning and code generation benchmarks show that our approach significantly outperforms previous RLVR methods, reaching or exceeding state-of-the-art performance among models of comparable size.

1 INTRODUCTION

Large Language Models (LLMs) have shown strong capabilities across various domains, as demonstrated by models like OpenAI’s “o” series (OpenAI, 2024a;b) and DeepSeek-R1 (DeepSeek-AI et al., 2025). While supervised pre-training enables LLMs to acquire vast amounts of world knowledge, post-training techniques such as Reinforcement Learning (RL) (DeepSeek-AI et al., 2025; Kimi Team et al., 2025; Yang et al., 2025a) and test-time scaling (Snell et al., 2025; Liu et al., 2025b) are crucial for enhancing their reasoning abilities. Compared to approaches like Monte Carlo Tree Search (MCTS) (Wan et al., 2024) and Process Reward Modeling (Lightman et al., 2024; Wang et al., 2024; Zhao et al., 2025), Reinforcement Learning with Verifiable Rewards (RLVR) has emerged as a simple yet effective way to further improve the reasoning abilities of LLMs (Shao et al., 2024; Yu et al., 2025).

Recent studies have revealed that RL mainly improves reasoning by better integrating and organizing the model’s existing abilities, such as reflection and planning, rather than directly changing the model’s factual memory or basic skills (e.g., arithmetic) (Gandhi et al., 2025; Vassoyan et al., 2025; Li et al., 2025). Wang et al. (2025) also show that high-entropy tokens, which often act as logical connectors, are the main focus of RLVR adjustment, while low-entropy tokens mostly capture factual or domain knowledge. These findings together suggest an important principle: during RLVR training, the behavior of tokens tied to factual knowledge (low-entropy tokens) should change little compared to the base model, while tokens related to logical reasoning (high-entropy tokens) require stronger learning signals and greater exploration.

To account for different token types in RLVR, Wang et al. (2025); Cui et al. (2025b) use gradient masking to exclude low-entropy or high-covariance tokens from updates. Meanwhile, Yang et al. (2025b) introduces an asynchronous training method, updating different token types in separate gradient steps. While these methods agree on the need to treat low-entropy and high-entropy tokens differently based on metrics like entropy or token probability, we argue that these strategies have basic limitations: **tokens within a sentence and sentences within a response are closely related through semantic and syntactic dependencies, and require coordinated learning dynamics**. Completely stopping updates to low-entropy tokens breaks these dependencies, which in turn reduces the effective optimization of high-entropy reasoning tokens, as shown in Section 4.3.1 and 4.3.2.

To address these issues, we propose a synchronized, entropy-aware framework for differentiated token training. We use a response-level entropy criterion to group tokens into two types: (1) **knowledge-related tokens**, which mainly contain factual or domain-specific knowledge, and (2) **reasoning-related tokens**, which serve as logical connectors and guide step-by-step reasoning. Unlike earlier works that use masking or asynchronous updates, our method **synchronously** updates all tokens but applies **dual-token constraints** during training. Specifically, we set a higher clip threshold and weaker KL regularization for reasoning tokens to promote exploration and learning logical patterns, while for knowledge tokens, we use a lower clip threshold and stronger regularization to maintain factual accuracy.

We evaluate our approach on challenging mathematical reasoning and code generation benchmarks. Our experiments show significant performance improvements across different tasks. Compared to the standard DAPO algorithm (Yu et al., 2025), our dual-token constraints method achieves notable gains: +6.6 Pass@1 on AIME24, +5.2 on AIME25, +3.4 on LiveCodeBench v5, and +2.6 on LiveCodeBench v6. When compared with RL-trained models with the same base model, our approach achieves state-of-the-art performance on both mathematical and coding benchmarks. Beyond *pass@1* results, further analysis shows that our method also performs better on *pass@K* metrics, indicating a higher potential for reasoning abilities. In summary, our main contributions are:

- We propose an entropy-aware dual-token constraints framework that applies different clip and KL constraints in a synchronous update manner. This preserves knowledge on low-entropy tokens while improving reasoning ability on high-entropy tokens.
- Our empirical results show that the method achieves strong performance on challenging math and code reasoning tasks, outperforming DAPO and achieving better results than similarly sized models.
- We provide a systematic study of how KL weights and clip ranges affect the balance between preserving factual knowledge and encouraging reasoning exploration, showing how they can be used to control trade-offs in RL training.

2 PRELIMINARIES

2.1 GROUP RELATIVE POLICY OPTIMIZATION

Group Relative Policy Optimization (GRPO) (Shao et al., 2024) proposes an alternative to the value-based advantage estimation used in Proximal Policy Optimization (PPO) (Schulman et al., 2017). Instead of learning a value model, GRPO estimates advantages by sampling multiple rollouts per prompt. Specifically, for a given prompt q , GRPO generates a group of responses $\{o^1, o^2, \dots, o^G\}$ and computes the corresponding rewards $\{R^1, R^2, \dots, R^G\}$. The advantage is then calculated as:

$$\hat{A}_t^i = \frac{R^i - \text{mean}(\{R^i\}_{i=1}^G)}{\text{std}(\{R^i\}_{i=1}^G)}, \quad (1)$$

where both the mean and standard deviation are computed within the sampled group. The GRPO loss is computed as:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{o^i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o^i|} \sum_{t=1}^{|o^i|} \left(\min \left(r_t^i(\theta) \hat{A}_t^i, \text{clip} \left(r_t^i(\theta), 1 - \varepsilon, 1 + \varepsilon \right) \hat{A}_t^i \right) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right) \right], \quad (2)$$

where $r_t^i = \frac{\pi_{\theta}(o_t^i|q, o_{<t}^i)}{\pi_{\theta_{\text{old}}}(o_t^i|q, o_{<t}^i)}$ denotes the importance sampling ratio, and β is a coefficient weighting the Kullback–Leibler (KL) divergence between the current policy π_{θ} and the reference policy π_{ref} .

2.2 DECOUPLE CLIP AND DYNAMIC SAMPLING POLICY OPTIMIZATION

Decouple Clip and Dynamic Sampling Policy Optimization (DAPO) (Yu et al., 2025) enhances GRPO by integrating four key techniques: Clip-Higher, Dynamic Sampling, Token-Level Policy Gradient Loss, and Overlong Reward Shaping. Similar to GRPO, DAPO samples multiple responses

per prompt and optimizes the policy using the following objective:

$$\begin{aligned} \mathcal{J}_{\text{DAPO}}(\theta) = & \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o^i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \\ & \left[\frac{1}{\sum_{i=1}^G |o^i|} \sum_{i=1}^G \sum_{t=1}^{|o^i|} \min \left(r_t^i(\theta) \hat{A}_t^i, \text{clip} \left(r_t^i(\theta), 1 - \varepsilon_{\text{low}}, 1 + \varepsilon_{\text{high}} \right) \hat{A}_t^i \right) \right] \quad (3) \\ \text{s.t. } & 0 < |\{i \in \{1, \dots, G\} \mid \text{is_equivalent}(o^i, a)\}| < G, \end{aligned}$$

where ε_{low} and $\varepsilon_{\text{high}}$ denote the lower and upper bounds of the clipping range.

3 METHOD

In this section, we introduce Archer, a novel RLVR approach with entropy-aware dual-token constraints. We begin by describing entropy-based method for identifying critical tokens (Section 3.1). Next, we discuss the limitations of prior methods in handling low-entropy tokens and motivate our approach for response-level entropy statistics (Section 3.1.1). We then analyze the necessity of joint training of high-entropy and low-entropy tokens (Section 3.2.1). Finally, we detail how Archer improves upon core constraints (clipping and KL) in previous RL algorithms by disentangling token-level optimization (Section 3.2.2).

3.1 CRITICAL TOKENS IDENTIFICATION VIA RESPONSE-LEVEL ENTROPY

Prior RL approaches like GRPO (Shao et al., 2024) and DAPO (Yu et al., 2025) typically adopt a uniform token-level optimization strength to all output tokens. This undifferentiated treatment fails to account for the distinct functional roles that different tokens play in the reasoning process (e.g., factual recall vs. logical decision points). Recent work shows that RL-driven improvements in LLM reasoning stem mainly from enhancing logical behaviors such as reflection and planning, which **integrate existing model capabilities**, rather than directly modifying the model’s factual memory or primitive skills (Yue et al., 2025a; Wen et al., 2025). Thus, during RL training, tokens associated with *factual knowledge* or *base-level skills* should largely retain their original distributions, while tokens involved in *logical reasoning and decision-making* require stronger learning signals and targeted exploration. Identifying these critical reasoning tokens is therefore a crucial first step. To address this issue, a crucial first step is to identify critical reasoning tokens.

Entropy-based Token Identification. Recent work proposes entropy as an effective signal for identifying critical tokens, observing that high-entropy tokens frequently appear at logical transition points between reasoning segments (Wang et al., 2025). In contrast, low-entropy tokens typically complete ongoing statements or syntactic structures. This observation aligns with our hypothesis that entropy discriminates between reasoning-oriented and knowledge-oriented tokens. To empirically verify this, we analyze token entropy distributions of 1024 responses (each prompt 16 times) generated by DeepSeek-R1-Distill-Qwen-1.5B during training on mathematical tasks. Following Wang et al. (2025), we visualize the top-100 highest entropy tokens and the top-100 lowest entropy tokens and retain tokens that appear more than 100 times. The visualization in Figure 1 shows that high-entropy tokens are mainly reasoning-related tokens, while most low-entropy tokens are related to factual knowledge or the suffix part of a word. These findings are also validated by recent studies (Yang et al., 2025b; Cheng et al., 2025). In summary, token entropy serves as an effective metric to distinguish between reasoning-oriented and knowledge-oriented tokens.

3.1.1 RESPONSE-LEVEL ENTROPY STATISTICS

To distinguish token types, prior works compute token entropy quantiles or covariance statistics at the batch level (Wang et al., 2025; Cui et al., 2025b). However, we find this suboptimal due to substantial entropy variation across responses from different prompts, as shown in Figure 2. For instance, some prompts yield responses with average entropy far above/below the batch mean (Figure 2 (a)); even within a single prompt, entropy can vary across sampled responses significantly (Figure 2 (b)).

Therefore, batch-level statistics for token classification introduce a key drawback: if a response’s overall entropy is low, even critical reasoning tokens may be misclassified as low-entropy, resulting in effective training. For example, using the 80th percentile as a threshold can result in only 4.34% of tokens being labeled as high-entropy in low-entropy responses. Conversely, for high-entropy responses, the proportion of high-entropy tokens may be abnormally inflated. To mitigate this, we adopt a **response-level** entropy statistics method for token classification, computing entropy quantiles

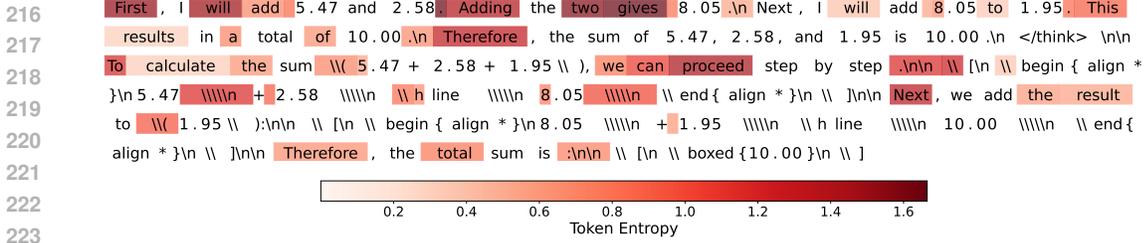


Figure 3: Visualization of high-entropy tokens (i.e., tokens with entropy larger than 80th percentile) within a response.

Clipping Constraint. To control the magnitude of policy updates at each step, we apply stricter clip ranges to knowledge-type (low-entropy) tokens to preserve the base model’s capabilities and looser clip ranges to reasoning-type (high-entropy) tokens to encourage exploratory behavior. Given a batch of responses, we first compute the entropy quantile τ_ρ^i of token entropy within each response using equation 4. Based on the computed entropy threshold, we categorize tokens into different types and assign distinct clipping ranges to each type accordingly.

$$\varepsilon(e_t^i) = \begin{cases} \varepsilon^r & \text{if } e_t^i \geq \tau_\rho^i, \\ \varepsilon^k & \text{otherwise.} \end{cases} \quad (5)$$

KL Constraint. In RL training, the KL divergence penalty is commonly used to constrain the overall deviation of the trained policy from a reference policy (Shao et al., 2024). Although recent works (Yu et al., 2025; Liu et al., 2025c; Hu et al., 2025; Chu et al., 2025; Yue et al., 2025b; He et al., 2025) advocate removing the KL divergence penalty, ProRL (Liu et al., 2025a) argues that this typically holds for base models without extensive SFT and using the KL penalty is crucial for training stability. Our experimental results also confirm that fully removing the KL penalty leads to training collapse and degraded performance, as shown in Section 4.3.1. Moreover, applying uniform KL penalties across all tokens, including high-entropy ones, significantly slows learning and reduces final performance.

Therefore, we extend the conventional KL penalty by adapting it based on the functional type of each token. Specifically, we apply a stronger KL penalty (i.e., a larger KL weight) to knowledge-type tokens (low entropy) to preserve the base model’s factual knowledge. In contrast, we apply a weaker KL penalty (i.e., a smaller KL weight) to reasoning-type tokens (high entropy), enabling greater flexibility in critical reasoning regions. The coefficients of KL constraints are as follows:

$$\beta(e_t^i) = \begin{cases} \beta^r, & \text{if } e_t^i \geq \tau_\rho^i, \\ \beta^k, & \text{otherwise.} \end{cases} \quad (6)$$

Finally, the overall objective of our algorithm is formulated as follows:

$$\mathcal{J}_{\text{TDPO}}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o^i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[\frac{1}{\sum_{i=1}^G |o^i|} \sum_{i=1}^G \sum_{t=1}^{|o^i|} \left(\min \left(r_t^i(\theta) \hat{A}_t^i, \text{clip} \left(r_t^i(\theta), 1 - \varepsilon(e_t^i), 1 + \varepsilon(e_t^i) \right) \hat{A}_t^i \right) - \beta(e_t^i) \mathbb{D}_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right) \right] \quad (7)$$

s.t. $0 < |\{i \in \{1, \dots, G\} \mid \text{is_equivalent}(o^i, a)\}| < G,$

where differentiated clipping and KL constraints are denoted using red color. The full algorithm of Archer is shown in Algorithm 1.

4 EXPERIMENTS

4.1 SETUP

Models and Baselines. We adopt DeepSeek-R1-Distill-Qwen-1.5B as the base model, which is distilled from DeepSeek-R1 (DeepSeek-AI et al., 2025) using Qwen2.5-1.5B (Yang et al.,

270 2024) as the backbone and fine-tuned on 800k high-quality reasoning data. To ensure a fair com-
 271 parison, we compare Archer against the following methods: (1) **Base Model**: The raw distilled
 272 model without further training. (2) **DAPO** (Yu et al., 2025): A RLVR algorithm that improves
 273 upon GRPO (Shao et al., 2024). (3) **DeepScaleR-1.5B** (Luo et al., 2025b): A 1.5B model trained
 274 on mathematical tasks with iterative context length expansion. (4) **DeepCoder-1.5B** (Luo et al.,
 275 2025a): A 1.5B model trained on code datasets, also utilizing context expansion strategies. (5)
 276 **FastCuRL-1.5B-V3** (Song et al., 2025): A strong 1.5B model with curriculum RL training. (6)
 277 **Nemotron-1.5B** (Liu et al., 2025a): Currently the best 1.5B reasoning model that RL-trained with
 278 DeepSeek-R1-Distill-Qwen-1.5B as the base model.

279 **Training Data.** For code domain, we construct a high-quality code training dataset from three
 280 publicly available sources: DeepCoder (Luo et al., 2025a), CodeContests (Li et al., 2022), and
 281 CodeForces (Penedo et al., 2025). Notably, CodeContests and CodeForces augment original problems
 282 with extensive test cases, which reduces false positives (i.e., incorrect solutions that pass test cases).
 283 Therefore, we prioritize these two datasets over DeepCoder in cases of duplication. After rigorous
 284 cleaning and filtering steps (detailed in Appendix C.1), we obtain a final corpus of **6,753** programming
 285 problems. For mathematics domain, we use datasets from DeepScaleR (Luo et al., 2025b), Skywork-
 286 OR1 (He et al., 2025), and DAPO (Yu et al., 2025). We merge these datasets and apply N-gram overlap
 287 removal to eliminate duplicates. After additional verification and filtering steps (see Appendix C.1),
 288 we derive a final mathematics training set of **51,800** problems.

289 **Evaluation and Metrics.** We conduct evaluation on both mathematical and coding benchmarks.
 290 For mathematics, we use six challenging datasets: AIME24 (MAA, 2024), AIME25 (MAA, 2025),
 291 AMC23 (MAA, 2023), MATH-500 (Lightman et al., 2024), Minerva Math (Lewkowycz et al.,
 292 2022), and OlympiadBench (He et al., 2024). For coding, we adopt the widely used LiveCodeBench
 293 v5 (2024.08.01-2025.02.01) and v6 (2025.02.01-2025.05.01) (Jain et al., 2025), which emphasize
 294 reasoning-intensive code generation. We use `vLLM` (Kwon et al., 2023) with temperature set to 0.8,
 295 `top_p` set to 1.0, and maximum output length set to 32,768 tokens for inference. Due to the high
 296 variance of the outputs from reasoning models, we report $avg@K$ (pass@1 performance averaged
 297 over K outputs) and $pass@K$ for each benchmark. For benchmarks with few samples (AIME24/25
 298 and AMC23), we set a larger $K=64$. We use $K=16$ for LiveCodeBench v6, $K=8$ for LiveCodeBench
 299 v5 and Minerva, and $K=4$ for MATH-500 and OlympiadBench. To ensure accurate evaluation,
 300 we adopt the verification functions from both DeepScaleR and `Math-Verify`¹ for mathematics
 301 problems.

302 **Implementation Details.** We perform RL training using the `verl` framework (Sheng et al., 2024).
 303 For DAPO-based baselines, we use clipping thresholds of $\epsilon_{low} = 0.2$ and $\epsilon_{high} = 0.28$. KL penalty
 304 loss and entropy regularization loss are omitted from the loss function. During training, we sample 16
 305 rollouts per prompt, with a temperature of 1.0 and a maximum response length of 32,768 tokens. The
 306 batch size is set to 64, the mini-batch size to 32, and the learning rate to 1×10^{-6} . For Archer, we
 307 set $\rho = 0.8$ following Wang et al. (2025). For clipping ranges and KL coefficients, we use $\epsilon^r = 0.5$,
 308 $\epsilon^k = 0.2$, $\beta^r = 0.0$, and $\beta^k = 0.001$. All experiments are conducted on 2 compute nodes, each
 309 equipped with $8 \times$ NVIDIA H800 80GB GPUs.

310 4.2 MAIN RESULTS

311 **Comparison with Base Model and DAPO.** The results in Table 1 and 2 show that our dual-token
 312 constraint training strategy leads to significant improvements on both mathematical and coding tasks.
 313 Compared to the original base model, the average accuracy increases by **18.1%** on AIME24 and
 314 **10.3%** on AIME25, resulting in an average gain of **12.3%**. On coding benchmarks, the accuracy
 315 rises by **12.7%** on LiveCodeBench v5 and **13.0%** on LiveCodeBench v6. When applying our method
 316 upon DAPO, the performance consistently exceeds that of DAPO across all benchmarks, with average
 317 gains of **5.6%** and **3.0%** for mathematical and coding tasks, respectively. These results demonstrate
 318 the effectiveness of our optimization approach.

319 **Comparison with SOTA Reasoning Models.** We also compare Archer with SOTA reasoning mod-
 320 els trained with RL using DeepSeek-R1-Distill-Qwen-1.5B as the base model. For coding
 321 tasks, our approach outperforms all comparable models, including the programming-specialized
 322 **DeepCoder-1.5B** and the general-purpose **Nemotron-1.5B**. On mathematical reasoning, our model
 323 achieves the highest average accuracy, surpassing both math-specialized models (**DeepScaleR-1.5B**,

¹<https://github.com/huggingface/Math-Verify>

Table 1: Evaluation results on mathematical benchmarks. The results of Archer are shaded and the highest values are bolded.

Method	AIME24		AIME25		AMC23		MATH-500		Minerva		Olympiad		Avg.
	avg@64	pass@64	avg@64	pass@64	avg@64	pass@64	avg@4	pass@4	avg@8	pass@8	avg@4	pass@4	
DeepSeek-R1-1.5B	30.6	80.0	23.5	63.3	70.7	100.0	83.6	92.4	27.6	48.2	44.6	59.4	46.8
DAPO	42.1	80.0	28.6	56.7	80.3	97.5	87.6	94.6	29.2	46.3	53.2	65.8	53.5
DeepScaleR-1.5B	42.0	83.3	29.0	63.3	81.3	100.0	87.7	93.6	30.3	51.1	50.7	61.0	53.5
FastCuRL-1.5B-V3	48.1	80.0	32.7	60.0	86.4	95.0	89.8	94.0	33.6	50.0	55.3	64.3	57.7
Nemotron-1.5B	48.0	76.7	33.1	60.0	86.1	97.5	90.6	93.6	35.3	47.8	59.2	66.8	58.7
Archer-Math-1.5B	48.7	83.3	33.8	70.0	86.0	97.5	90.8	94.4	35.7	51.1	59.3	67.1	59.1

Table 2: Evaluation results on code benchmarks. The results of Archer are shaded and the highest values are bolded.

Method	LCB v5 (2024.08.01-2025.02.01)		LCB v6 (2025.02.01-2025.05.01)		Avg.
	avg@8	pass@8	avg@16	pass@16	
DeepSeek-R1-1.5B	16.7	29.0	17.2	34.4	17.0
DAPO	26.0	40.5	27.6	43.5	26.8
DeepCoder-1.5B	23.3	39.1	22.6	42.0	23.0
Nemotron-1.5B	26.1	35.5	29.5	42.8	27.8
Archer-Code-1.5B	29.4	43.7	30.2	45.8	29.8

FastCuRL-1.5B-V3) and **Nemotron-1.5B**. We report the training costs of Archer and these open-source reasoning models, including the number of training steps, stages, and GPU hours in Table 3. Notably, our model achieves the best results with only **single-stage** training and **fewer** GPU hours, without the complex multi-round training used by the other methods. In addition to improvements in **pass@1**, our model also shows advantages in **pass@K** metrics, which suggests stronger reasoning diversity and higher capability limits of our method.

Table 3: Computational efficiency comparison between Archer and the baselines.

Method	Training Steps	Training Stages	GPU Hours
<i>Math RL</i>			
DeepScaleR-1.5B	1750	3	3,800 A100
FastCuRL-1.5B-V3	2620	5	—
Nemotron-1.5B	2500	8	16,000 H100
Archer-Math-1.5B	520	1	1,900 H800
<i>Code RL</i>			
DeepCoder-1.5B	—	—	—
Nemotron-1.5B	2500	8	16,000 H100
Archer-Code-1.5B	320	1	1,000 H800

4.3 ANALYSIS

4.3.1 IMPACT OF DIFFERENT KL WEIGHTS

In this part, we empirically study the impact of changing the KL penalty weight applied to low-entropy tokens during training. In addition to the default weight of 0.001 from earlier experiments, we conduct experiments with weights of 0.0 (i.e., no KL penalty) and 0.005. We calculate the average n-gram repetition ratio in generated outputs over training, which serves as a proxy for model collapse severity. The experimental results are shown in Table 4, with the corresponding training dynamics and repetition ratios visualized in Figure 5. Our results show that both the absence of KL regularization and an excessively high weight reduce performance. Specifically,

- **When KL weight = 0.0:** Figure 5(b) shows that model entropy drops rapidly, and there is a notable increase in repetition rate in Figure 5(c), indicating severe model collapse. Although performance on the LiveCodeBench v5 test set improves quickly at first, it soon levels off and shows limited gains in later stages. The final model not only underperforms the KL-regularized baseline but also falls below the standard DAPO method.

- **When KL weight = 0.005:** Entropy decreases more slowly, and repetition grows at a more gradual rate, better preserving the base model’s characteristics. However, this setting slows down learning progress, resulting in smaller performance gains.

Table 4: Model performance on LiveCodeBench v5 with varying KL weights on low-entropy tokens.

KL Weight	LiveCodeBench v5 (avg@8)
0.0	26.6
0.001	29.4
0.005	26.2

In summary, both too little and too much KL regularization hurt the final model quality. Insufficient weighting accelerates learning but makes collapse more likely, which ends up reducing performance. In contrast, excessive weighting limits learning on low-entropy tokens and thus restricts the model’s capabilities. These results **highlight the need for KL regularization on low-entropy tokens** to keep the model close to the base policy, which helps prevent collapse and retain key abilities. These observations further support our view that low-entropy tokens should be included in training, as masking them negatively affects overall learning.

4.3.2 IMPACT OF CLIP RANGES ON DIFFERENT TOKEN TYPES

We introduce different clip thresholds for different token types in equation 5. To investigate how the thresholds influence model performance, we vary the clip ranges for both high-entropy (ϵ^f) and low-entropy tokens (ϵ^k) and the results are shown in Table 5 and Figure 6, and Figure 7.

Table 5: Performance on LiveCodeBench v5 with different clip thresholds of low/high-entropy tokens.

Low-Entropy Token Clip ϵ^k	High-Entropy Clip ϵ^f	LiveCodeBench v5 (avg@8)
<i>Varying Low-Entropy Token Clip</i>		
0.1	0.4	24.6
0.2	0.4	28.7
0.3	0.4	26.0
<i>Varying High-Entropy Token Clip</i>		
0.2	0.2	27.7
0.2	0.4	28.7
0.2	0.5	29.4
0.2	0.6	26.0

Different Low-Entropy Token Clip Thresholds. As shown in Figure 6, we observe that increasing the clip threshold for low-entropy tokens produces effects similar to reducing their KL penalty weight: the model’s entropy decreases more rapidly, which leads to faster learning and earlier performance improvements. However, this also causes the repetition ratio to rise more quickly, making the model more susceptible to overfitting or collapse, which harms final performance.

On the other hand, lowering the clip threshold for low-entropy tokens has effects similar to increasing their KL weight: improvements on LiveCodeBench v5 are slower and tend to converge a lower level. Interestingly, we observe an counterintuitive entropy dynamic during training. Instead of a consistently slow decline, as seen with higher KL weights, entropy initially drops sharply, then plateaus and remains stable.

These results indicate that adjusting the clip threshold for low-entropy tokens strongly affects both the training process and the final model performance. In contrast, the model is much less sensitive to changes in the clip threshold for high-entropy tokens.

Different High-Entropy Token Clip Thresholds. As illustrated in Figure 7, increasing the clip threshold for high-entropy tokens encourages more exploration in the model’s reasoning. This leads to a slightly faster reduction in entropy during training and can improve the performance. However, these differences become more noticeable mainly in the later stages of training. In the early stages, training dynamics and LiveCodeBench v5 performance show little difference across various high-entropy clip values.

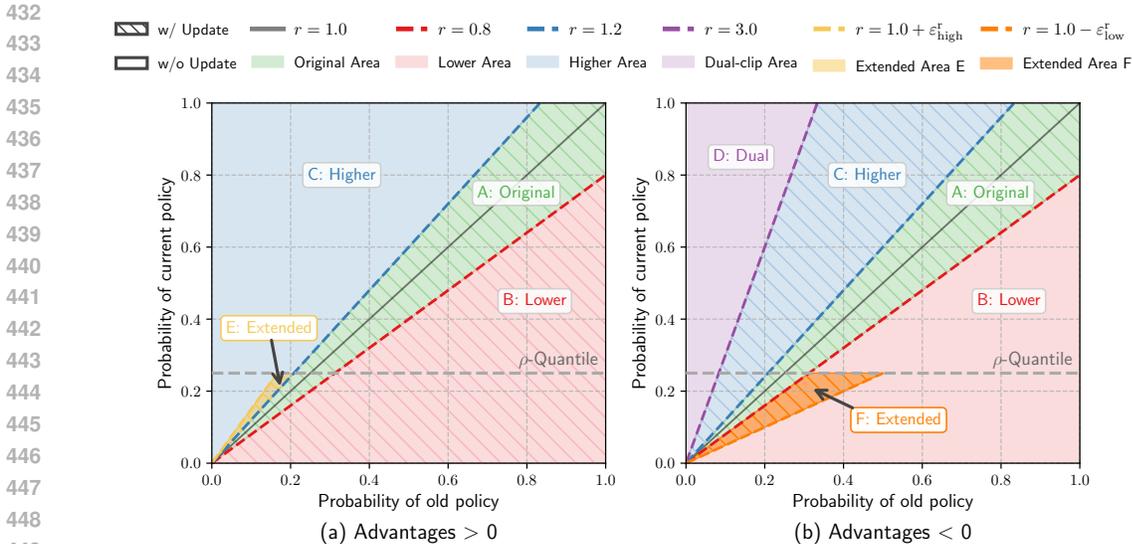


Figure 4: Visualization of PPO clip regions. The x-axis shows the sampled probability of a specific token $\pi_{\theta_{old}}$ during generation, and the y-axis shows the probability of the token under the current policy π_{θ} . Region A represents the optimization area for original GRPO. Regions B and C represent areas below and above the clipping threshold, respectively. Region D is the area for dual-clip (Ye et al., 2020). (a) When advantages > 0 , Archer optimizes region E. (b) When advantages < 0 , Archer optimizes region F.

4.3.3 VISUALIZATION OF RL OPTIMIZATION REGIONS

To better clarify the mechanism of our method, we visualize the optimization regions produced by the GRPO loss for different token types in Figure 4. Each data point in the coordinate system represents the importance sampling ratio r_t^i between the current and old policy probabilities. Figure 4(a) shows tokens with positive advantage values ($\hat{A}_t^i > 0$), while Figure 4(b) shows tokens with negative advantages ($\hat{A}_t^i < 0$). The colored regions mark the areas divided by the clipping thresholds. The shaded areas (Regions A, B for $\hat{A}_t^i > 0$ and Regions A, C for $\hat{A}_t^i < 0$) indicate where GRPO updates the model. Our method **extends the clipping boundaries for high-entropy tokens**, which are typically low-probability but are important for reasoning. As shown in Figure 4, Regions E and F correspond to the **newly extended optimization areas** introduced by Archer. Region E provides **additional reward signals** to high-entropy tokens when $\hat{A}_t^i > 0$, while Region F applies **stronger penalties** to high-entropy tokens when $\hat{A}_t^i < 0$. This design **increases the model’s focus on learning reasoning-critical tokens**.

5 CONCLUSION

In this work, we propose an entropy-aware, synchronized training framework that updates all tokens simultaneously while applying different regularization and clipping strategies depending on the type of token. By encouraging exploration on reasoning-related tokens and preserving factual correctness for knowledge-related tokens, our method balances the goals of keeping factual accuracy and improving logical reasoning. Extensive experiments on mathematical and code reasoning benchmarks show that our approach improves over the base model and outperforms existing SOTA models. These results indicate that coordinating the learning processes of different token types through entropy-aware constraints improves the reasoning abilities of LLMs. We believe this work highlights the interaction between factual knowledge and reasoning processes during RL training of LLMs, and suggests future research directions for fine-grained, token-level optimization strategies that respect the inherent structural dependencies in natural language generation.

486 REPRODUCIBILITY STATEMENT

487 We have discuss the implementation details in Section 4 and Appendix C. The code will be open-
488 sourced after paper acceptance.

490 REFERENCES

491 Shivam Agarwal, Zimin Zhang, Lifan Yuan, Jiawei Han, and Hao Peng. The unreasonable effective-
492 ness of entropy minimization in llm reasoning. *arXiv preprint arXiv:2505.15134*, 2025.

493 Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn
494 Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson
495 Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez,
496 Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario
497 Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan.
498 Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv*
499 *preprint arXiv:2204.05862*, 2022.

500 Yang Chen, Zhuolin Yang, Zihan Liu, Chankyu Lee, Peng Xu, Mohammad Shoeybi, Bryan Catanzaro,
501 and Wei Ping. Acereason-nemotron: Advancing math and code reasoning through reinforcement
502 learning. *arXiv preprint arXiv:2505.16400*, 2025.

503 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu
504 Wei. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*, 2025.

505 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario
506 Amodei. Deep reinforcement learning from human preferences. In I. Guyon, U. Von
507 Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Ad-*
508 *vances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,
509 2017. URL [https://proceedings.neurips.cc/paper_files/paper/2017/](https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf)
510 [file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf).

511 Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. Gpg: A simple and strong
512 reinforcement learning baseline for model reasoning. *arXiv preprint arXiv:2504.02546*, 2025.

513 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
514 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint*
515 *arXiv:2502.01456*, 2025a.

516 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen
517 Fan, Huayu Chen, Weize Chen, Zhiyuan Liu, Hao Peng, Lei Bai, Wanli Ouyang, Yu Cheng, Bowen
518 Zhou, and Ning Ding. The entropy mechanism of reinforcement learning for reasoning language
519 models. *arXiv preprint arXiv:2505.22617*, 2025b.

520 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,
521 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,
522 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao
523 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,
524 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,
525 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,
526 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang
527 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong,
528 Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao,
529 Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang,
530 Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang,
531 Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L.
532 Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang,
533 Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng
534 Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng
535 Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan
536 Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang,
537 Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen,

- 540 Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li,
541 Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang,
542 Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan,
543 Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia
544 He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong
545 Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha,
546 Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang,
547 Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li,
548 Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen
549 Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv
550 preprint arXiv:2501.12948*, 2025.
- 551 Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive
552 behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv
553 preprint arXiv:2503.01307*, 2025.
- 554
555 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han,
556 Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. OlympiadBench: A
557 challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific
558 problems. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd
559 Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp.
560 3828–3850, Bangkok, Thailand, August 2024. Association for Computational Linguistics. doi:
561 10.18653/v1/2024.acl-long.211. URL [https://aclanthology.org/2024.acl-long.
562 211/](https://aclanthology.org/2024.acl-long.211/).
- 563 Jujie He, Jiakai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang
564 Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint
565 arXiv:2505.22312*, 2025.
- 566
567 Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum.
568 Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base
569 model. *arXiv preprint arXiv:2503.24290*, 2025.
- 570 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
571 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free eval-
572 uation of large language models for code. In *The Thirteenth International Conference on Learning
573 Representations*, 2025. URL <https://openreview.net/forum?id=chfJJYC3iL>.
- 574
575 Amirhossein Kazemnejad, Milad Aghajohari, Eva Portelance, Alessandro Sordoni, Siva Reddy,
576 Aaron Courville, and Nicolas Le Roux. VinePPO: Refining credit assignment in RL training of
577 LLMs. In *Forty-second International Conference on Machine Learning*, 2025. URL [https:
578 //openreview.net/forum?id=Myx2kJFzAn](https://openreview.net/forum?id=Myx2kJFzAn).
- 579 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun
580 Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1.5: Scaling reinforcement learning with llms.
581 *arXiv preprint arXiv:2501.12599*, 2025.
- 582
583 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.
584 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
585 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating
586 Systems Principles*, 2023.
- 587
588 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski,
589 Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo,
590 Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving Quan-
591 titative Reasoning Problems with Language Models. In S. Koyejo, S. Mohamed,
592 A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Informa-
593 tion Processing Systems (NeurIPS)*, volume 35, pp. 3843–3857. Curran Associates, Inc.,
2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/
file/18abbef8cfe9203fdf9053c9c4fe191-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/18abbef8cfe9203fdf9053c9c4fe191-Paper-Conference.pdf).

- 594 Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh
595 Hakhmaneshi, Shishir G Patil, Matei Zaharia, et al. Llms can easily learn to reason from
596 demonstrations structure, not content, is what matters! *arXiv preprint arXiv:2502.07374*, 2025.
597
- 598 Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa
599 Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong,
600 Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. NuminaMath.
601 [<https://github.com/project-numina/aimo-progress-prize>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.
602
603
- 604 Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom
605 Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien
606 de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven
607 Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson,
608 Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code
609 generation with alphacode. *Science*, 378(6624):1092–1097, 2022. doi: 10.1126/science.abq1158.
610 URL <https://www.science.org/doi/abs/10.1126/science.abq1158>.
- 611 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
612 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *International
613 Conference on Learning Representations (ICLR)*, 2024. URL [https://openreview.net/
614 forum?id=v8L0pN6EOi](https://openreview.net/forum?id=v8L0pN6EOi).
- 615 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
616 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models.
617 *arXiv preprint arXiv:2505.24864*, 2025a.
- 618 Runze Liu, Fengshuo Bai, Yali Du, and Yaodong Yang. Meta-reward-net: Implicitly dif-
619 ferentiable reward learning for preference-based reinforcement learning. In S. Koyejo,
620 S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neu-
621 ral Information Processing Systems*, volume 35, pp. 22270–22284. Curran Associates, Inc.,
622 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/
623 file/8be9c134bb193d8bd3827d4df8488228-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/8be9c134bb193d8bd3827d4df8488228-Paper-Conference.pdf).
- 624 Runze Liu, Junqi Gao, Jian Zhao, Kaiyan Zhang, Xiu Li, Biqing Qi, Wanli Ouyang, and Bowen
625 Zhou. Can 1b llm surpass 405b llm? rethinking compute-optimal test-time scaling. *arXiv preprint
626 arXiv:2502.06703*, 2025b.
627
- 628 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min
629 Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*,
630 2025c.
- 631 Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang
632 Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, Ce Zhang, Li Er-
633 ran Li, Raluca Ada Popa, and Ion Stoica. Deepcoder: A fully open-source
634 14b coder at o3-mini level. [https://pretty-radio-b75.notion.site/
635 DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349](https://pretty-radio-b75.notion.site/DeepCoder-A-Fully-Open-Source-14B-Coder-at-O3-mini-Level-1cf81902c14680b3bee5eb349)
636 2025a. Notion Blog.
- 637 Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin
638 Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. Deepscaler: Surpassing o1-
639 preview with a 1.5b model by scaling rl. [https://pretty-radio-b75.notion.site/
640 DeepScaler-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005b](https://pretty-radio-b75.notion.site/DeepScaler-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005b)
641 2025b. Notion Blog.
- 642 MAA. American mathematics contest 12 (amc 12), November 2023. URL [https:
643 //artofproblemsolving.com/wiki/index.php/AMC_12_Problems_and_
644 Solutions](https://artofproblemsolving.com/wiki/index.php/AMC_12_Problems_and_Solutions).
- 645 MAA. American invitational mathematics examination (aime), February 2024. URL
646 [https://artofproblemsolving.com/wiki/index.php/AIME_Problems_
647 and_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions).

- 648 MAA. American invitational mathematics examination (aime), February 2025. URL
649 [https://artofproblemsolving.com/wiki/index.php/AIME_Problems_](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions)
650 [and_Solutions](https://artofproblemsolving.com/wiki/index.php/AIME_Problems_and_Solutions).
651
- 652 OpenAI. Learning to reason with llms, 2024a. URL [https://openai.com/index/](https://openai.com/index/learning-to-reason-with-llms)
653 [learning-to-reason-with-llms](https://openai.com/index/learning-to-reason-with-llms).
654
- 655 OpenAI. Openai o3-mini, 2024b. URL <https://openai.com/index/openai-o3-mini>.
656
- 657 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
658 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser
659 Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan
660 Leike, and Ryan Lowe. Training language models to follow instructions with human feed-
661 back. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Ad-*
662 *vancess in Neural Information Processing Systems*, volume 35, pp. 27730–27744. Curran Asso-
663 ciates, Inc., 2022. URL [https://proceedings.neurips.cc/paper_files/paper/](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf)
664 [2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf).
665
- 666 Guilherme Penedo, Anton Lozhkov, Hynek Kydlíček, Loubna Ben Allal, Edward Beeching,
667 Agustín Piqueres Lajarín, Quentin Gallouédec, Nathan Habib, Lewis Tunstall, and Leandro von
668 Werra. Codeforces. <https://huggingface.co/datasets/open-rl/codeforces>,
669 2025.
670
- 671 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
672 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
673
- 674 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
675 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
676 mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
677
- 678 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
679 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint*
680 *arXiv: 2409.19256*, 2024.
681
- 682 Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute
683 optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth Inter-*
684 *national Conference on Learning Representations*, 2025. URL [https://openreview.net/](https://openreview.net/forum?id=4FWAwZtd2n)
685 [forum?id=4FWAwZtd2n](https://openreview.net/forum?id=4FWAwZtd2n).
686
- 687 Mingyang Song, Mao Zheng, Zheng Li, Wenjie Yang, Xuan Luo, Yue Pan, and Feng Zhang. Fastcurl:
688 Curriculum reinforcement learning with stage-wise context scaling for efficient training rl-like
689 reasoning models. *arXiv preprint arXiv:2503.17287*, 2025.
690
- 691 Jean Vassoyan, Nathanaël Beau, and Roman Plaud. Ignore the kl penalty! boosting exploration on
692 critical tokens to enhance rl fine-tuning. *arXiv preprint arXiv:2502.06533*, 2025.
693
- 694 Ziyu Wan, Xidong Feng, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun
695 Wang. AlphaZero-like tree-search can guide large language model decoding and training. In Ruslan
696 Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and
697 Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*,
698 volume 235 of *Proceedings of Machine Learning Research*, pp. 49890–49920. PMLR, 21–27 Jul
699 2024. URL <https://proceedings.mlr.press/v235/wan24c.html>.
700
- 701 Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. In
Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
1: Long Papers), pp. 9426–9439, 2024.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,
Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen
Yu, Gao Huang, and Junyang Lin. Beyond the 80/20 rule: High-entropy minority tokens drive
effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025.

- 702 Xumeng Wen, Zihan Liu, Shun Zheng, Zhijian Xu, Shengyu Ye, Zhirong Wu, Xiao Liang, Yang
703 Wang, Junjie Li, Ziming Miao, et al. Reinforcement learning with verifiable rewards implicitly
704 incentivizes correct reasoning in base llms. *arXiv preprint arXiv:2506.14245*, 2025.
705
- 706 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
707 Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
708 Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang,
709 Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia,
710 Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu
711 Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*,
712 2024.
- 713 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
714 Gao, Chengen Huang, Chenxu Lv, Chujiu Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu,
715 Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin
716 Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang,
717 Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui
718 Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang
719 Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger
720 Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan
721 Qiu. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.
- 722 Zhihe Yang, Xufang Luo, Zilong Wang, Dongqi Han, Zhiyuan He, Dongsheng Li, and Yunjian Xu.
723 Do not let low-probability tokens over-dominate in rl for llms. *arXiv preprint arXiv:2505.12929*,
724 2025b.
- 725 Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang,
726 Xipeng Wu, Qingwei Guo, Qiaobo Chen, Yinyuting Yin, Hao Zhang, Tengfei Shi, Liang Wang,
727 Qiang Fu, Wei Yang, and Lanxiao Huang. Mastering complex control in moba games with deep
728 reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):
729 6672–6679, Apr. 2020. doi: 10.1609/aaai.v34i04.6144. URL [https://ojs.aaai.org/
730 index.php/AAAI/article/view/6144](https://ojs.aaai.org/index.php/AAAI/article/view/6144).
- 731 Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
732 Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi
733 Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiase Chen, Jiangjie Chen, Chengyi
734 Wang, Hongli Yu, Weinan Dai, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying
735 Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source
736 llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025.
- 737 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. Does
738 reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv
739 preprint arXiv:2504.13837*, 2025a.
740
- 741 Yu Yue, Yufeng Yuan, Qiyang Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiase Chen, Chengyi
742 Wang, TianTian Fan, Zhengyin Du, Xiangpeng Wei, Xiangyu Yu, Gaohong Liu, Juncai Liu,
743 Lingjun Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu,
744 Ru Zhang, Xin Liu, Mingxuan Wang, Yonghui Wu, and Lin Yan. Vapo: Efficient and reliable
745 reinforcement learning for advanced reasoning tasks. *arXiv preprint arXiv:2504.05118*, 2025b.
- 746 Kaiwen Zha, Zhengqi Gao, Maohao Shen, Zhang-Wei Hong, Duane S Boning, and Dina Katabi.
747 RL tango: Reinforcing generator and verifier together for language reasoning. *arXiv preprint
748 arXiv:2505.15034*, 2025.
- 749 Jian Zhao, Runze Liu, Kaiyan Zhang, Zhimu Zhou, Junqi Gao, Dong Li, Jiafei Lyu, Zhouyi Qian,
750 Biqing Qi, Xiu Li, et al. Genprm: Scaling test-time compute of process reward models via
751 generative reasoning. *arXiv preprint arXiv:2504.00891*, 2025.
752
753
754
755

A FULL ALGORITHM

Algorithm 1 Archer

Input: Base model π_{base} , prompt dataset \mathcal{D} , quantile level ρ , clipping thresholds $\varepsilon^r, \varepsilon^k$, KL coefficients β^r, β^k

- 1: Initialize policy model $\pi_\theta \leftarrow \pi_{\text{base}}$ and reference model $\pi_{\text{ref}} \leftarrow \pi_{\text{base}}$
- 2: **for** step = 1, 2, ..., T **do**
- 3: Sample a batch of prompts \mathcal{D}_b from \mathcal{D}
- 4: Generate responses $\{o^i\}_{i=1}^G$ for each prompt q in the batch
- 5: **for** each response $|o^i|$ **do**
- 6: Compute the ρ -quantile of token entropy τ_ρ^i with equation 4
- 7: Compute clipping thresholds and coefficients of KL penalty with equation 5 and equation 6
- 8: **end for**
- 9: Update the policy model π_θ using equation 7
- 10: **end for**

B RELATED WORK

B.1 REINFORCEMENT LEARNING FOR LARGE LANGUAGE MODELS

Previous works have shown that RL, particularly Reinforcement Learning from Human Feedback (RLHF) (Christiano et al., 2017; Liu et al., 2022), is an effective tool for aligning LLMs with human preferences (Ouyang et al., 2022; Bai et al., 2022). With the recent success of scaling RL in LLMs (OpenAI, 2024a; DeepSeek-AI et al., 2025; Kimi Team et al., 2025), RLVR has emerged as an effective method to improve the reasoning ability of LLMs using rule-based rewards. However, approaches like GRPO (Shao et al., 2024) and its extensions (Yu et al., 2025; Liu et al., 2025c; Chu et al., 2025; Yue et al., 2025b; He et al., 2025) rely on response-level learning signals, which uniformly assign the same advantage value to all tokens within a response. This uniform treatment overlooks the distinct roles tokens play during reasoning (e.g., factual recall vs. logical inference), potentially leading to suboptimal learning at critical reasoning steps and limiting overall performance gains. Although process-based RL (Kazemnejad et al., 2025; Cui et al., 2025a; Zha et al., 2025) and unsupervised RL (Agarwal et al., 2025; Cheng et al., 2025) provide fine-grained rewards for RL optimization, they still lack consideration for the functions of different tokens.

B.2 CRITICAL TOKEN ANALYSIS IN RL FOR REASONING

Several recent studies have provided token-level analyses of RLVR training (Yang et al., 2025b; Cui et al., 2025b; Wang et al., 2025; Cheng et al., 2025). Yang et al. (2025b) observe that low-probability tokens, often exhibiting high entropy, dominate the RL updates and the update of high-probability tokens are suppressed. Cui et al. (2025b) show that changes in policy entropy are linked to the covariance between action probabilities and advantages. Wang et al. (2025) identify high-entropy tokens, referred to as “forking tokens”, as logical connectors. Cheng et al. (2025) further associate high-entropy tokens with reasoning-related behaviors, such as logical transitions and self-reflection. Unlike prior works that either completely isolate low-entropy tokens (Wang et al., 2025) or high-covariance tokens (Wang et al., 2025; Cui et al., 2025b), or train them separately (Yang et al., 2025b), our approach employs joint training. While we similarly utilize entropy to distinguish between logic-oriented and knowledge-oriented tokens, we avoid direct filtering or separation. Instead, we apply differentiated training constraints, enabling us to preserve the capabilities of the base model while simultaneously encouraging more effective exploration during training.

C EXPERIMENTAL DETAILS

C.1 DATASET

C.1.1 CODE DOMAIN

Data Sources and Integration. The code dataset is compiled from three publicly available sources: DeepCoder, CodeContests, and CodeForces. Notably, CodeContests and CodeForces extend their original problem sets with a larger number of test cases, improving the reliability of evaluation and reducing the incidence of false positives—i.e., incorrect code that inadvertently passes tests. As such,

810 these two datasets are prioritized. In cases of duplication with DeepCoder, we retain the entries from
811 either CodeContests or CodeForces.

812 **Data Cleaning and Filtering Pipeline.** We apply a rigorous multi-stage cleaning and selection
813 process to ensure dataset quality:
814

- 815 1. **Test Case Preprocessing:** We remove illustrative test cases embedded in problem descriptions
816 and discard problems with fewer than five test cases, which are more susceptible to false
817 positives.
- 818 2. **Model Validation and Difficulty Filtering:** Each problem is evaluated using 8-sample gen-
819 eration with a strong language model (Qwen3-30B-A3B (Yang et al., 2025a)). We exclude
820 problems for which all samples fail verification, filtering out flawed questions (e.g., with invalid
821 test cases), overly long I/O problems beyond the verifier’s capacity, or those that are excessively
822 difficult—even for strong models. This reduces potential false negatives.
- 823 3. **Problem Deduplication:** We perform N-gram-level deduplication to eliminate duplicate ques-
824 tions within the training corpus.
- 825 4. **Test Set Contamination Prevention:** To prevent data leakage, we remove any overlapping
826 problems by conducting N-gram-level deduplication against the evaluation set of LiveCodeBench
827 v5.
828
- 829 5. **Sampling Stability Filtering:** Using a warm-start model
830 (DeepSeek-R1-Distill-Qwen-1.5B), we generate 8 additional samples per problem.
831 We remove problems where all generations are either completely correct or completely incorrect,
832 thereby ensuring sufficient learning signal and gradient diversity.

833 **Data Standardization.** All retained code problems are reformatted into either *function-call* or
834 *stdin/stdout* formats, enabling consistent and automated validation via a code verifier.
835

836 **Final Dataset.** Following the aforementioned pipeline, we construct a high-quality code training
837 dataset consisting of **6,753** problems.

838 C.1.2 MATHEMATICS DOMAIN

839 **Data Sources and Integration.** For the mathematics domain, we leverage existing curated datasets
840 rather than raw symbolic corpora such as NuminaMath (LI et al., 2024). Specifically, we integrate
841 three high-quality, verifiable datasets: DeepScaleR, Skywork-OR1, and DAPO. The datasets are
842 merged and deduplicated using N-gram overlap removal to eliminate redundancy.

843 **Data Cleaning and Filtering Pipeline.**
844

- 845 1. **Model Validation and Filtering:** Each math problem undergoes 8-sample generation using
846 the Qwen3-30B-A3B model, followed by verification using a mathematical logic verifier.
847 Problems for which all samples fail are excluded to remove noise, overly complex items, or
848 verification bottlenecks that might cause false negatives.
- 849 2. **Sampling Stability Filtering:** We repeat the 8-sample generation process using a warm-start
850 model (DeepSeek-R1-Distill-Qwen-1.5B) and discard problems with homogeneous
851 sampling outcomes (i.e., all correct or all incorrect).
- 852 3. **Test Set Contamination Prevention:** To avoid contamination of evaluation benchmarks, we
853 perform N-gram deduplication against the AMC competition datasets (AIME24 and AIME25),
854 ensuring zero overlap.
855

856 **Final Dataset.** After rigorous verification and filtering, we obtain a final mathematics training
857 corpus comprising approximately **51,800** high-quality problems suitable for reinforcement learning.
858

859 D ADDITIONAL EXPERIMENTAL RESULTS

860 D.1 IMPACT OF CLIP RANGES ON HIGH-ENTROPY TOKENS

861 D.2 MUTUAL ENHANCEMENT BETWEEN MATH RL AND CODE RL

862 Figure 8 shows results on AIME24, AIME25, and LiveCodeBench v5, comparing RL applied to
863 math tasks (math RL) and code tasks (code RL). We observe that RL training in either domain

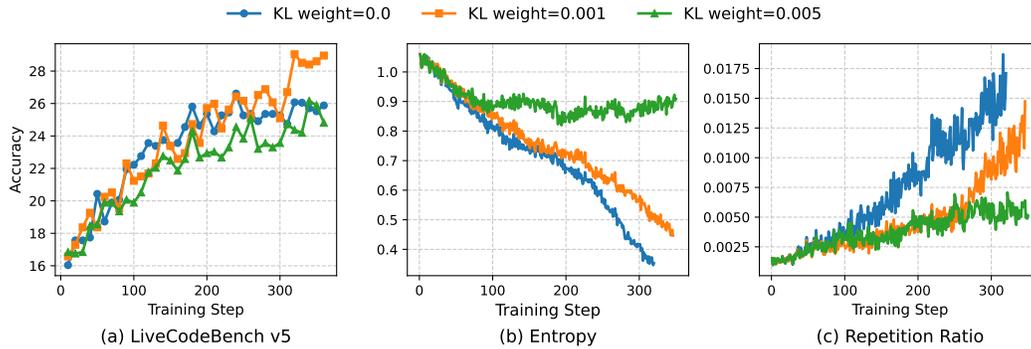


Figure 5: Effects of varying KL weights on (a) model performance on LiveCodeBench v5, (b) model entropy, and (c) repetition ratio.

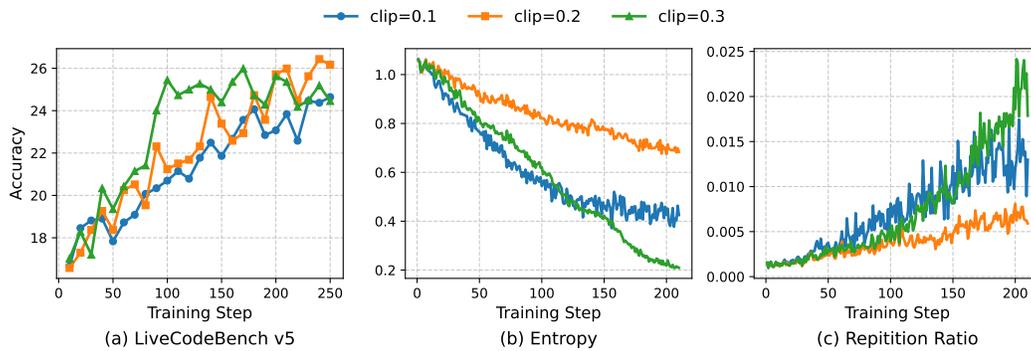


Figure 6: Effects of varying the clip threshold of low-entropy tokens on (a) model performance on LiveCodeBench v5, (b) model entropy, and (c) repetition ratio.

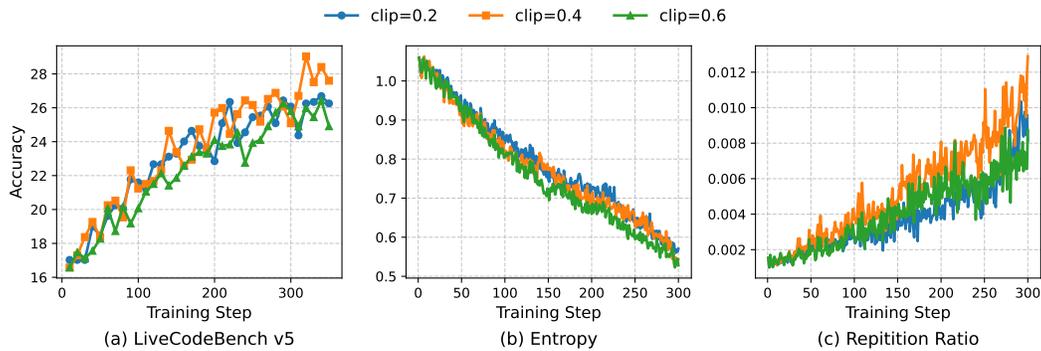


Figure 7: Effects of varying clip value on high-entropy tokens on (a) model performance on LiveCodeBench v5, (b) model entropy, and (c) repetition ratio.

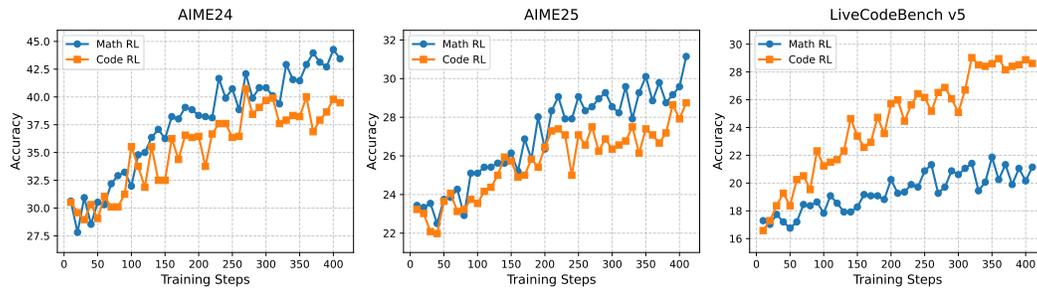


Figure 8: Model performance on AIME24, AIME25, and LiveCodeBench v5 of math RL and code RL.

leads to significant performance improvements not only in-domain but also on out-of-domain (OOD) benchmarks.

To analyze the source of these cross-domain improvements, we evaluate the base model and its math/code RL variants on OOD benchmarks (LiveCodeBench v5 and AIME24/25), measuring problem-level accuracy across all tasks. Unlike AceReason-Nemotron (Chen et al., 2025), which attributes the benefits of math RL on code tasks primarily to the presence of math-related subdomains (e.g., Algebra, Counting, Combinatorics), our results suggest a different explanation: performance improvements correlate more strongly with the intrinsic difficulty of the problems rather than their topical categories. Specifically, problems where the base model already achieves relatively high accuracy tend to benefit most from RL training, as shown in Figure 9 and Figure 10.

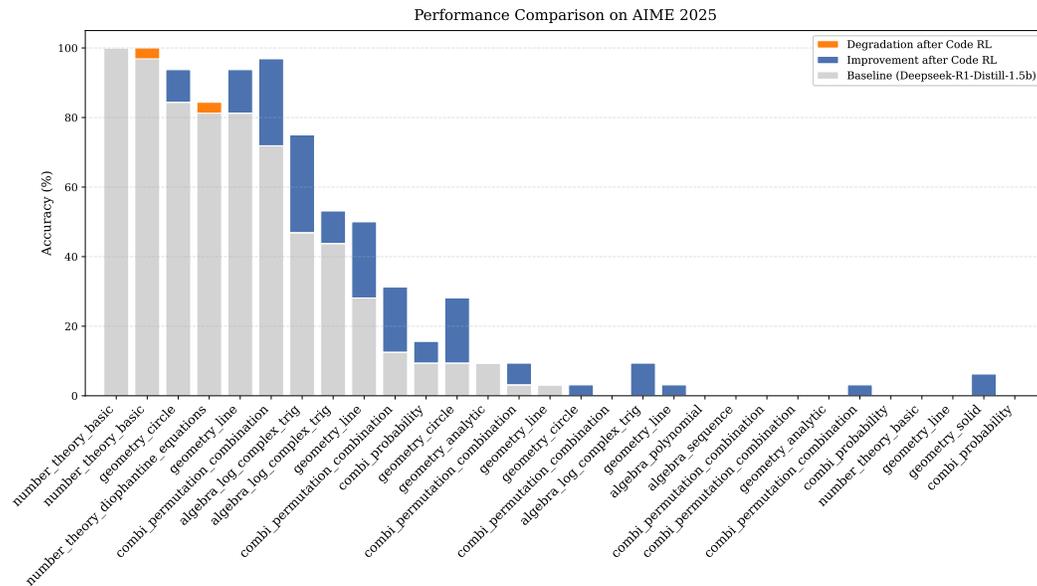
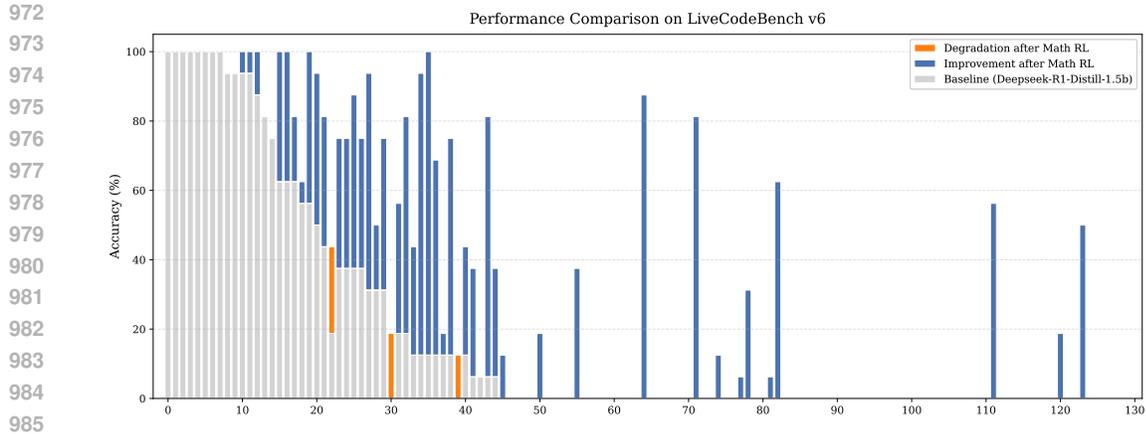


Figure 9: Problem-level accuracy comparison between the base model and RL-trained model.

A closer analysis of the problems with notable improvement in Figure 9 shows that RL training **does not introduce fundamentally new knowledge** beyond what is already present in the base model’s outputs. This observation applies to both less challenging problems (where the base model already performs well) and more challenging ones. Instead, the improvements **mainly result from enhanced reasoning capabilities**. We identify three main areas of improvement:

- **Enhanced Structural Organization:** Responses demonstrate a clearer logical flow and improved structural coherence.



986 Figure 10: Problem-level accuracy comparison on LiveCodeBench v6 between the base model and
 987 Math RL trained model.

- 988
- 989
- 990
- 991
- 992
- 993
- 994
- 995
- 996
- **Increased Attention to Details:** Models are more careful with edge cases and boundary conditions. This effect is especially clear in the Code-RL model, likely because boundary handling is important in programming tasks.
 - **Improved Contextual Consistency:** RL-trained models are more accurate at integrating and summarizing previous reasoning steps. In contrast, the base model sometimes produces final answers based on incorrect intermediate reasoning even if some steps are correct, which leads to inconsistencies.

997 These findings further support our main claim: the main way RL improves model capability **is not**
 998 **by changing stored knowledge or basic skills** (such as arithmetic), **but by better integrating and**
 999 **optimizing existing abilities** through structured logical behavior such as reflection and planning.
 1000 At the same time, this provides empirical support for the effectiveness of our proposed dual-token
 1001 constraint training strategy.

1002 E THE USE OF LARGE LANGUAGE MODELS

1004 We only use LLMs to polish some sections instead of directly using them for writing all words in this
 1005 paper.

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025