SLOW-FAST POLICY OPTIMIZATION: REPOSITION-BEFORE-UPDATE FOR LLM REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Reinforcement learning (RL) has become central to enhancing reasoning in large language models (LLMs). Yet on-policy algorithms such as Group Relative Policy Optimization (GRPO) often suffer in early training: noisy gradients from low-quality rollouts lead to unstable updates and inefficient exploration. We introduce Slow-Fast Policy Optimization (SFPO), a simple yet efficient mechanism to address the above limitations via decomposing each iteration into three stages: a short fast trajectory of inner steps on the same batch, a reposition step to control off-policy drift, and a final slow correction. This reposition-before-update design preserves the objective and rollout process unchanged, making SFPO plug-compatible with existing policy-gradient pipelines. Extensive experiments demonstrate that SFPO consistently improves stability, reduces number of rollouts, and accelerates convergence of reasoning RL training. Specifically, it outperforms GRPO by up to 2.80 points on math reasoning benchmarks. It also achieves up to 4.93× fewer rollouts and a 4.19× reduction in wall-clock time to match GRPO's best accuracy.

1 Introduction

Large language models (LLMs) have recently achieved remarkable progress on complex multi-step reasoning tasks, especially in mathematics problem solving and scientific question answering (OpenAI et al., 2024; DeepSeek-AI et al., 2025b). A central driver of these advances has been reinforcement learning (RL), which fine-tunes LLMs using reward signals tied to semantic correctness and solution quality. For example, DeepSeekMath (Shao et al., 2024) demonstrates that Group Relative Policy Optimization (GRPO) can substantially improve open-source LMs on mathematical reasoning, while DeepSeek-R1 (Guo et al., 2025) shows that RL signals alone can induce emergent reasoning behaviors across various domains.

Despite these successes, GRPO inherits structural inefficiencies that make training fragile in the context of LLM reasoning. During the early stages of training, when rollouts are weak or uninformative, stochastic rewards induce high-variance gradients that destabilize updates. Prior studies corroborate this issue: early rollouts often fail to provide useful signals (Zheng et al., 2025), reward shaping can trigger sudden collapse (Dai et al., 2025), and noisy group normalization further impairs learning under imbalanced responses (Shen et al., 2025). Although group normalization partially mitigates variance, GRPO still applies only a single gradient update per batch, discarding information that multiple inner steps could exploit (Schulman et al., 2017; Engstrom et al., 2020). As a result, sample efficiency remains low and early exploration highly brittle, with progress overly sensitive to noise and rollout quality in LLM reasoning tasks.

In this work, we introduce **Slow–Fast Policy Optimization** (**SFPO**), a simple yet efficient mechanism that augments on-policy policy-gradient methods such as GRPO while keeping the loss, rollout collection, and KL/clip regularization intact. Each iteration is restructured into three coordinated stages: (i) **fast**—multiple inner updates on the same batch to stabilize the search direction; (ii) **reposition**—interpolation back toward the starting point to control off-policy drift; and (iii) **slow**—an extra gradient correction to align with local curvature. This reposition-before-update design transforms noisy one-shot updates into structured trajectories, yielding more stable optimization, higher sample efficiency, and faster convergence. We summarize our contributions as follows:

• We propose SFPO, a plug-compatible update rule that reuses rollout batches via a fast-reposition-slow decomposition, with theoretical insights for each stage.

Figure 1: Pipeline of SFPO at iteration s. Starting from the current policy $\pi_{\theta^{s,0}}$, we first generate rollouts for training. **Stage I (Fast Trajectory):** apply K successive gradient updates on the same batch to obtain $\theta^{s,K}$. **Stage II (Reposition):** interpolate between $\theta^{s,K}$ and the starting point $\theta^{s,0}$ to form $\widetilde{\theta}^{s,K}$, controlling off-policy drift. **Stage III (Slow Correction):** perform one additional update on $\widetilde{\theta}^{s,K}$, yielding $\pi_{\theta^{s+1,0}}$ for the next iteration.

- SFPO introduces no changes to the underlying objective, rollout generation, or regularization, enabling drop-in integration into existing LLM reasoning pipelines.
- Through extensive experiments on math reasoning benchmarks, we show that SFPO consistently improves stability, sample efficiency, and convergence speed over GRPO.

2 Preliminaries

Group Relative Policy Optimization (GRPO) (Shao et al., 2024) is a policy gradient method that has been widely applied in large-scale LLM training pipelines. GRPO dispenses with an explicit value function and instead normalizes rewards across a group of responses to the same prompt. Given an input q, the policy π_{θ} generates G candidate sequences $\{o_i\}_{i=1}^G$ with corresponding rewards $\{r_i\}_{i=1}^G$. The normalized advantage for each candidate is defined as

$$\widehat{A}_{i} = \frac{r_{i} - mean(\{r_{i}\}_{i=1}^{G})}{std(\{r_{i}\}_{i=1}^{G})}.$$
(1)

This group-based normalization can be interpreted as a form of reward shaping: by emphasizing relative differences among candidate sequences for the same input, GRPO strengthens the reliability of the gradient signal. Instead of embedding a KL penalty inside the reward, GRPO directly regularizes the policy by including an explicit KL term between the learned policy and a reference policy. The training objective is

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q \sim P(\mathcal{Q}), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(O|q)} \\
\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \min(r_{i,t}(\theta) \widehat{A}_{i,t}, clip(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \widehat{A}_{i,t}) - \beta D_{KL}[\pi_{\theta} \| \pi_{ref}] \tag{2}$$

where $r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t}|q,o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q,o_{i,<t})}$ is the token-level importance ratio between the new and old policies, ϵ is the clipping range that prevents overly aggressive updates, β controls the KL regularization strength, and π_{ref} is a fixed reference policy. Here, q denotes the input prompt, o_i a generated sequence, and $o_{i,t}$ its t-th token.

Limitations of GRPO. Despite its popularity, GRPO inherits structural drawbacks from the underlying on-policy update rule. At each iteration, the parameters are updated via a **single stochastic gradient step** estimated from a batch of rollouts. The randomness of rewards leads to high-variance gradient estimates, making updates unstable. Group normalization partially dampens this effect but remains sensitive to fluctuations within each batch. Moreover, restricting each batch to a single update discards potentially useful gradient information across inner steps, leading to inefficient data use and limited variance reduction. These drawbacks underscore the need for an update mechanism that can stabilize gradient directions while making more effective use of available samples.

```
108
              Algorithm 1 SFPO: unified fast-reposition-slow update.
109
              Require: Initial policy \pi_{\theta^{0,0}}, dataset \mathcal{D}, hyperparameters S, K, \alpha_0, \eta, \omega, \tau, loss \mathcal{L}(\theta)
110
                  Initialize rolling buffer \mathcal{H} \leftarrow \emptyset, stats (\mu, \sigma) \leftarrow (0, 1), trigger index s^* \leftarrow +\infty
111
                  for s = 0, 1, ..., S - 1 do
112
                      Generate rollouts with the current policy \pi_{\theta^s,0} on prompts from \mathcal{D}.
113
                      \alpha \leftarrow \alpha_0 \cdot \mathbf{1}[s < s^*] // Set \alpha for this iteration from past trigger (non-anticipatory)
114
                      if \alpha = 0 then
                          \widetilde{\widetilde{\theta}}^{s,K} \leftarrow \theta^{s,0}
                                                  // Skip fast trajectory & reposition
115
116
                           \begin{array}{l} \textbf{for } k = 0, 1, \dots, K-1 \textbf{ do} \\ \theta^{s,k+1} \leftarrow \theta^{s,k} - \eta \nabla_{\theta} \mathcal{L}(\theta^{s,k}) & \textit{ // Stage I: Fast Trajectory} \end{array} 
117
118
119
                          \tilde{\theta}^{s,K} \leftarrow \theta^{s,0} + \alpha(\theta^{s,K} - \theta^{s,0}) // Stage II: Reposition
120
121
                      \theta^{s+1,0} \leftarrow \widetilde{\theta}^{s,K} - \eta \nabla_{\theta} \mathcal{L}(\widetilde{\theta}^{s,K}) // Stage III: Slow Correction
122
                      Compute entropy H_s; update rolling buffer \mathcal{H} (keep last \omega ones) and (\mu_s, \sigma_s).
123
                      Z_s \leftarrow \frac{H_s - \mu_s}{\sigma_s + \varepsilon} (\varepsilon for numerical stability)
124
                      if s^* = +\infty and |Z_s| \ge \tau then
125
                          s^* \leftarrow s + 1 // will set \alpha = 0 for all future s' \ge s^*
126
                      end if
127
                  end for
128
                  return final policy \pi_{\theta^{S,0}}
```

3 Method

To address the instability and inefficiency of one-shot policy updates, we propose **Slow-Fast Policy Optimization (SFPO)**, a simple yet general update mechanism that reuses rollouts more effectively while remaining plug-compatible with standard on-policy policy gradient algorithms. Each iteration consists of three stages: (i) a *fast trajectory* of multiple inner updates on the same batch, (ii) a *reposition* step that interpolates back toward the on-policy point to control drift, and (iii) a *slow correction* via an extra-gradient update. As illustrated in Fig. 1 and Alg. 1, this fast–reposition–slow design transforms noisy one-shot updates into well-structured update trajectories, yielding more stable optimization and higher sample efficiency without additional rollouts.

Notation. To align with standard gradient-based optimization notation, we define $\mathcal{L}(\theta) = -\mathcal{J}(\theta)$, so that maximizing the objective \mathcal{J} is equivalent to minimizing the loss \mathcal{L} .

3.1 STAGE I: FAST TRAJECTORY

In standard on-policy policy-gradient methods such as GRPO, each outer iteration is updated by a single stochastic gradient:

$$\theta^{s+1} = \theta^s - \eta \nabla_{\theta} \mathcal{L}(\theta^s), \tag{3}$$

where $\nabla_{\theta} \mathcal{L}(\theta^s)$ is estimated from one batch of rollouts. Such one-shot updates suffer from high variance and often drive the policy in unstable directions, especially during early training. SFPO mitigates this by performing multiple inner updates on the *same* batch or rollouts.

Formally, starting from parameters $\theta^{s,0}$ at the beginning of iteration s, we execute a short fast trajectory of K inner updates:

$$\theta^{s,k+1} = \theta^{s,k} - \eta \nabla_{\theta} \mathcal{L}(\theta^{s,k}), \qquad k = 0, \dots, K - 1.$$
(4)

This produces a sequence $\theta^{s,0} \to \theta^{s,1} \to \cdots \to \theta^{s,K}$, where each step refines the gradient direction using the same rollout data.

Intuition. Unlike one-shot updates that rely on a single noisy gradient, the displacement

$$\theta^{s,K} - \theta^{s,0} = -\eta \sum_{k=0}^{K-1} \nabla_{\theta} \mathcal{L}(\theta^{s,k})$$
 (5)

captures the cumulative effect of K sequential corrections. Even if some individual steps are perturbed by noise, their composition tends to damp idiosyncratic fluctuations and align with the underlying gradient direction. Geometrically, this can be viewed as integrating the local gradient field along a short trajectory in parameter space rather than trusting a single noisy vector at $\theta^{s,0}$. As a result, the update direction at the end of Stage I is typically more stable and less sensitive to randomness in any single gradient estimate.

Compact Mathematical Intuition. Let $\theta_k := \theta^{s,k}$ and consider the K inner steps $\theta_{k+1} = \theta_k - \eta \nabla \mathcal{L}(\theta_k)$ starting at $\theta_0 = \theta^{s,0}$. In a small neighborhood, linearizing the gradient field, $\nabla \mathcal{L}(\theta) \approx g_0 + H_0(\theta - \theta_0)$ with $g_0 = \nabla \mathcal{L}(\theta_0)$ and $H_0 = \nabla^2 \mathcal{L}(\theta_0)$, yields the closed-form displacement

$$\theta_K - \theta_0 \approx -[I - (I - \eta H_0)^K] H_0^{\dagger} g_0,$$
 (6)

where H_0^\dagger denotes the (pseudo)inverse on the range of H_0 (the $\lambda \to 0$ case is understood by continuity). Spectrally, along an eigen-direction with curvature $\lambda \geq 0$, the scalar gain is $\left(1-\left(1-\eta\lambda\right)^K\right)/\lambda$: for small λ it behaves like $K\eta$ (steadily accumulating progress in gentle directions), while for larger λ it saturates below 1 (damping stiff-direction oscillations). Thus the fast trajectory acts as a curvature-aware low-pass filter that stabilizes the endpoint direction relative to a one-shot step. We use this as local intuition under a sufficiently small step size (e.g., $\eta < 1/\|H_0\|_2$) and in neighborhoods where positive-curvature directions dominate; in practice, KL/clip regularization and Stage II reposition mitigate adverse effects of negative curvature.

3.2 STAGE II: REPOSITION

While the fast trajectory of Stage I improves stability, it also *changes the nature of the update from on-policy to off-policy*. Since all inner steps $\theta^{s,1},\ldots,\theta^{s,K}$ reuse the same rollouts generated at $\theta^{s,0}$, the endpoint $\theta^{s,K}$ no longer corresponds to the distribution that produced those samples. This *distribution mismatch* is a fundamental drawback of off-policy learning, as it biases gradient estimates and can destabilize training.

Inspired by Lookahead Optimization (Zhang et al., 2019), SFPO introduces a *reposition step* that interpolates the fast trajectory back toward its starting point:

$$\widetilde{\theta}^{s,K} = \theta^{s,0} + \alpha(\theta^{s,K} - \theta^{s,0}), \qquad \alpha \in [0,1]. \tag{7}$$

Here α regulates the degree of off-policy drift: smaller values keep the update close to the original on-policy iterate, while larger values rely more on the fast trajectory at the risk of greater mismatch.

Intuition. The interpolation is equivalent to solving a linearized proximal subproblem around $\theta^{s,0}$:

$$\min_{\theta} \left\langle \sum_{k=0}^{K-1} \nabla_{\theta} \mathcal{L}(\theta^{s,k}), \ \theta - \theta^{s,0} \right\rangle + \frac{\lambda}{2} \|\theta - \theta^{s,0}\|^2, \tag{8}$$

whose unique solution coincides with $\widetilde{\theta}^{s,K}$ for $\lambda = \frac{1}{\alpha\eta}$. Thus, α acts as an *implicit trust-region radius*: smaller α implies a larger proximal weight λ , enforcing stronger contraction toward the on-policy point.

3.3 STAGE III: SLOW CORRECTION

After repositioning, SFPO applies one more (slow) correction step at the interpolated point:

$$\theta^{s+1} = \widetilde{\theta}^{s,K} - \eta \nabla_{\theta} \mathcal{L}(\widetilde{\theta}^{s,K}). \tag{9}$$

This yields a *predictor—corrector* structure: Stage I produces a stabilized *fast trajectory*, Stage II tempers off-policy drift via *reposition*, and Stage III applies a *slow correction* aligned with the local curvature at the update point.

 $^{^1}$ A heuristic alternative is obtained if one replaces $\sum_k \nabla_\theta \mathcal{L}(\theta^{s,k})$ by the averaged gradient \bar{g} and uses the approximation $\theta^{s,K} - \theta^{s,0} \approx -\eta K \, \bar{g}$, leading to $\lambda \approx 1/(\alpha \eta K)$.

Theoretical intuition. Under L-smoothness and sufficiently small η , the descent lemma implies

$$\mathbb{E}[\mathcal{L}(\theta^{s+1})] \le \mathcal{L}(\theta^{s,0}) - c \, \eta \|\nabla \mathcal{L}(\theta^{s,0})\|^2 + O(\eta^2 L \cdot \mathcal{F}(K,\alpha)), \tag{10}$$

where $\mathcal{F}(K,\alpha)$ represents the combined effect of Stage I (fast trajectory length K) and Stage II (reposition factor α). Intuitively, $\mathcal{F}(K,\alpha)$ reflects a balance between exploiting more gradient information and controlling distributional drift: increasing K leverages the same rollout data across multiple steps but also amplifies off-policy mismatch, while larger α interpolates more aggressively toward the fast trajectory at the risk of greater instability. In practice, since increasing K raises both wall-clock cost and $\mathcal{F}(K,\alpha)$, we adopt small K with moderately large α , while the slow correction in Stage III mitigates overshooting and preserves progress along the stabilized trajectory direction.

3.4 Scheduling α

A nonzero α is essential to exploit the stabilized fast trajectory when K>0: if $\alpha=0$, the reposition collapses to $\theta^{s,0}$ and the fast trajectory is discarded, eliminating the benefit of Stage I. However, the same aggressiveness that helps early can be counter-productive near a minimizer. When $\|\nabla \mathcal{L}\|$ is large, a nonzero α accelerates progress by moving along the stabilized fast direction; but as we approach a minimum, the signal weakens while curvature and stochastic noise dominate, so a large α amplifies drift and instability. This motivates an *adaptive schedule for* α .

Why adapt α rather than K? K controls both stabilization and wall-clock runtime: increasing K reduces oscillations but incurs proportional compute cost, making mid-training changes impractical. By contrast, α is a *soft trust parameter*: it decides how much of the stabilized fast trajectory is exploited, without changing runtime or discarding the already-computed K steps. Thus α is the natural lever to adapt dynamically.

Adaptive rule in practice. In our implementation, α is scheduled online at each iteration. After generating rollouts with the current policy, we compute the policy entropy H_s and maintain a rolling buffer of the past ω entropy values (window size ω). Let μ_s and σ_s denote the mean and standard deviation within this buffer. We define the one-sided z-score

$$Z_s = \frac{H_s - \mu_s}{\sigma_s}. (11)$$

If $|Z_s| \ge \tau$ for a threshold τ , we mark the current step s^* and set $\alpha = 0$ for all $s \ge s^*$. Otherwise we keep $\alpha = \alpha_0$. Intuitively, sharp entropy fluctuations signal that the policy is close to a local optimum where noise dominates, so interpolation should be disabled. This entropy-triggered schedule exploits the fast trajectory in the high-signal early phase, while reverting to pure on-policy updates near convergence for stability.²

3.5 Unified SFPO Update

Collecting the three stages, the unified SFPO update for each step s is:

$$\theta^{s+1} = \theta^{s,0} - \eta \left[\underbrace{\alpha \sum_{k=0}^{K-1} \nabla_{\theta} \mathcal{L}(\theta^{s,k})}_{\text{fast trajectory & reposition}} + \underbrace{\nabla_{\theta} \mathcal{L}(\widetilde{\theta}^{s,K})}_{\text{slow correction}} \right], \tag{12}$$

As shown in Alg. 1, SFPO unifies the three stages into a single update rule and serves as a plug-compatible drop-in replacement for on-policy policy-gradient methods such as GRPO. Our experiments show that this structural change consistently improves stability and sample efficiency on diverse math reasoning benchmarks, with minimal engineering overhead.

²Although a decaying schedule on α could be a formal solution, we find little empirical difference, as shown in Fig. 7 of Sec. 5. Hence, for simplicity and efficiency, we downgrade SFPO to GRPO once the iteration reaches $s \ge s^*$.

Table 1: Performance on math reasoning benchmarks with DAPO and Math training dataset.

Model	Method	Math-500	AIME24	AIME25	AMC	Minerva	Olympiad	Avg
Qwen2.5-Math-1.5B	Base	55.55	9.17	5.83	37.65	17.74	28.45	25.73
	GRPO	77.15	16.67	11.67	53.31	31.89	39.42	38.35
	SFPO	78.35	20.00	15.00	56.02	32.07	39.72	40.19
Qwen2.5-Math-7B	Base	71.65	21.67	9.17	53.61	27.02	38.54	36.94
	GRPO	82.50	34.20	20.83	71.08	36.76	44.80	48.36
	SFPO	82.30	35.00	20.83	74.49	36.94	45.59	49.19
DS-distilled-Qwen-1.5B	Base	73.80	16.67	20.00	47.59	28.86	36.94	37.27
	GRPO	84.65	30.00	23.33	66.86	31.71	49.85	47.73
	SFPO	86.10	32.50	30.83	70.28	32.81	50.67	50.53
DS-distilled-Qwen-7B	Base	83.60	28.33	25.00	61.75	41.73	48.89	48.21
	GRPO	91.7	50.00	35.83	80.42	43.65	61.24	60.47
	SFPO	92.60	54.17	37.50	83.75	44.49	65.73	63.04
Qwen3-4B-Base	Base	45.25	2.50	0.83	20.48	15.99	20.66	17.62
	GRPO	83.35	16.67	17.50	59.03	38.78	48.59	43.99
	SFPO	84.30	21.67	20.83	57.23	40.81	48.67	45.59

4 EXPERIMENTS

4.1 EXPERIMENTAL SETTINGS

Models. We conduct the reasoning rl training with our proposed SFPO on a wide range of models: Qwen2.5-Math-1.5B (Yang et al., 2024), DeepSeek-R1-Distill-Qwen-1.5B (DeepSeek-AI et al., 2025a), Qwen3-4B-Base (Yang et al., 2025), Qwen2.5-Math-7B (Yang et al., 2024), and DeepSeek-R1-Distill-Qwen-7B (DeepSeek-AI et al., 2025a). For Qwen2.5-Math-1.5B/7B, we set the context length in the training process as 4096 the same as their maximum context length; while for Qwen3-4B-Base and DeepSeek-R1-Distill-Qwen-1.5B/7B, we set the context length as 8192 for the better accuracy-efficiency tradeoff.

Reasoning RL Training. We conduct the reasoning RL training on two different training datasets to test the effectiveness of SFPO on different scales of dataset. The first one is the combination of DAPO training dataset (Yu et al., 2025a) and Math training dataset (Hendrycks et al., 2021), which is a total of approximate 24K data. The second one is the Skywork-OR1 Math RL training dataset (He et al., 2025b) with 105k data. The training batch size is set to 256, and the number of responses for each question is set to 8 by default. The total trianing step is set to 400 by default. All the training experiments are done based on verl (Sheng et al., 2025) with a single 8xH100 GPUs node.

Baseline and Evaluation. We compare SFPO with vanilla GRPO and the base model without rl training on six commonly used mathematical reasoning benchmarks with variant difficulty: Math500 (Hendrycks et al., 2021), AIME24 (Art of Problem Solving, 2024a), AIME25 (MAA), AMC (Art of Problem Solving, 2024b), MinervaMath (Lewkowycz et al., 2022), Olympiad Bench (He et al., 2024). Each benchmark is evaluated multiple times with rollout temperature being 1, and we report the average Pass@1 accuracy by default.

4.2 MAIN RESULTS

4.2.1 MATH REASONING BENCHMARKS

As illustrated in Table 1, our proposed SFPO consistently outperforms vanilla GRPO across all base models and benchmarks. Specifically, for small-scale models such as Qwen2.5-Math-1.5B and DS-distilled-Qwen-1.5B, SFPO demonstrates superior performance enhancements on math reasoning benchmarks, raising the average accuracy from 38.35 to 40.19 with an absolute gain of +1.84, and from 47.73 to 50.53 with a gain of +2.80, respectively. The improvements are particularly pronounced on challenging tasks such as AIME24 and AIME25, where DS-distilled-Qwen-1.5B achieves an absolute gain of +7.5 on AIME25. The larger models also exhibit similar performance gains. For Qwen2.5-Math-7B, SFPO raises the average accuracy from 48.36 to 49.19 with an absolute gain of +1.80. For DS-distilled-Qwen-7B, SFPO boosts the average accuracy from 60.47 to 63.04, corresponding to an absolute gain of +0.8. For Qwen3-4B-Base model, SFPO improves

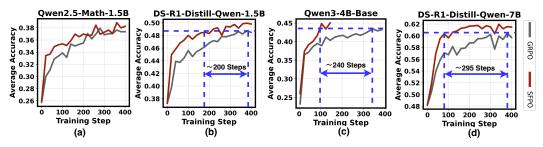


Figure 2: Average validation accuracy of different base models throughout the learning process.

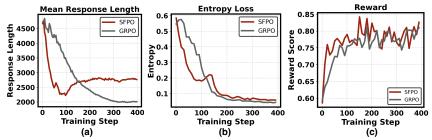


Figure 3: Comparison of training behaviors in terms of response length, entropy loss, and reward.

average accuracy from 43.99 to 45.59, an absolute gain of **+1.60**, highlighting its robustness across various models. Moreover, Table 2 demonstrates that SFPO can consistently achieve better training performance compared to GRPO for the larger training dataset Skywork-or1, proving the robustness of SFPO across different scales and distributions of training datasets.

4.2.2 Training Dynamics.

We compare the training dynamics between SFPO and GRPO to better understand the differences in optimization behavior as shown in Fig. 2 and Fig. 3.

Validation. From Fig. 2, we can clearly spot that SFPO consistently outperforms GRPO across all base models throughout the training process. Not only does SFPO achieve faster convergence in the early stages, but it also sustains higher global accuracy by the end of training. For example, Qwen3-4B-Base model achieves a sharper rise and stabilizes at

Table 2: Performance on AIME24/25 with Skywork-or1 training dataset.

Model	Method	AIME24	AIME25	
DS-Qwen-1.5B	Base	20.40	17.90	
	GRPO	32.92	25.83	
	SFPO	34.17	27.50	
DS-Qwen-7B	Base	25.42	22.92	
	GRPO	42.50	29.20	
	SFPO	43.75	30.00	

a higher accuracy within only 150 training steps, while vanilla GRPO cannot surpass this accuracy even after 400 steps.

Response Length. Moreover, distinct training behaviors between SFPO and GRPO for DeepSeek-R1-Distill-Qwen-7B are shown in Fig. 3 including the response length, entropy loss, and the reward throughout the training process. Specifically, GRPO gradually collapses to overly short responses while SFPO quickly converges to a stable range of around 2700 tokens with better accuracy, highlighting SFPO's ability to regulate response length more effectively and avoid overthinking with verbose responses (Liu et al., 2025).

Entropy. From Fig. 3(b), we can observe that SFPO makes model's entropy loss lower compared to GRPO. Typically, lower entropy means weak exploration ability for reasoning models; however, the entropy reduction under SFPO mainly reflects the model's ability to eliminate unproductive search paths early, rather than suppressing exploration altogether as evidenced by its sustained accuracy gains. In fact, the model still explores a sufficiently broad set of reasoning trajectories; therefore, the lower entropy observed under SFPO should be viewed as a sign of more efficient exploration rather than a sign of limited exploration.

Reward Score. SFPO also achieves a higher and more stable reward throughout the training process compared to GRPO, indicating stronger alignment with the reward function and more robust convergence. This is further reflected in its superior accuracy and well-controlled response length.

4.2.3 EFFICIENCY ANALYSIS.

We evaluate the efficiency gains of SFPO over GRPO by comparing the total number of rollouts and the wall-clock time required to reach the same benchmark accuracy. The results in Fig. 4 illustrate that SFPO consistently outperforms GRPO in both rollout efficiency and training time across all model scales. To be specific, SFPO requires 3.21×, 3.50×, and 4.93× fewer rollouts than GRPO for DS-Qwen-1.5B, Qwen3-4B-Base, and DS-Qwen-7B, respectively, to reach the same best accuracy. This advantage directly translates into reduced training time, where SFPO achieves 2.62×, 2.65×, and 4.19×

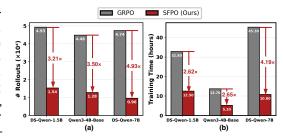


Figure 4: Comparison of GRPO and SFPO. (a) Number of rollouts required to achieve the best accuracy of GRPO. (b) Corresponding training time.

speedups over GRPO for the same models, significantly lowering the training cost. Note that SFPO does not introduce extra GPU memory overhead as it does not need to store the heavy optimizer status. The detailed profiling results for GPU memory usage in the training process can be found in Appendix B. These significant efficiency gains align with our expectations, since the primary bottleneck in the training process lies in rollout generation, which accounts for more than 70% of the overall inference time (He et al., 2025a). By substantially reducing the number of rollouts required and harnessing the reposition mechanism, SFPO alleviates this bottleneck and achieves faster training.

5 ANALYSIS AND ABLATION STUDY

Impact of α and K. As discussed in Sec. 3, α and K strongly affect the stability of SFPO. To better understand their impacts, we investigate two scenarios: (i) a small K=3 under varying α , and (ii) a large K=7 under varying α , with results shown in Fig. 5. With small K, SFPO stays stable across α values and consistently outperforms GRPO, consistent with our intuition. With large K, however, fast weights drift substantially; a large α further amplifies this mismatch, injecting noise into updates. A smaller α mitigates this drift and restores stability.

Interpolation Against Off-Policy Overfitting. When $\alpha=1$ (no interpolation), performance initially rises but steadily declines as training progresses. Without interpolation, the model quickly adapts by overfitting to small batches of rollouts, yielding short-term gains but injecting increasing noise into gradient updates over time, which leads to instability and degraded long-term performance. This issue is amplified when K is large, as shown in Fig. 5(b): fast weights drift substantially, and with $\alpha=1$ the original weights fully adopt these noisy trajectories, causing sharp performance degradation.

The Importance of Stage III: Slow Correction. As shown in Fig. 5(d), incorporating slow correction consistently improves stability and accuracy over GRPO. Without this stage, the reposition stage leaves the iterate at an interpolated point that may deviate from the true descent direction. Intuitively, slow correction provides a curvature-aware adjustment that realigns updates with the correct optimization trajectory.

Necessity of Entropy Control. We evaluate the effect of adaptive entropy control (EC) on SFPO using the DS-Qwen-7B model. As shown in Fig. 6, removing EC leads to an accuracy drop after roughly 100 steps. This degradation coincides with a rapid divergence of entropy loss, indicating that the policy becomes unstable and overfits to noisy rollouts. These results highlight entropy control as a crucial component for ensuring the stability and reliability of SFPO.

Strategies for Decaying α **.** We further evaluate different strategies for reducing the value of α once the stop step is identified by the z-score. In particular, we compare our default approach of directly setting α to zero with a linear decay schedule. As shown in Fig. 7, the choice of decay strategy has negligible impact on the overall performance.

6 RELATED WORKS

RL for LLM Reasoning. OpenAI O1 (OpenAI) introduced a paradigm shift in LLM reasoning by extending the reasoning horizon before final responses. DeepSeek-R1 (Guo et al., 2025) further

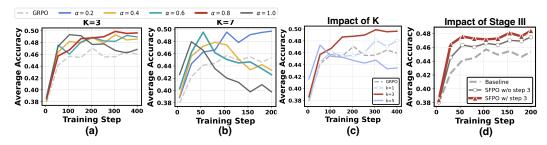
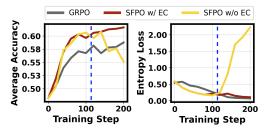


Figure 5: Average training accuracy of different base models throughout the learning process.



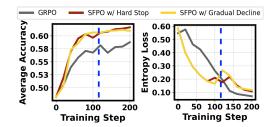


Figure 6: Comparison between SFPO w/ and w/o entropy control (EC). The blue dashed line indicates the stop step identified by z-score.

Figure 7: Comparison between SFPO with different α decay strategies. The blue dashed line indicates the stop step identified by z-score.

advanced this line by open-sourcing both its training algorithm, the value-model-free GRPO (Shao et al., 2024), and model weights, achieving performance comparable to O1. Subsequent work has focused on stabilizing and simplifying GRPO: DAPO (Yu et al., 2025b) identifies entropy collapse as a key challenge and proposes effective remedies, while Dr. GRPO (Liu et al., 2025) removes normalization terms without sacrificing performance. In contrast, SFPO is orthogonal to both families: it leaves the objective unchanged yet restructures the update itself into a fast–reposition–slow decomposition, improving variance reduction, stability, and sample efficiency in LLM reasoning.

Data Efficiency for LLM Reasoning. Despite focusing on designing novel training pipelines, a complementary line of work (Ivison et al., 2025; Xia et al., 2024; Muennighoff et al., 2025; Ye et al., 2025) improves the efficiency of LLM training through data filtering. One direction focuses on pruning data for supervised fine-tuning (Xia et al., 2024; Chen et al., 2023; Ivison et al., 2022). Another direction targets reinforcement learning, where studies (Muldrew et al., 2024; Liu et al., 2024; Das et al., 2024; Li et al., 2025; Fatemi et al., 2025; Wang et al., 2025) show that GRPO requires only a small subset of the training data to improve reasoning ability. However, these methods largely optimize which data is used rather than how updates are performed. SFPO is orthogonal: it assumes no change to the data pipeline but restructures the policy update itself. By converting one-shot updates into a fast–reposition–slow trajectory, SFPO reduces variance and stabilizes learning, thereby yielding higher data efficiency even under random rollout sampling.

7 Conclusion

We proposed **Slow-Fast Policy Optimization** (**SFPO**), a simple three-stage update mechanism that stabilizes early training by combining a fast trajectory, a reposition step, and a slow correction. SFPO directly addresses the instability of one-shot GRPO updates, achieving higher reasoning accuracy and fewer generated tokens without added wall-clock cost. Beyond these gains, its modular structure naturally opens avenues for curriculum or meta-learning extensions, and our entropy-triggered α schedule suggests richer adaptive rules. We believe SFPO offers not only an effective plug-in for current policy optimization, but also a foundation for scalable and data-aware optimization strategies in the LLM reasoning area.

REPRODUCIBILITY STATEMENT

All results reported in Sec. 4 are fully reproducible. We will release the code and experiment scripts upon acceptance.

REFERENCES

486

487 488

489

490 491

492 493

494

495

496

497

498 499

500

501

502

503

504 505

506

507

508

509

510

511

512

513

514

515

516

517

519

520

521

522

523

524

525

526

527

528

529

530

531

532

533

534 535

536

537

538

- Art of Problem Solving. Aime problems and solutions. https://artofproblemsolving.com/wiki/index.php/AIME, 2024a. Accessed: 2025-04-20.
- Art of Problem Solving. Amc problems and solutions. https://artofproblemsolving.com/wiki/index.php?title=AMC, 2024b. Problems and Solutions, 2024b. Accessed: 2025-04-20.
 - Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, et al. Alpagasus: Training a better alpaca with fewer data. *arXiv* preprint arXiv:2307.08701, 2023.
 - Muzhi Dai, Shixuan Liu, and Qingyi Si. Stable reinforcement learning for efficient reasoning. *arXiv* preprint arXiv:2505.18086, 2025. URL https://arxiv.org/abs/2505.18086.
 - Nirjhar Das, Souradip Chakraborty, Aldo Pacchiano, and Sayak Ray Chowdhury. Active preference optimization for sample efficient rlhf. *arXiv preprint arXiv:2402.10500*, 2024.
 - DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, 2025a. URL https://arxiv.org/abs/2501.12948.
 - DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song,

541

543

544

545

546

547

548

549

550

551

552

553

554

556

558

559

560 561

562

563

564

565

566 567

568

569

570

571

572

573574

575

576

577

578

579

580

581 582

583

584

585

586

588

589

590

592

Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. Deepseek-v3 technical report, 2025b. URL https://arxiv.org/abs/2412.19437.

- Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. Implementation matters in deep policy gradients: A case study on ppo and trpo. *arXiv preprint arXiv:2005.12729*, 2020.
- Mehdi Fatemi, Banafsheh Rafiee, Mingjie Tang, and Kartik Talamadupula. Concise reasoning via reinforcement learning. *arXiv preprint arXiv:2504.05185*, 2025.
- D. Guo et al. Deepseek-r1: Incentivizing reasoning capability in llms through rl. *Nature*, 2025. URL https://doi.org/10.1038/s41586-025-09422-z. To appear / in press.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, et al. OlympiadBench: A challenging benchmark for promoting AGI with olympiad-level bilingual multimodal scientific problems. *arXiv* preprint *arXiv*:2402.14008, 2024. URL https://arxiv.org/abs/2402.14008.
- Jingkai He, Tianjian Li, Erhu Feng, Dong Du, Qian Liu, Tao Liu, Yubin Xia, and Haibo Chen. History rhymes: Accelerating llm reinforcement learning with rhymerl, 2025a. URL https://arxiv.org/abs/2508.18588.
- Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, Siyuan Li, Liang Zeng, Tianwen Wei, Cheng Cheng, Bo An, Yang Liu, and Yahui Zhou. Skywork open reasoner 1 technical report, 2025b. URL https://arxiv.org/abs/2505.22312.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021. URL https://arxiv.org/abs/2103.03874.
- Hamish Ivison, Noah A Smith, Hannaneh Hajishirzi, and Pradeep Dasigi. Data-efficient finetuning using cross-task nearest neighbors. *arXiv preprint arXiv:2212.00196*, 2022.
- Hamish Ivison, Muru Zhang, Faeze Brahman, Pang Wei Koh, and Pradeep Dasigi. Large-scale data selection for instruction tuning. *arXiv preprint arXiv:2503.01807*, 2025.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and et al. Solving quantitative reasoning problems with language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, pp. 3843–3857, 2022.

595

596

597

598

600

601

602

603

604 605

606

607

608

609

610 611

612

613

614

615

616

617

618

619

620

621

622

623

625

626

627

630

631

632

633

634

635

636

637

638

640

641

642

644

645

646

Xuefeng Li, Haoyang Zou, and Pengfei Liu. Limr: Less is more for rl scaling. *arXiv preprint arXiv:2502.11886*, 2025.

Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*, 2025.

Zijun Liu, Boqun Kou, Peng Li, Ming Yan, Ji Zhang, Fei Huang, and Yang Liu. Enabling weak llms to judge response reliability via meta ranking. *arXiv preprint arXiv:2402.12146*, 2024.

MAA. American invitational mathematics examination (aime). https://maa.org/math-competitions/aime. Mathematics Competition Series; n.d.a.

Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time scaling. *arXiv preprint arXiv:2501.19393*, 2025.

William Muldrew, Peter Hayes, Mingtian Zhang, and David Barber. Active preference learning for large language models. *arXiv preprint arXiv:2402.08114*, 2024.

OpenAI. Learning to reason with llms: Openai o1 preview. URL https://openai.com/index/introducing-openai-o1-preview.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie

Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL https://arxiv.org/abs/2303.08774.

- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Si Shen, Peijun Shen, Wenhua Zhao, and Danhao Zhu. Mitigating think-answer mismatch in llm reasoning through noise-aware advantage reweighting. *arXiv preprint arXiv:2508.05928*, 2025. URL https://arxiv.org/abs/2508.05928.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, EuroSys '25, pp. 1279–1297. ACM, March 2025. doi: 10.1145/3689031.3696075. URL http://dx.doi.org/10.1145/3689031.3696075.
- Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, et al. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*, 2025.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. Less: Selecting influential data for targeted instruction tuning. *arXiv preprint arXiv:2402.04333*, 2024.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement, 2024. URL https://arxiv.org/abs/2409.12122.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025. URL https://arxiv.org/abs/2505.09388.
- Yixin Ye, Zhen Huang, Yang Xiao, Ethan Chern, Shijie Xia, and Pengfei Liu. Limo: Less is more for reasoning. *arXiv preprint arXiv:2502.03387*, 2025.
- Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Hongli Yu, Yuxuan Song, Xiangpeng Wei, Hao Zhou, Jingjing Liu, Wei-Ying Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yonghui Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale, 2025a. URL https://arxiv.org/abs/2503.14476.
- Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*, 2025b.

Michael Zhang, James Lucas, Jimmy Ba, and Geoffrey E Hinton. Lookahead optimizer: k steps forward, 1 step back. *Advances in neural information processing systems*, 32, 2019.

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

A LLM USAGE

Large language models (LLMs) were used solely to improve the writing of this paper, including grammar, clarity, and readability. They were not used for generating ideas, designing experiments, conducting analyses, or producing scientific content. All research contributions, technical claims, and conclusions are entirely the work of the authors.

B GPU MEMORY PROFILING RESULTS

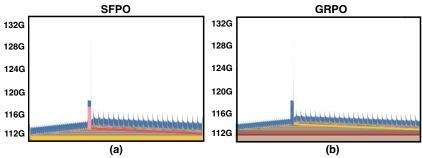


Figure 8: Comparison of GPU memory during one training step between SFPO and GRPO for Deepseek-R1-Distill-Qwen-1.5B model.

From Fig. 8, we can clearly observe that SFPO and GRPO demonstrate similar GPU memory consumption during one training step. This aligns with our expectation: SFPO does not need to store the heavy optimizer states and parameters but only need to store one copy of the model weight, which does not introduce significant overhead, especially when the model is sharded across GPUs.