

From $\log \pi$ to π : Taming Divergence in Soft Clipping via Bilateral Decoupled Decay of Probability Gradient Weight

Xiaoliang Fu^{1,2,*}, Jiaye Lin^{1,*†}, Yangyi Fang^{1,3,*}, Chaowen Hu¹,
Cong Qin^{1,4}, Zekai Shao², Binbin Zheng^{1,5}, Lu Pan¹, Ke Zeng¹

¹Meituan ²Fudan University ³Tsinghua University

⁴Peking University ⁵University of Science and Technology of China

{fuxiaoliang04, linjiaye}@meituan.com

Abstract

Reinforcement Learning with Verifiable Rewards (RLVR) has catalyzed a leap in Large Language Model (LLM) reasoning, yet its optimization dynamics remain fragile. Standard algorithms like GRPO enforce stability via “hard clipping”, which inadvertently stifles exploration by discarding gradients of tokens outside the trust region. While recent “soft clipping” methods attempt to recover these gradients, they suffer from a critical challenge: relying on *log-probability gradient* ($\nabla_{\theta} \log \pi_{\theta}$) yields divergent weights as probabilities vanish, destabilizing LLM training. We rethink this convention by establishing *probability gradient* ($\nabla_{\theta} \pi_{\theta}$) as the superior optimization primitive. Accordingly, we propose **Decoupled Gradient Policy Optimization (DGPO)**, which employs a decoupled decay mechanism based on importance sampling ratios. By applying asymmetric, continuous decay to boundary tokens, DGPO resolves the conflict between stability and sustained exploration. Extensive experiments across DeepSeek-R1-Distill-Qwen series models (1.5B/7B/14B) demonstrate that DGPO consistently outperforms strong baselines on various mathematical benchmarks, offering a robust and scalable solution for RLVR. Our code and implementation are available at: <https://github.com/FlyTune/DGPO-RL>.

1 Introduction

Reinforcement Learning (RL) has emerged as a transformative paradigm for aligning Large Language Models (LLMs) with human intent, shifting focus from imitation learning to goal-driven optimization (Ouyang et al., 2022; Shao et al., 2024; Guo et al., 2025). In reasoning-intensive domains like mathematics, Reinforcement Learning with Verifiable Rewards (RLVR) achieves remarkable success by leveraging ground-truth feedback to enhance accuracy and logical coherence (Lightman




Figure 1: Schematic overview of our DGPO algorithm. While “hard clipping” methods (GRPO/ASPO) discard gradients at boundaries, and prior “soft clipping” approaches (CISPO/GPPO/CE-GPPO) risk divergence on the left boundary and limit exploration on the right boundary, DGPO optimizes the gradient decay mechanism accordingly. By enforcing a controlled “Slow Down” behavior for stability and a “Slow Down Gently” behavior to sustain exploration, DGPO effectively resolves the exploration-stability conflict while maintaining minimal bias against the true policy gradient.

et al., 2023; Shao et al., 2024). Beyond standard benchmarks, stronger reasoning has also shown promise in optimization modeling, tool-augmented mathematical problem solving, scientific explanation, and narrative generation (Liu et al., 2026; Luo et al., 2025; Shao et al., 2025b,a).

Despite this promise, RLVR optimization remains fragile from the conflict between exploration and stability. Algorithms like PPO (Schulman et al., 2015) and GRPO (Shao et al., 2024) enforce a trust region by “hard clipping” the Importance Sampling (IS) ratio ($\pi_{\theta}/\pi_{\theta_{\text{old}}}$). While preventing destructive policy shifts, this inadvertently induces a *vanishing gradient for exploration*. Tokens drifting outside the trust region – representing valuable exploratory steps – receive zero updates, leading to entropy collapse and premature convergence (Williams and Peng, 1991; Eysenbach and Levine, 2021).

Recently, “soft clipping” approaches attempt to preserve gradients for out-of-bound tokens (Chen et al., 2025a; Su et al., 2025a,b). However, we

* Equal contribution. † Corresponding author.

identify a fundamental challenge: these methods predominantly operate on *log-probability gradient* ($\nabla_{\theta} \log \pi_{\theta}$). As probability approaches zero, the gradient weight of log-probability-based methods grows divergently. This causes catastrophic instability at the left boundary (where the IS ratio is extremely low), disproportionately penalizing tokens with vanishing probabilities.

Accordingly, we propose a paradigm shift: establishing *probability gradient* ($\nabla_{\theta} \pi_{\theta}$) as the primary optimization primitive over *log-probability gradient* ($\nabla_{\theta} \log \pi_{\theta}$). By analyzing the gradient landscape in probability space, we derive criteria for optimal weighting: (1) sustain exploration for clipped tokens, (2) ensure stability via convergent weights at boundaries, and (3) maximize alignment with the unbiased policy gradient.

In this paper, we propose **Decoupled Gradient Policy Optimization (DGPO)** guided by these principles. DGPO replaces “hard clipping” with a decoupled decay mechanism applied to the probability gradient weight. It applies a polynomial decay to tokens with *low IS ratios* (left boundary) for stability, and a reciprocal radical decay to tokens with *high IS ratios* (right boundary) to foster exploration. This mathematically guarantees gradient continuity and prevents the weight divergence seen in prior methods. Our contributions are summarized as:

- We introduce a novel perspective, establishing the gradient of probability, rather than log-probability, as the superior optimization primitive in LLMs, with two key insights: the inherent alignment of RL objectives and the geometric symmetry of probability space, which facilitates stable gradient design.
- We propose DGPO, which leverages a decoupled adaptive decay mechanism to reconcile the exploration-stability conflict. Crucially, this design preserves gradients for clipped tokens while rigorously preventing weight divergence.
- Comprehensive experiments against competitive baselines across mathematical reasoning benchmarks demonstrate the effectiveness of DGPO. Further results on diverse models scales confirm its scalability and robustness.

2 Related Works

2.1 RLVR in LLMs

RLVR utilizes deterministic signals (e.g., correct answers) rather than learned reward models to en-

hance LLM reasoning (Uesato et al., 2022). Recent advancements like DeepSeek-Math (Shao et al., 2024) popularized GRPO, which efficiently normalizes rewards within a sampled group, eliminating the need for a critic. However, GRPO inherits the clipping-induced exploration limitations of PPO.

2.2 Importance Sampling and Clipping

IS enables off-policy training by correcting distribution shifts via the ratio $\pi_{\theta}/\pi_{\theta_{\text{old}}}$ (Precup et al., 2000; Schulman et al., 2015). To prevent variance explosion, PPO (Schulman et al., 2017) clips this ratio within $[1 - \epsilon, 1 + \epsilon]$. This “hard clipping” mechanism zeros out gradients for outlier tokens, prioritizing exploitation while neglecting low-probability tokens essential for exploration, causing rapid entropy decay (O’Donoghue et al., 2016; Yu et al., 2025; Bai et al., 2026; Zhou et al., 2026). To mitigate exploration losses, several methods dynamically adjust clipping bounds (Yu et al., 2025; Yang et al., 2025a), allowing more updates for specific tokens. However, they still rely on hard boundaries, inevitably discarding gradient information for tokens beyond the adjusted thresholds.

2.3 Soft Clipping Policy Optimization

More recent approaches replace “hard clipping” with soft schemes. For instance, CISPO (Chen et al., 2025a) combines “soft clipping” with soft dual clip (Ye et al., 2020), while GPPO (Su et al., 2025a) retains a constant log-probability gradient weight for out-of-bound tokens. Crucially, both CISPO and GPPO suffer from left boundary instability: as $\pi_{\theta} \rightarrow 0$, the gradient grows indefinitely. Without proper decay mechanisms, this results in divergent updates that destabilize training. CE-GPPO (Su et al., 2025b) attempts to scale boundary gradients via hyperparameters but fails to resolve the underlying divergence. Distinctly, ASPO (Wang et al., 2025) proposes a reversed ratio to balance updates. Compared with these approaches, DGPO redefines the optimization target in probability space to ensure theoretical continuity and stability.

3 Preliminary

Problem Definition. In this paper, we focus on RLVR settings. Given a query q from dataset \mathcal{D} , a policy π_{θ} generates a response o . The rule-based reward function $r(q, o) \in \{-1, 1\}$ evaluates the correctness of each response. Following GRPO (Shao

et al., 2024), for each query q , a group of G outputs $\{o_i\}_{i=1}^G$ is sampled from the old policy $\pi_{\theta_{\text{old}}}$. The advantage \hat{A}_i is computed by normalizing the group-level rewards: $\hat{A}_i = (r(q, o_i) - \mu_R)/\sigma_R$, where μ_R and σ_R denote the mean and standard deviation of the rewards within the group.

Unified Gradient Formulation. To align off-policy training with reward maximization while constraining policy shifts for stability, most RLVR algorithms employ IS and clipping mechanisms. Thus, we formulate a *unified gradient estimator* that encompasses these methods. Let $w_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}$ denotes the token-level IS ratio, and the gradient is then expressed as:

$$\nabla_{\theta} \mathcal{J}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \mathcal{F}_{i,t}(\theta) \hat{A}_i \nabla_{\theta} \log \pi_{\theta}(o_{i,t}|q, o_{i,<t}). \quad (1)$$

The behavior of $\mathcal{F}_{i,t}(\theta)$ typically depends on five distinct regions, which are simultaneously defined by the clipping boundaries ($1 - \varepsilon_{\text{low}}$ and $1 + \varepsilon_{\text{high}}$) and the sign of the advantage \hat{A}_i . We formally define these regions in Table 1.

Table 1: Definition of clipping regions based on IS ratio $w_{i,t}(\theta)$ and advantage \hat{A}_i . Abbrev. denotes the abbreviation of five cases under different conditions.

Case	Abbrev.	Condition
Left Boundary (Low ratio, Neg. adv.)	LN	$w_{i,t}(\theta) < 1 - \varepsilon_{\text{low}}$ $\wedge \hat{A}_i < 0$
Right Boundary (High ratio, Pos. adv.)	HP	$w_{i,t}(\theta) > 1 + \varepsilon_{\text{high}}$ $\wedge \hat{A}_i > 0$
Reverse Left Boundary (Low ratio, Pos. adv.)	LP	$w_{i,t}(\theta) < 1 - \varepsilon_{\text{low}}$ $\wedge \hat{A}_i > 0$
Reverse Right Boundary (High ratio, Neg. adv.)	HN	$w_{i,t}(\theta) > 1 + \varepsilon_{\text{high}}$ $\wedge \hat{A}_i < 0$
In-Boundary (Medium ratio)	M	$\neg \text{LN} \wedge \neg \text{HP}$ $\wedge \neg \text{LP} \wedge \neg \text{HN}$

Instantiations of Clipping Strategies. Various existing methods can be interpreted as specific instantiations of $\mathcal{F}_{i,t}(\theta)$. Standard **PPO** (Schulman et al., 2017) and **GRPO** (Shao et al., 2024) apply “hard clipping” to penalize excessive updates:

$$\mathcal{F}_{i,t}^{\text{GRPO}}(\theta) = \begin{cases} 0, & \text{if LN} \vee \text{HP}, \\ w_{i,t}(\theta), & \text{otherwise.} \end{cases} \quad (2)$$

CISPO (Chen et al., 2025a) preserves exploration for these tokens by the following gradient weight:

$$\mathcal{F}_{i,t}^{\text{CISPO}}(\theta) = \begin{cases} 1 - \varepsilon_{\text{low}}, & \text{if LN} \vee \text{LP}, \\ 1 + \varepsilon_{\text{high}}, & \text{if HP} \vee \text{HN}, \\ w_{i,t}(\theta), & \text{otherwise.} \end{cases} \quad (3)$$

This method inherently employs soft dual clip (Ye et al., 2020), which clips the loss values for LP and HN (reverse cases) while preserving gradients. Subsequently, **GPPO** (Su et al., 2025a) introduces a gradient weight more aligned with PPO, focusing solely on LN and HP cases:

$$\mathcal{F}_{i,t}^{\text{GPPO}}(\theta) = \begin{cases} 1 - \varepsilon_{\text{low}}, & \text{if LN}, \\ 1 + \varepsilon_{\text{high}}, & \text{if HP}, \\ w_{i,t}(\theta), & \text{otherwise.} \end{cases} \quad (4)$$

CE-GPPO (Su et al., 2025b) further refines GPPO by introducing hyperparameters (β_1 and β_2) to control gradient scaling at both boundaries:

$$\mathcal{F}_{i,t}^{\text{CE}}(\theta) = \begin{cases} \beta_1(1 - \varepsilon_{\text{low}}), & \text{if LN}, \\ \beta_2(1 + \varepsilon_{\text{high}}), & \text{if HP}, \\ w_{i,t}(\theta), & \text{otherwise.} \end{cases} \quad (5)$$

To address the imbalance in updating positive-advantage tokens, **ASPO** (Wang et al., 2025) reverses gradient weight and adopts soft dual clip:

$$\mathcal{F}_{i,t}^{\text{ASPO}}(\theta) = \begin{cases} 0, & \text{if LN} \vee \text{HP}, \\ 1 - \varepsilon'_{\text{low}}, & \text{if LP}', \\ 1 + \varepsilon'_{\text{high}}, & \text{if HN}', \\ \frac{1}{w_{i,t}(\theta)}, & \text{otherwise.} \end{cases} \quad (6)$$

The values of $1 - \varepsilon'_{\text{low}}$ and $1 + \varepsilon'_{\text{high}}$ for cases LP' and HN' in ASPO may be more extreme.

4 Methodology

4.1 Shifting Focus: From $\log \pi$ to π

Classical policy gradient methods (Williams, 1992) typically employ the log-derivative trick to reformulate the gradient into an expectation, which establishes $\nabla_{\theta} \log \pi_{\theta}$ as the canonical term. This has led to a pervasive focus on log-probability gradient in subsequent research. However, we rethink that *probability gradient* constitutes a superior analytical and design target for two primary reasons: (1) probability, rather than log-probability, acts as the superior optimization primitive in LLM training, and (2) probability exhibits superior geometric symmetry within its value range.

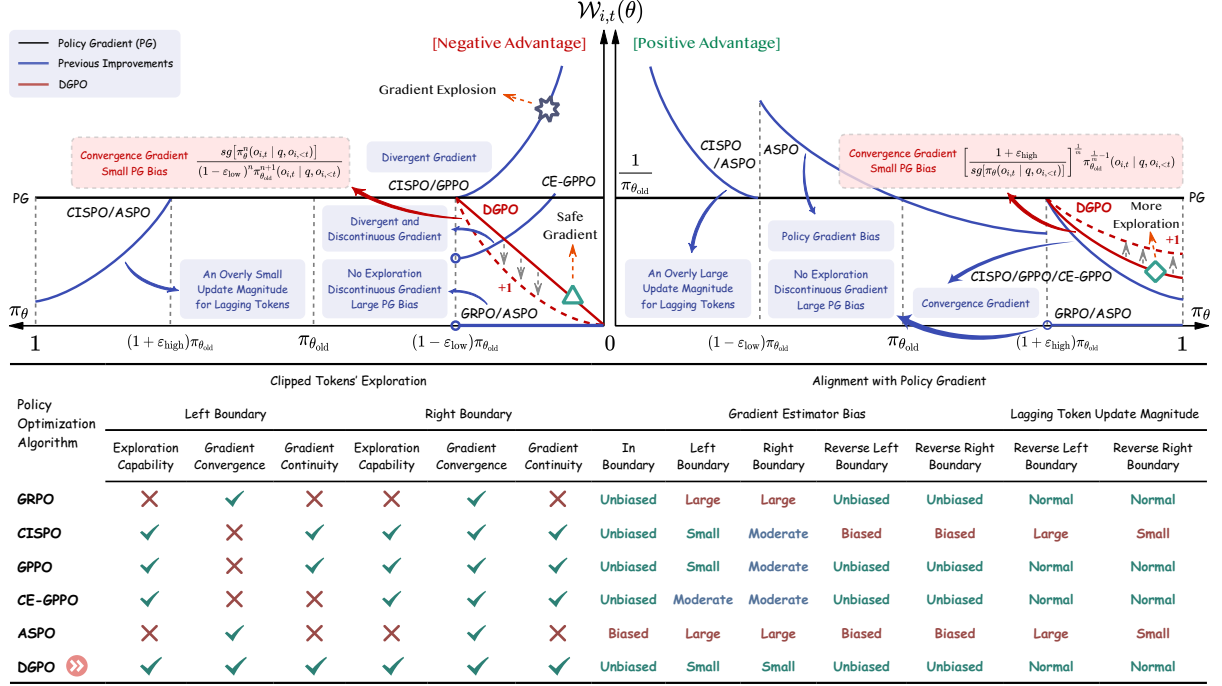


Figure 2: Comparative analysis of gradient dynamics. We systematically contrast DGPO with the standard GRPO, prior “soft clipping” enhancements (CISPO, GPPO, and CE-GPPO), and importance sampling improvements (ASPO). The visualization highlights the theoretical properties regarding the exploration capability of clipped tokens and the alignment with the true policy gradient, demonstrating DGPO’s superior stability and gradient consistency.

Probability as the Optimization Primitive. To substantiate this claim, we contrast the token-level objectives of Supervised Fine-Tuning (SFT) and RL. SFT maximizes the mean of expert token log-probabilities, with the gradient estimator as:

$$\nabla_{\theta} \mathcal{J}_{\text{SFT}}(\theta) = \mathbb{E}_{\mathcal{D}} \sum_{t=1}^{|\mathcal{O}|} \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}). \quad (7)$$

Conversely, RL (under binary advantage assumptions) equivalently maximizes the mean of expert token *probabilities* (derivation provided in Appendix A.1), yielding the estimator:

$$\nabla_{\theta} \mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{\mathcal{D}} \sum_{t=1}^{|\mathcal{O}|} \nabla_{\theta} \pi_{\theta}(o_t | q, o_{<t}). \quad (8)$$

Equations (7) and (8) reveal a fundamental distinction: SFT operates on log-probability, whereas RL inherently operates on probability. This distinction renders the SFT objective mathematically a lower bound of the RL objective (Qin and Springenberg, 2025), explaining the empirical performance ceiling often observed in SFT compared with RL (Wu et al., 2025). Consequently, we identify probability as the superior optimization primitive. Since the probability aligns more closely with the core requirements of LLM training, its gradient should be prioritized in algorithm design.

Geometric Symmetry and Boundedness. In LLM training, token probabilities reside in the symmetric and bounded interval $(0, 1)$. This boundedness facilitates the analysis of gradient impacts on probability values and enables the design of symmetric gradient mechanisms. In contrast, log-probabilities span the asymmetric and unbounded interval $(-\infty, 0)$, complicating the gradient design.

4.2 Decoupled Gradient Policy Optimization

Instability in Soft Clipping. While clipped tokens often contain critical information essential for model performance (Liu et al., 2025b; Su et al., 2025a), preserving their gradients requires a principled approach. Prior “soft clipping” works (Chen et al., 2025a; Su et al., 2025a) maintain constant log-probability gradient weights in LP and HN cases. However, the gradient weights are present asymmetrically at the boundaries, contradicting the symmetric nature of probability values.

Thus, we define probability gradient weight as $\mathcal{W}_{i,t}(\theta) = \frac{\nabla_{\theta} \mathcal{J}(\theta)}{\hat{A}_i \nabla_{\theta} \pi_{\theta}(o_{i,t} | q, o_{i,<t})}$. As illustrated in Figure 2, while the right-boundary gradient weight decreases convergently (promoting stability), the left-boundary weight grows divergently. This causes negative-advantage tokens to shrink disproportionately—the smaller the probability, the larger the up-

Table 2: Comparative analysis of gradient bias. We evaluate the bias magnitudes of DGPO and other baselines relative to the true policy gradient under different boundary conditions.

Bias Condition	Magnitude Comparison
In-Boundary Bias	$0 = \text{Bias}_{\text{DGPO}}^{\text{M}} = \text{Bias}_{\text{GRPO}}^{\text{M}} = \text{Bias}_{\text{CISPO}}^{\text{M}} = \text{Bias}_{\text{GPPO}}^{\text{M}} = \text{Bias}_{\text{CE}}^{\text{M}} \leq \text{Bias}_{\text{ASPO}}^{\text{M}}$
Left-Boundary Bias	$\begin{cases} 0 < \text{Bias}_{\text{DGPO}}^{\text{LN}} < \text{Bias}_{\text{CISPO}}^{\text{LN}} = \text{Bias}_{\text{GPPO}}^{\text{LN}} < \text{Bias}_{\text{CE}}^{\text{LN}} < \text{Bias}_{\text{GRPO}}^{\text{LN}} = \text{Bias}_{\text{ASPO}}^{\text{LN}}, & n = 1 \\ 0 < \text{Bias}_{\text{CISPO}}^{\text{LN}} = \text{Bias}_{\text{GPPO}}^{\text{LN}} < \text{Bias}_{\text{DGPO}}^{\text{LN}} < \text{Bias}_{\text{CE}}^{\text{LN}} < \text{Bias}_{\text{GRPO}}^{\text{LN}} = \text{Bias}_{\text{ASPO}}^{\text{LN}}, & n > 1 \end{cases}$
Right-Boundary Bias	$0 < \text{Bias}_{\text{DGPO}}^{\text{HP}} \leq \text{Bias}_{\text{CISPO}}^{\text{HP}} = \text{Bias}_{\text{GPPO}}^{\text{HP}} = \text{Bias}_{\text{CE}}^{\text{HP}} < \text{Bias}_{\text{GRPO}}^{\text{HP}} = \text{Bias}_{\text{ASPO}}^{\text{HP}}$
Reverse Left-Boundary Bias	$0 = \text{Bias}_{\text{DGPO}}^{\text{LP}} = \text{Bias}_{\text{GRPO}}^{\text{LP}} = \text{Bias}_{\text{GPPO}}^{\text{LP}} = \text{Bias}_{\text{CE}}^{\text{LP}} < \text{Bias}_{\text{CISPO}}^{\text{LP}} = \text{Bias}_{\text{ASPO}}^{\text{LP}}$
Reverse Right-Boundary Bias	$0 = \text{Bias}_{\text{DGPO}}^{\text{HN}} = \text{Bias}_{\text{GRPO}}^{\text{HN}} = \text{Bias}_{\text{GPPO}}^{\text{HN}} = \text{Bias}_{\text{CE}}^{\text{HN}} < \text{Bias}_{\text{CISPO}}^{\text{HN}} = \text{Bias}_{\text{ASPO}}^{\text{HN}}$

date magnitude—leading to training instability (Su et al., 2025b; Zhai et al., 2026; Zhang et al., 2026b). Although CE-GPPO attempts to mitigate this via hyperparameters, it does not resolve the divergent growth, leaving the risk of collapse.

DGPO Formulation. To address these limitations, we propose DGPO in this paper, designed to: (1) preserve gradient for clipped tokens, (2) stabilize exploration via adaptive gradient decay, and (3) minimize policy gradient bias for alignment.

We define the weighting function $\mathcal{W}_{i,t}^{\text{DGPO}}(\theta)$ based on the regions defined in Table 1:

$$\mathcal{W}_{i,t}^{\text{DGPO}}(\theta) = \begin{cases} C_{\text{left}} \cdot \text{sg}^n[\pi_{\theta}(o_{i,t}|q, o_{i,<t})], & \text{if LN,} \\ C_{\text{right}} \cdot \text{sg}^{-\frac{1}{m}}[\pi_{\theta}(o_{i,t}|q, o_{i,<t})], & \text{if HP,} \\ \frac{1}{\pi_{\theta_{\text{old}}}}, & \text{otherwise.} \end{cases} \quad (9)$$

Then, the objective function is formulated as:

$$\mathcal{J}_{\text{DGPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \mathcal{W}_{i,t}^{\text{DGPO}}(\theta) \hat{A}_i \pi_{\theta}(o_{i,t}|q, o_{i,<t}), \quad (10)$$

where $n, m \in \mathbb{Z}^+$ are hyperparameters controlling the decay rate, and $\text{sg}[\cdot]$ denotes the stop-gradient operator. $C_{\text{left}} = (1 - \varepsilon_{\text{low}})^{-n} \pi_{\theta_{\text{old}}}^{-(n+1)}(o_{i,t}|q, o_{i,<t})$ and $C_{\text{right}} = (1 + \varepsilon_{\text{high}})^{\frac{1}{m}} \pi_{\theta_{\text{old}}}^{\frac{1}{m}-1}(o_{i,t}|q, o_{i,<t})$ are constants ensuring continuity of gradient weights (derivation provided in Appendix A.2).

4.3 Theoretical Analysis and Advantages

Gradient Preserving and Exploration. Following prior gradient-preserving methods, DGPO adopts “soft clipping” to maintain gradient of clipped tokens. By retaining adaptively reduced gradients in both LN (Low Ratio, Negative Advantage) and HP (High Ratio, Positive Advantage) scenarios, DGPO enables sustained exploration for clipped tokens, enhancing the model’s exploration potential and raising its convergence upper bound.

Symmetric Stability Control. To ensure training stability while adhering to the principle of trust region optimization (limiting divergence between π_{θ} and $\pi_{\theta_{\text{old}}}$), gradient weights must decrease as the policy deviates. DGPO achieves this via a decoupled design: (1) **Left Boundary:** A positive integer power function of π_{θ} , ensuring weights decay as probability decreases. (2) **Right Boundary:** A reciprocal radical power function of π_{θ} , ensuring weights decay as probability increases.

Furthermore, DGPO manages entropy dynamics by controlling the “openness” of gradient weights. Since LN tokens drive entropy reduction and HP tokens drive entropy increase (Cui et al., 2025; Su et al., 2025b), the default symmetric setting ($m = 1$ and $n = 1$) may lead to rapid entropy decay. We balance this by adjusting n (reducing left-boundary openness) and m (increasing right-boundary openness). Additionally, multiplying gradient weights by C_{left} and C_{right} ensures continuity at the boundaries, ensuring a smooth transition between stable updates and decelerated exploration.

Minimal Bias Estimation. A critical challenge of existing RL algorithms is their deviation from the distribution correction principle of IS, resulting in significant bias relative to the true policy gradient. DGPO minimizes this bias to establish a stronger theoretical guarantee, which is the key factor for higher convergence bounds of model performance. Table 2 presents a comparative analysis of bias (proofs provided in Appendix A.3).

While CISPO and GPPO theoretically achieve minimal left-boundary bias when $n > 1$, they suffer from divergent gradient weights, constituting the root of training collapse. DGPO balances these objectives: it ensures gradient continuity and adaptive convergence while achieving minimal bias at $n = 1$ and maintaining a constant bias ratio relative to CISPO/GPPO for $n > 1$, thus offering a trade-off between stability and correctness.

Table 3: Comparison results of different methods on various benchmarks. Avg@32 (%) and Pass@32 (%) are abbreviated as A@32 and P@32. The best results are **bold**, and the second-best results are underlined, respectively.

Method	AIME24		AIME25		AMC23		MATH500		Minerva		Olympiad		Avg.	
	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32
DEEPSEEK-R1-DISTILL-QWEN-1.5B														
GRPO	33.2	71.8	27.7	49.9	79.5	94.8	<u>77.6</u>	90.8	26.1	48.8	<u>46.3</u>	64.7	48.4	70.1
CISPO	34.8	69.1	25.8	53.3	76.9	<u>94.9</u>	76.8	91.8	26.5	54.2	45.8	<u>65.8</u>	47.8	<u>71.5</u>
GPPO	29.6	60.5	23.5	51.9	73.5	94.1	76.3	89.1	26.6	50.0	43.9	64.2	45.6	68.3
CE-GPPO	35.1	70.2	27.7	<u>55.1</u>	82.5	95.0	76.7	90.2	<u>27.8</u>	<u>50.5</u>	45.6	63.1	<u>49.2</u>	70.7
ASPO	<u>36.4</u>	<u>73.2</u>	<u>28.3</u>	51.5	<u>83.1</u>	94.7	74.6	90.5	26.0	49.8	44.9	63.7	48.9	70.6
Ours	43.3	79.3	32.8	56.1	86.0	95.0	77.9	<u>91.0</u>	28.2	50.4	48.0	66.4	52.7	73.0
DEEPSEEK-R1-DISTILL-QWEN-7B														
GRPO	48.2	82.5	37.4	60.5	88.1	<u>96.6</u>	<u>84.8</u>	92.4	37.4	57.2	<u>57.2</u>	<u>73.9</u>	58.9	<u>77.2</u>
CISPO	51.6	76.6	<u>38.2</u>	<u>65.4</u>	90.6	<u>96.6</u>	82.1	91.6	38.7	56.5	54.3	69.9	<u>59.3</u>	76.1
GPPO	43.1	72.5	31.7	62.5	85.6	94.9	83.2	95.4	33.1	59.3	53.2	74.3	55.0	76.5
CE-GPPO	48.7	76.9	36.4	60.4	<u>90.5</u>	95.0	84.3	93.3	<u>39.0</u>	55.4	54.9	72.5	59.0	75.6
ASPO	<u>51.8</u>	79.6	37.1	54.1	90.0	97.2	83.8	<u>94.9</u>	37.0	<u>59.2</u>	54.1	72.8	59.0	76.3
Ours	55.5	<u>81.9</u>	43.1	68.0	90.6	<u>96.6</u>	85.4	92.0	39.8	56.7	57.7	72.0	62.0	77.9

5 Experiments

5.1 Experimental Setup

Configurations. For all experiments, we employ DeepSeek-R1-Distill-Qwen (Guo et al., 2025) with various scales as backbone models to validate the effectiveness of DGPO. The open-source DAPO-Math-17K dataset (Yu et al., 2025) is utilized for model training. During training, rule-based rewards are computed with `math_verify` (Kydliček, 2024), while a conservative method `prime_math` is used for evaluation (Lightman et al., 2023).

Moreover, to ensure consistent convergence dynamics across different model scales, we calibrate the learning rate to maintain a *constant total gradient variance* (derivation provided in Appendix A.4). Accordingly, the learning rates are set as follows: 1.0×10^{-6} for 1.5B, 4.63×10^{-7} for 7B, and 3.27×10^{-7} for 14B. We set the batch size to 512 and the mini-batch size to 32. This configuration enables 16 off-policy updates per IS step, sufficiently amplifying the influence of boundary tokens. ϵ_{low} and ϵ_{high} are fixed to 0.2. Detailed configurations are available in Appendix B.

Benchmarks and Metrics. We assess the generalization of reasoning capabilities on widely used mathematical benchmarks, i.e., AIME24 (MAA, 2024), AIME25 (MAA, 2025), AMC23 (MAA, 2023), MATH500 (Hendrycks et al., 2021), Min-

erva (Lewkowycz et al., 2022), and Olympiad-Bench (He et al., 2024). More details about benchmarks are described in Appendix B. We report Avg@32 to measure the expected performance and Pass@32 to gauge the potential capability.

Baselines. We benchmark DGPO against standard GRPO and variants employing distinct gradient weighting strategies, including “soft clipping” (CISPO and GPPO), scaled “soft clipping” (CE-GPPO, adopting the recommended $\beta_1 = 0.75$ and $\beta_2 = 1$), and reverse gradient weight (ASPO).

5.2 Main Results

Overall Performance. Table 3 presents the performance comparison across 1.5B and 7B scales. DGPO demonstrates superior performance across the majority of benchmarks on both scales. On the 1.5B model, DGPO surpasses the vanilla GRPO by +4.3% and the best baseline (CE-GPPO) by +3.5% in average Avg@32. On the 7B model, the improvement is also significant, with DGPO outperforming GRPO by +3.1% and CISPO by +2.7%. As shown in Figure 3(a,g), DGPO consistently maintains higher Pass@K scores on AIME2025 compared with all baselines, demonstrating superior potential capability. Detailed results of Pass@K ($K = 1, \dots, 32$) on AIME24 and AIME25 across all model scales are provided in Appendix C.3.

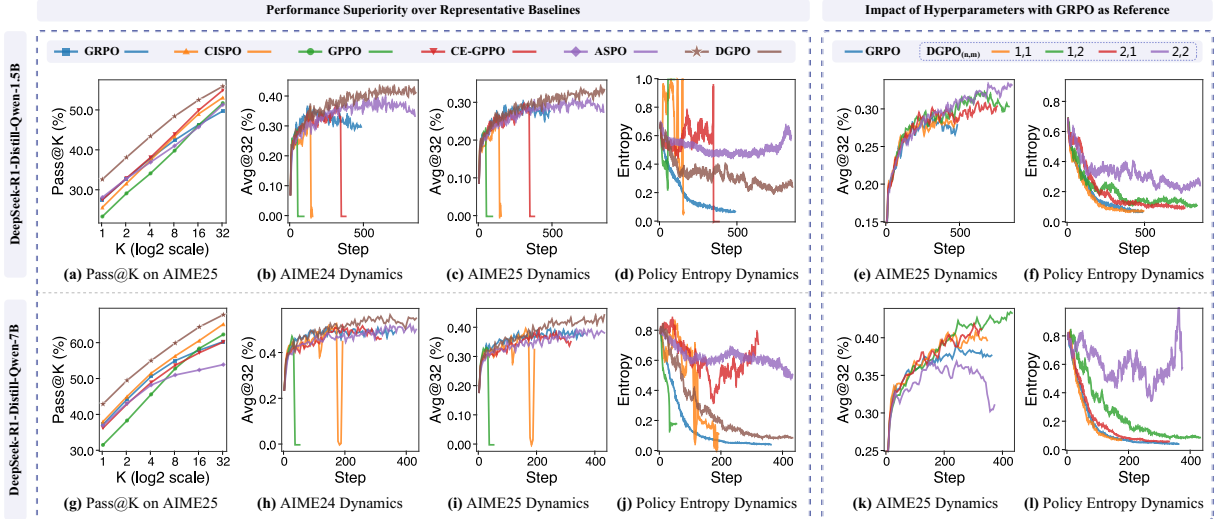


Figure 3: Comprehensive comparison of training dynamics and performance. **Top row:** DeepSeek-R1-Distill-Qwen-1.5B results. **Bottom row:** DeepSeek-R1-Distill-Qwen-7B results. **Columns (L-R):** Pass@K on AIME25, Avg@32 on AIME24/25, policy entropy, followed by hyperparameter analysis for Avg@32 and entropy.

Training Dynamics. Figure 3(b,c,h,i) illustrates the training dynamics of Avg@32 on AIME2024 and AIME2025. DGPO consistently outperforms all baselines in the mid-to-late training stages. Notably, algorithms with divergent gradient weights at the left boundary (CISPO, GPPO, and CE-GPPO) succumb to training collapse, and algorithms with larger policy gradient bias (GRPO and ASPO) exhibit suboptimal convergence.

Figure 3(d,j) depicts the entropy dynamics. GRPO exhibits an entropy drop early in the training process, indicating premature exploitation that limits the achievable performance ceiling. Conversely, ASPO maintains excessively high entropy, reflecting over-exploration and insufficient exploitation, while CISPO, GPPO, and CE-GPPO show unstable entropy patterns leading to eventual collapse. DGPO demonstrates a moderate and controlled entropy reduction, signifying an optimal balance between exploration and exploitation.

5.3 Hyperparameter Analysis

We investigate the impact of hyperparameters n and m by expanding from the baseline ($n = 1$, $m = 1$) to four configurations: (1, 1), (1, 2), (2, 1), and (2, 2). Figure 3(e,k) shows their performance on AIME25 across 1.5B and 7B models.

Robustness and Patterns. All DGPO configurations outperform GRPO (with the exception of $n = 2$ and $m = 2$ on the 7B model), demonstrating our algorithm’s robustness. However, the optimal configuration varies by scale: (2, 2) for 1.5B and

(1, 2) for 7B. Analyzing Figure 3(e,f,k,l), which reveals a consistent pattern: increasing n or m generally yields: (1) improved performance, (2) elevated overall entropy levels, but (3) reduced entropy stability. Notably, significant entropy volatility is observed only in the 7B model under the (2, 2) setting that negates the performance benefits.

Tuning Guideline. Based on these observations, we propose a heuristic for hyperparameter tuning: *Enhance exploration by increasing n and m as long as entropy remains stable. Upon observing instability, revert to the preceding stable configuration.* Empirically, we recommend $n = 1$ and $m = 2$ as a robust and conservative baseline configuration.

Table 4: Scalability analysis of average Avg@32 (A) and Pass@32 (P) across 1.5B, 7B, and 14B models.

Method	1.5B	7B	14B ¹
	(A / P)	(A / P)	(A / P)
GRPO	48.4 / 70.1	58.9 / 77.2	53.6 / 67.4
DGPO	52.7 / 73.0	62.0 / 77.9	56.7 / 70.4
<i>Improvement</i>	<i>+4.3 / +2.9</i>	<i>+3.1 / +0.7</i>	<i>+3.1 / +3.0</i>

Scaling to Larger Models. The instability of (2, 2) on 7B (absent in 1.5B) suggests that larger models exhibit higher intrinsic entropy volatility, necessitating more conservative hyperparameters.

¹The performance fluctuation is attributed to the base model difference: 1.5B/7B models are derived from Qwen-2.5-Math, whereas the 14B model is derived from Qwen-2.5-Base.

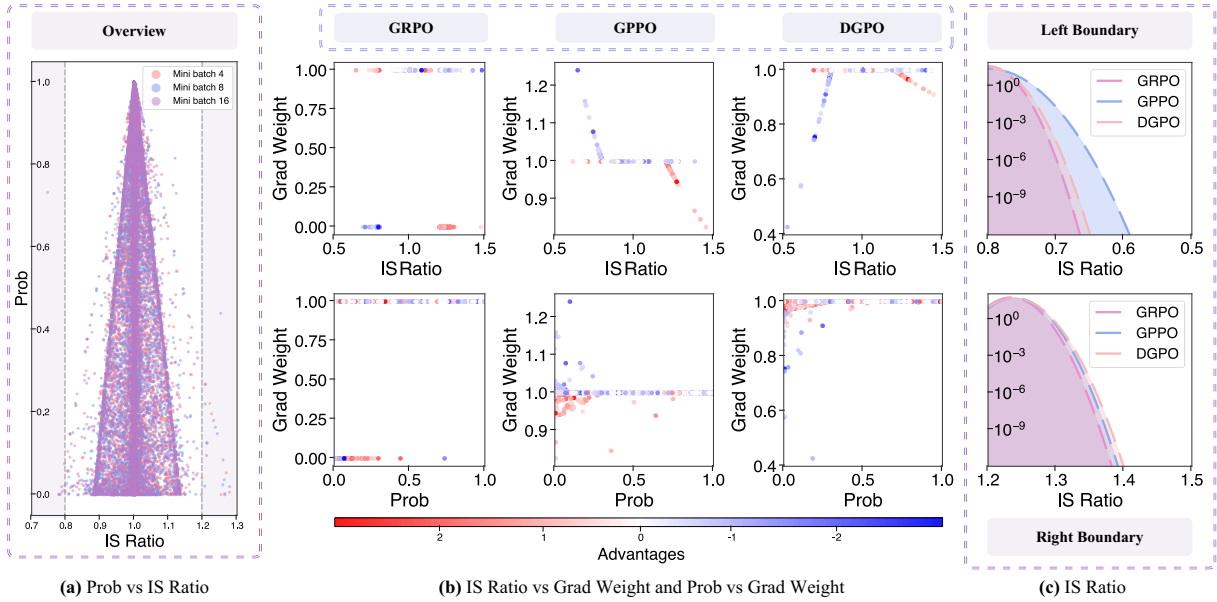


Figure 4: Comparison of gradient weight distributions (Prob: Probability, Grad: Gradient). (a) Overall distribution. (b) Detailed scatter plots of grad weight vs. IS ratio (top row) and prob (bottom row) across three methods: GRPO (left), GPPO (middle), and DGPO (right). Points are colored by advantage. (c) Boundary distribution analysis.

To validate this pattern, we apply the optimal 7B configuration ($n = 1, m = 2$) to the 14B model.

As summarized in Table 4, DGPO consistently outperforms GRPO across all model scales in both Avg@32 and Pass@32, which confirms that the benefits of our decoupled gradient mechanism effectively transfer to larger models. Furthermore, the training dynamics and entropy variations of the 14B model (shown in Appendix C.1) exhibit stable convergence patterns similar to 7B under DGPO.

5.4 Mechanistic Analysis via Visualization

Mechanism of Stability (Left Boundary). *Why does DGPO prevent collapse?* Figure 4(a) visualizes the joint distribution of probability and IS ratios in the final mini-batch. Crucially, tokens at both boundaries are predominantly low-probability tokens. Figure 4(b) further details the relationship between probability, IS ratio, advantage, and relative gradient weight (normalized by $\pi_{\theta_{\text{old}}}$ probability) for GRPO (zero weight), GPPO (divergent weight), and DGPO (convergent weight). As illustrated, GRPO’s zero gradients at the left boundary lead to the narrowest ratio distribution and insufficient exploration. Conversely, GPPO’s divergent weights induce an excessively broad distribution, eventually precipitating training collapse due to instability (Figure 3(b,c,h,i)) (Yang et al., 2025b). In contrast, DGPO maintains convergent relative weights, resulting in a ratio distribution

only slightly wider than GRPO that successfully balances stability with effective exploration.

Mechanism of Improvement (Right Boundary).

Why does DGPO perform better? Compared with GRPO (where right-boundary tokens have zero gradient and the narrowest ratio distribution), DGPO maintains gradients to foster exploration. Compared with GPPO (which uses a reciprocal standard weight equivalent to DGPO’s $m = 1$), DGPO with $m = 2$ employs a reciprocal radical weight. This design induces the widest ratio distribution on the right boundary (Figure 4(c)), significantly enhancing performance. This finding aligns with prior research suggesting that increasing the contribution of positive samples improves performance (Yang et al., 2025b; Xi et al., 2025), and is consistent with our hyperparameter analysis in Section 5.3 where increasing m boosts results.

6 Conclusion

In this paper, we revisit the fundamental optimization primitive in RLVR and establish probability—rather than log-probability—as the essential alignment target. We argue that the prevalent focus on log-probability gradients in prior “soft clipping” methods constitutes a misalignment with the true RL objective, leading to divergent gradient weights at boundaries and consequent training instability. To address this critical issue, we propose DGPO, which directly optimizes probability gra-

dients via a decoupled decay mechanism based on IS ratios. This approach effectively resolves the inherent conflict between training stability and gradient preservation. Extensive experiments across the DeepSeek-R1-Distill-Qwen series models with multiple scales (1.5B, 7B, and 14B) demonstrate that DGPO consistently outperforms competitive baselines on various mathematical reasoning benchmarks, validating that aligning with the probability objective is crucial for unlocking the full potential of LLM reasoning capabilities through RL.

7 Limitations

We discuss two practical limitations of our work. (1) **Domain Specificity:** In line with most RLVR studies, our experiments are primarily conducted on mathematical reasoning datasets (e.g., AIME, MATH500) where verifiable rewards are readily available. While we believe the principles of probability gradients are generalizable, the efficacy of DGPO in domains with sparse or subjective rewards (e.g., creative writing) remains to be verified. We anticipate that the principles underlying DGPO generalize beyond the current setting: applying them to adjacent challenges—such as training data quality improvement (Liu et al., 2023; Liu, 2024; Liu et al., 2025a), parameter-efficient fine-tuning (Dong et al., 2025), structured reasoning (Dong et al., 2026; Jiang et al., 2026; Zhang et al., 2026c), agentic exploration (Zhang et al., 2026a), and non-textual domains (Li et al., 2026; Chen et al., 2025b, 2026)—constitutes a promising direction for future work. (2) **Computational Constraints:** Due to limited computational resources, our scaling laws analysis is restricted to models up to 14B parameters. Although we observed consistent performance gains and predictable hyperparameter patterns from 1.5B to 14B, validating these findings on significantly larger foundation models (e.g., 70B+) would provide further insights into the method’s scalability.

8 Ethical Considerations

We have carefully considered the ethical implications of our research and provide the following statements: (1) **Compliance and Transparency:** Throughout this study, we have strictly followed established ethical guidelines. Our findings are reported honestly, and we provide comprehensive theoretical derivations to ensure transparency. (2) **Data Usage:** The datasets employed in our ex-

periments (e.g., DAPO-Math-17k) originate from publicly available sources. No private, sensitive, or confidential user information was used at any stage of our research. (3) **Reproducibility:** We offer detailed descriptions of the hyperparameter configurations, including the specific tuning guidelines for n and m , to ensure the reproducibility of our results. (4) **Open Source:** In the interest of openness and to facilitate future research in the RLVR community, we have made our code available anonymously and will fully open-source it upon the acceptance of this paper.

References

- Sikai Bai, Haoxi Li, Jie Zhang, Yongjiang Liu, and Song Guo. 2026. Ttvs: Boosting self-exploring reinforcement learning via test-time variational synthesis. *arXiv preprint arXiv:2604.08468*.
- Aili Chen, Aonian Li, Bangwei Gong, Binyang Jiang, Bo Fei, Bo Yang, Boji Shan, Changqing Yu, Chao Wang, Cheng Zhu, and 1 others. 2025a. Minimax-m1: Scaling test-time compute efficiently with lightning attention. *arXiv preprint arXiv:2506.13585*.
- Zhiwei Chen, Yupeng Hu, Zhiheng Fu, Zixu Li, Jiale Huang, Qinlei Huang, and Yinwei Wei. 2026. Intent: Invariance and discrimination-aware noise mitigation for robust composed image retrieval. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 20463–20471.
- Zhiwei Chen, Yupeng Hu, Zixu Li, Zhiheng Fu, Xueming Song, and Liqiang Nie. 2025b. Offset: Segmentation-based focus shift revision for composed image retrieval. In *ACM International Conference on Multimedia (ACM MM)*, pages 6113–6122.
- Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen Fan, Huayu Chen, Weize Chen, and 1 others. 2025. The entropy mechanism of reinforcement learning for reasoning language models. *arXiv preprint arXiv:2505.22617*.
- Haonan Dong, Kehan Jiang, Haoran Ye, Wenhao Zhu, Zhaolu Kang, and Guojie Song. 2026. **Neureasoner: Towards explainable, controllable, and unified reasoning via mixture-of-neurons.** *Preprint*, arXiv:2604.02972.
- Haonan Dong, Wenhao Zhu, Guojie Song, and Liang Wang. 2025. **AuroRA: Breaking low-rank bottleneck of loRA with nonlinear mapping.** In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Benjamin Eysenbach and Sergey Levine. 2021. Maximum entropy rl (provably) solves some robust rl problems. *arXiv preprint arXiv:2103.06257*.

- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint arXiv:2402.14008*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Kehan Jiang, Haonan Dong, Zhaolu Kang, Zhengzhou Zhu, and Guojie Song. 2026. Foe: Forest of errors makes the first solution the best in large reasoning models. *Preprint*, arXiv:2604.02967.
- Hynek Kydlíček. 2024. [Math-Verify: Math Verification Library](#).
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and 1 others. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857.
- Zixu Li, Yupeng Hu, Zhiwei Chen, Shiqi Zhang, Qinlei Huang, Zhiheng Fu, and Yinwei Wei. 2026. Habit: Chrono-synergia robust progressive learning framework for composed image retrieval. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 6762–6770.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Peiyang Liu. 2024. Unsupervised corrupt data detection for text training. *Expert Systems with Applications*, 248:123335.
- Peiyang Liu, Ziqiang Cui, Di Liang, and Wei Ye. 2025a. Who stole your data? a method for detecting unauthorized rag theft. *arXiv preprint arXiv:2510.07728*.
- Peiyang Liu, Jinyu Yang, Lin Wang, Sen Wang, Yunlai Hao, and Huihui Bai. 2023. Retrieval-based unsupervised noisy label detection on text data. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 4099–4104.
- Weiting Liu, Han Wu, Yufei Kuang, Xiongwei Han, Tao Zhong, Jianfeng Feng, and Wenlian Lu. 2026. Automated optimization modeling via a localizable error-driven perspective. *arXiv preprint arXiv:2602.11164*.
- Zheng Liu, Mengjie Liu, Siwei Wen, Mengzhang Cai, Bin Cui, Conghui He, and Wentao Zhang. 2025b. From uniform to heterogeneous: Tailoring policy optimization to every token’s nature. *arXiv preprint arXiv:2509.16591*.
- Haipeng Luo, Huawen Feng, Qingfeng Sun, Can Xu, Kai Zheng, Yufei Wang, Tao Yang, Han Hu, Yansong Tang, and Di Wang. 2025. Agentmath: Empowering mathematical reasoning for large language models via tool-augmented agent. *arXiv preprint arXiv:2512.20745*.
- MAA. 2023. [American mathematics competitions - amc](#). Accessed: 2023.
- MAA. 2024. [American invitational mathematics examination - aime](#). Accessed: 2024.
- MAA. 2025. [American invitational mathematics examination - aime](#). Accessed: 2025.
- Brendan O’Donoghue, Remi Munos, Koray Kavukcuoglu, and Volodymyr Mnih. 2016. Combining policy gradient and q-learning. *arXiv preprint arXiv:1611.01626*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Doina Precup, Richard S Sutton, and Satinder P Singh. 2000. Eligibility traces for off-policy policy evaluation. In *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 759–766.
- Chongli Qin and Jost Tobias Springenberg. 2025. Supervised fine tuning on curated data is reinforcement learning (and can be improved). *arXiv preprint arXiv:2507.12856*.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zekai Shao, Leixian Shen, Haotian Li, Yi Shan, Huamin Qu, Yun Wang, and Siming Chen. 2025a. Narrative player: Reviving data narratives with visuals. *IEEE Transactions on Visualization and Computer Graphics*, 31(10):6781–6795.

- Zekai Shao, Siyu Yuan, Lin Gao, Yixuan He, Deqing Yang, and Siming Chen. 2025b. Unlocking scientific concepts: How effective are llm-generated analogies for student understanding and classroom practice? In *Proceedings of the 2025 CHI conference on human factors in computing systems*, pages 1–19.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, pages 1279–1297.
- Zhenpeng Su, Leiyu Pan, Xue Bai, Dening Liu, Guanting Dong, Jiaming Huang, Wenping Hu, Fuzheng Zhang, Kun Gai, and Guorui Zhou. 2025a. Klear-reasoner: Advancing reasoning capability via gradient-preserving clipping policy optimization. *arXiv preprint arXiv:2508.07629*.
- Zhenpeng Su, Leiyu Pan, Minxuan Lv, Yuntao Li, Wenping Hu, Fuzheng Zhang, Kun Gai, and Guorui Zhou. 2025b. Ce-gppo: Coordinating entropy via gradient-preserving clipping policy optimization in reinforcement learning. *arXiv preprint arXiv:2509.20712*.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. Solving math word problems with process-and outcome-based feedback. *arXiv preprint arXiv:2211.14275*.
- Jiakang Wang, Runze Liu, Lei Lin, Wenping Hu, Xiu Li, Fuzheng Zhang, Guorui Zhou, and Kun Gai. 2025. Aspo: Asymmetric importance sampling policy optimization. *arXiv preprint arXiv:2510.06062*.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Ronald J Williams and Jing Peng. 1991. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268.
- Yongliang Wu, Yizhou Zhou, Zhou Ziheng, Yingzhe Peng, Xinyu Ye, Xinting Hu, Wenbo Zhu, Lu Qi, Ming-Hsuan Yang, and Xu Yang. 2025. On the generalization of sft: A reinforcement learning perspective with reward rectification. *arXiv preprint arXiv:2508.05629*.
- Zhiheng Xi, Xin Guo, Yang Nan, Enyu Zhou, Junrui Shen, Wenxiang Chen, Jiaqi Liu, Jixuan Huang, Zhihao Zhang, Honglin Guo, and 1 others. 2025. Bapo: Stabilizing off-policy reinforcement learning for llms via balanced policy optimization with adaptive clipping. *arXiv preprint arXiv:2510.18927*.
- Shihui Yang, Chengfeng Dou, Peidong Guo, Kai Lu, Qiang Ju, Fei Deng, and Rihui Xin. 2025a. Dcpo: Dynamic clipping policy optimization. *arXiv preprint arXiv:2509.02333*.
- Zhihe Yang, Xufang Luo, Zilong Wang, Dongqi Han, Zhiyuan He, Dongsheng Li, and Yunjian Xu. 2025b. Do not let low-probability tokens over-dominate in rl for llms. *arXiv preprint arXiv:2505.12929*.
- Deheng Ye, Zhao Liu, Mingfei Sun, Bei Shi, Peilin Zhao, Hao Wu, Hongsheng Yu, Shaojie Yang, Xipeng Wu, Qingwei Guo, and 1 others. 2020. Mastering complex control in moba games with deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, pages 6672–6679.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, and 1 others. 2025. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Naixin Zhai, Pengyang Shao, Binbin Zheng, Yonghui Yang, Fei Shen, Long Bai, and Xun Yang. 2026. Maximizing local entropy where it matters: Prefix-aware localized llm unlearning. *arXiv preprint arXiv:2601.03190*.
- Wenyuan Zhang, Xinghua Zhang, Haiyang Yu, Shuaiyi Nie, Bingli Wu, Juwei Yue, Tingwen Liu, and Yongbin Li. 2026a. [Expseek: Self-triggered experience seeking for web agents](#). *Preprint*, arXiv:2601.08605.
- Xinglang Zhang, Yunyao Zhang, ZeLiang Chen, Junqing Yu, Wei Yang, and Zikai Song. 2026b. [Logical phase transitions: Understanding collapse in llm logical reasoning](#). *Preprint*, arXiv:2601.02902.
- Yunyao Zhang, Xinglang Zhang, Junxi Sheng, Wenbing Li, Junqing Yu, Yi-Ping Phoebe Chen, Wei Yang, and Zikai Song. 2026c. [Semantic-aware logical reasoning via a semiotic framework](#). *Preprint*, arXiv:2509.24765.
- Yixiao Zhou, Yang Li, Dongzhou Cheng, Hehe Fan, and Yu Cheng. 2026. Look inward to explore outward: Learning temperature policy from llm internal states via hierarchical rl. *arXiv preprint arXiv:2602.13035*.

A Derivation and Proof

A.1 Derivation of RL Estimator in SFT Form

The token-level objective of RL is to maximize the arithmetic mean expectation of the estimated advantage, formulated as follows:

$$\mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta}(\cdot|q)} \left[\sum_{t=1}^{|o|} A_t \right] \quad (11)$$

Assume that $\pi_e(\cdot|q)$ represents the ideal expert distribution, which satisfies two key properties: (1)

its output follows a one-hot distribution; (2) every sampled token o_t achieves the maximum advantage A_t . By applying Importance Sampling, we obtain:

$$\mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_e(\cdot|q)} \left[\sum_{t=1}^{|o|} \frac{\pi_\theta(o_t | q, o_{<t})}{\pi_e(o_t | q, o_{<t})} \cdot A_t \right] \quad (12)$$

Since $\pi_e(\cdot|q)$ is a one-hot distribution, the probability $\pi_e(o_t | q, o_{<t})$ is always 100%, leading to:

$$\mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_e(\cdot|q)} \left[\sum_{t=1}^{|o|} \pi_\theta(o_t | q, o_{<t}) \cdot A_t \right] \quad (13)$$

Given that each action o_t attains the highest possible A_t , under the binary advantage setting, A_t is consistently 1, resulting in:

$$\mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{(q,o) \sim \mathcal{D}_e} \left[\sum_{t=1}^{|o|} \pi_\theta(o_t | q, o_{<t}) \right] \quad (14)$$

Consequently, the gradient is derived as:

$$\nabla_\theta \mathcal{J}_{\text{RL}}(\theta) = \mathbb{E}_{(q,o) \sim \mathcal{D}_e} \left[\sum_{t=1}^{|o|} \nabla_\theta \pi_\theta(o_t | q, o_{<t}) \right] \quad (15)$$

A.2 Derivation of Continuity Constants

In this section, we derive the constants C_{left} and C_{right} used in the DGPO weighting function. The primary objective is to ensure that the gradient estimator is continuous with respect to the policy probability π_θ at the clipping boundaries.

From equation (10), the gradient equation of DGPO can be explicitly expressed as:

$$\nabla_\theta \mathcal{J}_{\text{DGPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \mathcal{W}_{i,t}^{\text{DGPO}}(\theta) \hat{A}_i \nabla_\theta \pi_\theta(o_{i,t} | q, o_{i,<t}), \quad (16)$$

Using the identity $\nabla_\theta \pi_\theta = \pi_\theta \nabla_\theta \log \pi_\theta$, the effective coefficient applied to the standard score function $\nabla_\theta \log \pi_\theta$ is $\mathcal{F}_{i,t} = \mathcal{W}_{i,t}^{\text{DGPO}} \cdot \pi_\theta$. Since π_θ is continuous, ensuring the continuity of $\mathcal{W}_{i,t}^{\text{DGPO}}$ at the boundaries is sufficient to ensure the continuity of the entire gradient estimator.

Let $w_{i,t} = \frac{\pi_\theta}{\pi_{\theta_{\text{old}}}}$ denote the importance sampling ratio. The In-Boundary (M) weight is given by:

$$\mathcal{W}_{i,t}^{\text{M}} = \frac{1}{\pi_{\theta_{\text{old}}}}. \quad (17)$$

A.2.1 Left Boundary Derivation

The transition between the Left Boundary (LN) and the In-Boundary (M) region occurs when the importance sampling ratio is $w_{i,t} = 1 - \varepsilon_{\text{low}}$. At this boundary, the current policy probability is:

$$\pi_\theta = (1 - \varepsilon_{\text{low}}) \pi_{\theta_{\text{old}}}. \quad (18)$$

We then equate the weighting functions for the LN and M regions at this point:

$$\begin{aligned} \mathcal{W}_{i,t}^{\text{LN}} \Big|_{w_{i,t}=1-\varepsilon_{\text{low}}} &= \mathcal{W}_{i,t}^{\text{M}} \Big|_{w_{i,t}=1-\varepsilon_{\text{low}}} \\ C_{\text{left}} \cdot \pi_\theta^n &= \frac{1}{\pi_{\theta_{\text{old}}}} \\ C_{\text{left}} \cdot [(1 - \varepsilon_{\text{low}}) \pi_{\theta_{\text{old}}}]^n &= \frac{1}{\pi_{\theta_{\text{old}}}}. \end{aligned} \quad (19)$$

Solving for C_{left} :

$$\begin{aligned} C_{\text{left}} &= \frac{1}{\pi_{\theta_{\text{old}}} \cdot (1 - \varepsilon_{\text{low}})^n \pi_{\theta_{\text{old}}}^n} \\ &= \frac{1}{(1 - \varepsilon_{\text{low}})^n \pi_{\theta_{\text{old}}}^{n+1}}. \end{aligned} \quad (20)$$

Thus, $C_{\text{left}} = (1 - \varepsilon_{\text{low}})^{-n} \pi_{\theta_{\text{old}}}^{-(n+1)}$.

A.2.2 Right Boundary Derivation

Similarly, the transition between the Right Boundary (HP) and the In-Boundary (M) region occurs when $w_{i,t} = 1 + \varepsilon_{\text{high}}$. At this boundary:

$$\pi_\theta = (1 + \varepsilon_{\text{high}}) \pi_{\theta_{\text{old}}}. \quad (21)$$

We similarly equate the weighting functions for the HP and M regions at this point:

$$\begin{aligned} \mathcal{W}_{i,t}^{\text{HP}} \Big|_{w_{i,t}=1+\varepsilon_{\text{high}}} &= \mathcal{W}_{i,t}^{\text{M}} \Big|_{w_{i,t}=1+\varepsilon_{\text{high}}} \\ C_{\text{right}} \cdot \pi_\theta^{-\frac{1}{m}} &= \frac{1}{\pi_{\theta_{\text{old}}}} \\ C_{\text{right}} \cdot [(1 + \varepsilon_{\text{high}}) \pi_{\theta_{\text{old}}}]^{-\frac{1}{m}} &= \frac{1}{\pi_{\theta_{\text{old}}}}. \end{aligned} \quad (22)$$

Solving for C_{right} :

$$\begin{aligned} C_{\text{right}} &= \frac{(1 + \varepsilon_{\text{high}})^{\frac{1}{m}} \pi_{\theta_{\text{old}}}^{\frac{1}{m}}}{\pi_{\theta_{\text{old}}}} \\ &= (1 + \varepsilon_{\text{high}})^{\frac{1}{m}} \pi_{\theta_{\text{old}}}^{\frac{1}{m}-1}. \end{aligned} \quad (23)$$

This concludes the derivation of the constants.

Table 5: Definition of bias relative to Policy Gradient

Bias Type	Mathematical Definition
In-Boundary Bias	$\text{Bias}_{\text{Algo}}^{\text{M}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{Algo}}^{\text{M}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{M}}(\theta) \right\ $
Left-Boundary Bias	$\text{Bias}_{\text{Algo}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{Algo}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $
Right-Boundary Bias	$\text{Bias}_{\text{Algo}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{Algo}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $
Reverse Left-Boundary Bias	$\text{Bias}_{\text{Algo}}^{\text{LP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{Algo}}^{\text{LP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LP}}(\theta) \right\ $
Reverse Right-Boundary Bias	$\text{Bias}_{\text{Algo}}^{\text{HN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{Algo}}^{\text{HN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HN}}(\theta) \right\ $

A.3 Proof of Policy Gradient Bias

A.3.1 Standard Policy Gradient Estimator

The gradient estimator for the standard policy gradient method is formally given by:

$$\nabla_{\theta} \mathcal{J}_{\text{PG}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta}(\cdot|q)} [A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \quad (24)$$

By introducing Importance Sampling with the old policy, we obtain the following form:

$$\nabla_{\theta} \mathcal{J}_{\text{PG}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \left[\frac{\pi_{\theta}(o_t|q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t|q, o_{<t})} A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t}) \right] \quad (25)$$

We adopt Equation (25) as the standard policy gradient estimator in the subsequent proofs.

A.3.2 Decomposed Gradient Estimation

To distinguish gradient estimators under different conditions, we define the following binary variables based on the importance sampling ratio $r_t(\theta) = \frac{\pi_{\theta}(o_t|q, o_{<t})}{\pi_{\theta_{\text{old}}}(o_t|q, o_{<t})}$ and the advantage A_t . Note that while r_t is a function of θ , for the purpose of defining the estimator’s functional form, these regions are treated as piecewise conditions:

$$\begin{cases} v_t^{\text{LN}} = \mathbb{I}[r_t(\theta) < 1 - \varepsilon_{\text{low}} \wedge A_t < 0] \\ v_t^{\text{HP}} = \mathbb{I}[r_t(\theta) > 1 + \varepsilon_{\text{high}} \wedge A_t > 0] \\ v_t^{\text{LP}} = \mathbb{I}[r_t(\theta) < 1 - \varepsilon_{\text{low}} \wedge A_t > 0] \\ v_t^{\text{HN}} = \mathbb{I}[r_t(\theta) > 1 + \varepsilon_{\text{high}} \wedge A_t < 0] \\ v_t^{\text{M}} = 1 - (v_t^{\text{LN}} + v_t^{\text{HP}} + v_t^{\text{LP}} + v_t^{\text{HN}}) \end{cases} \quad (26)$$

The policy gradient estimator is formally decomposed and expressed as a sum of five distinct terms:

$$\nabla_{\theta} \mathcal{J}_{\text{PG}}(\theta) = \sum_{X \in \mathcal{X}} \nabla_{\theta} \mathcal{J}_{\text{PG}}^X(\theta), \quad (27)$$

where $\mathcal{X} \in \{\text{M}, \text{LN}, \text{LP}, \text{HP}, \text{HN}\}$, and the general mathematical form for each term is defined as:

$$\nabla_{\theta} \mathcal{J}_{\text{PG}}^X(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}(\cdot|q)} [v_t^X \cdot r_t A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})]. \quad (28)$$

Similarly, gradient estimators for GRPO, CISPO, GPPO, CE-GPPO, ASPO, and DGPO can be consistently decomposed into five-term sums. The bias is defined as the magnitude (L_2 norm) of the difference vector between the algorithm’s gradient estimator and the standard policy gradient under identical conditions, as shown in Table 5.

A.3.3 In-Boundary Bias

The gradient estimates for each algorithm under in-boundary conditions are analyzed below. While most algorithms (GRPO, CISPO, GPPO, CE, DGPO) maintain an unbiased estimator locally ($v_t^{\text{M}} = 1$), ASPO introduces a non-linear regularization term specifically when the advantage is positive ($A_t > 0$). The magnitude of the bias is then computed as shown in Table 6.

The relationship of in-boundary bias magnitudes is determined to be:

$$0 = \text{Bias}_{\text{DGPO}}^{\text{M}} = \text{Bias}_{\text{Others}}^{\text{M}} < \text{Bias}_{\text{ASPO}}^{\text{M}}.$$

A.3.4 Left-Boundary Bias

The gradient estimates for each algorithm under specific left-boundary conditions ($r_t < 1 - \varepsilon_{\text{low}}, A_t < 0$) are presented in Eq. (37). Let $r_0 = 1 - \varepsilon_{\text{low}}$. The magnitude of the bias is computed as shown in Table 7.

Since $0 < r_t < r_0 < 1$, it follows universally that $\text{Bias}_{\text{DGPO}}^{\text{LN}} < \text{Bias}_{\text{GRPO}}^{\text{LN}}$ and $\text{Bias}_{\text{DGPO}}^{\text{LN}} < \text{Bias}_{\text{ASPO}}^{\text{LN}}$. However, ranking other algorithms requires a quantitative analysis of the integrals.

Table 6: In-Boundary bias among various policy optimization algorithms

Algorithm	In-Boundary Bias Magnitude
GRPO, CISPO, GPPO, CE-GPPO DGPO	$\text{Bias}_X^M = \ \nabla_{\theta} \mathcal{J}_X^M(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^M(\theta)\ = 0$
ASPO	$\text{Bias}_{\text{ASPO}}^M = \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}(\cdot q)} \left[v_t^M \cdot \left(\frac{1}{r_t} - r_t \right) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $ $\approx \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}(\cdot q)} \left[v_t^M \cdot \underbrace{-2(r_t - 1)}_{\lambda_t} \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ \neq 0^*$

Decoupling Assumption and Quantitative Analysis. To rigorously compare the bias magnitudes, we apply a **decoupling assumption**: we assume the gradient norm is locally independent of the importance sampling ratio r_t within the small boundary region. Let $\delta = \mathbb{E}_{q \sim \mathcal{D}}[\|\nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})\|]$ denote the average gradient magnitude.

The explicit expressions analytically are derived as follows (integrals are over r_t):

$$\begin{aligned} \text{Bias}_{\text{GRPO}}^{\text{LN}} &= \text{Bias}_{\text{ASPO}}^{\text{LN}} \\ &= \int_0^{r_0} \underbrace{|0 - r_t|}_{\text{Coeff. Diff}} \cdot \delta \cdot k r_t^{\gamma} dr_t \\ &= \frac{k \delta r_0^{\gamma+2}}{\gamma + 2} \end{aligned} \quad (29)$$

$$\begin{aligned} \text{Bias}_{\text{CISPO}}^{\text{LN}} &= \text{Bias}_{\text{GPPO}}^{\text{LN}} \\ &= \int_0^{r_0} \underbrace{|r_0 - r_t|}_{\text{Coeff. Diff}} \cdot \delta \cdot k r_t^{\gamma} dr_t \\ &= k \delta r_0^{\gamma+2} \left(\frac{1}{\gamma + 1} - \frac{1}{\gamma + 2} \right) \end{aligned} \quad (30)$$

$$\begin{aligned} \text{Bias}_{\text{CE}}^{\text{LN}} &= \int_0^{r_0} \underbrace{|\beta_1 r_0 - r_t|}_{\text{Coeff. Diff}} \cdot \delta \cdot k r_t^{\gamma} dr_t \\ &= k \delta r_0^{\gamma+2} \left| \frac{\beta_1}{\gamma + 1} - \frac{1}{\gamma + 2} \right| \end{aligned} \quad (31)$$

$$\begin{aligned} \text{Bias}_{\text{DGPO}}^{\text{LN}} &= \int_0^{r_0} \underbrace{\left| \frac{r_t^{n+1}}{r_0^n} - r_t \right|}_{\text{Coeff. Diff}} \cdot \delta \cdot k r_t^{\gamma} dr_t \\ &= k \delta r_0^{\gamma+2} \left(\frac{1}{\gamma + 2} - \frac{1}{n + \gamma + 2} \right) \end{aligned} \quad (32)$$

where we utilize the property $r_t > r_t^{n+1}/r_0^n$ for $r_t < r_0$. In the extreme limiting case ($\gamma \rightarrow +\infty$), the resulting bias ratios are:

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{DGPO}}^{\text{LN}}}{\text{Bias}_{\text{CISPO}}^{\text{LN}}} &= \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{DGPO}}^{\text{LN}}}{\text{Bias}_{\text{GPPO}}^{\text{LN}}} \\ &= \lim_{\gamma \rightarrow \infty} \frac{n(\gamma + 1)}{\gamma + 2 + n} = n \end{aligned} \quad (33)$$

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{DGPO}}^{\text{LN}}}{\text{Bias}_{\text{CE}}^{\text{LN}}} &= \lim_{\gamma \rightarrow \infty} \frac{n(\gamma + 1)(\gamma + 2 + n)^{-1}}{|(1 - \beta_1)\gamma + 1 - 2\beta_1|} \\ &= 0 \end{aligned} \quad (34)$$

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{CISPO}}^{\text{LN}}}{\text{Bias}_{\text{CE}}^{\text{LN}}} &= \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{GPPO}}^{\text{LN}}}{\text{Bias}_{\text{CE}}^{\text{LN}}} \\ &= \lim_{\gamma \rightarrow \infty} \frac{1}{|(1 - \beta_1)\gamma + 1 - 2\beta_1|} \\ &= 0 \end{aligned} \quad (35)$$

$$\begin{aligned} \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{GRPO}}^{\text{LN}}}{\text{Bias}_{\text{CE}}^{\text{LN}}} &= \lim_{\gamma \rightarrow \infty} \frac{\text{Bias}_{\text{ASPO}}^{\text{LN}}}{\text{Bias}_{\text{CE}}^{\text{LN}}} \\ &= \lim_{\gamma \rightarrow \infty} \frac{\gamma + 1}{|(1 - \beta_1)\gamma + 1 - 2\beta_1|} \\ &= \frac{1}{|1 - \beta_1|} \end{aligned} \quad (36)$$

Assuming $n = 1$ (linear decay), we obtain the strict ranking: $0 < \text{Bias}_{\text{DGPO}}^{\text{LN}} < \text{Bias}_{\text{CISPO}}^{\text{LN}} = \text{Bias}_{\text{GPPO}}^{\text{LN}} < \text{Bias}_{\text{CE}}^{\text{LN}} < \text{Bias}_{\text{GRPO}}^{\text{LN}} = \text{Bias}_{\text{ASPO}}^{\text{LN}}$.

A.3.5 Right-Boundary Bias

The gradient estimates for the right-boundary case across algorithms are shown in Eq. (38). The bias magnitude is computed as shown in Table 8.

Given the right-boundary condition $1 < 1 + \varepsilon_{\text{high}} < r_t$, and adopting the CE-GPPO configuration $\beta_2 = 1$, we derive the ranking:

$$\begin{aligned} 0 < \text{Bias}_{\text{DGPO}}^{\text{HP}} &\leq \text{Bias}_{\text{CISPO}}^{\text{HP}} = \text{Bias}_{\text{GPPO}}^{\text{HP}} \\ &= \text{Bias}_{\text{CE}}^{\text{HP}} < \text{Bias}_{\text{GRPO}}^{\text{HP}} = \text{Bias}_{\text{ASPO}}^{\text{HP}}. \end{aligned}$$

Table 7: Left-Boundary bias among various policy optimization algorithms

Algorithm	Left-Boundary Bias Magnitude
GRPO	$\text{Bias}_{\text{GRPO}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
CISPO	$\text{Bias}_{\text{CISPO}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot (r_0 - r_t) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
GPPO	$\text{Bias}_{\text{GPPO}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot (r_0 - r_t) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
CE-GPPO	$\text{Bias}_{\text{CE}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left\{ v_t^{\text{LN}} \cdot [\beta_1 r_0 - r_t] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right\} \right\ $
ASPO	$\text{Bias}_{\text{ASPO}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
DGPO	$\text{Bias}_{\text{DGPO}}^{\text{LN}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{LN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot \left(r_t - \frac{r_t^{n+1}}{r_0^n} \right) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $

$$\begin{cases}
 \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot 0 \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot (1 - \varepsilon_{\text{low}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot (1 - \varepsilon_{\text{low}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot \beta_1 (1 - \varepsilon_{\text{low}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot 0 \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\
 \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{LN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{LN}} \cdot \frac{r_t^{n+1}}{(1 - \varepsilon_{\text{low}})^n} \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right]
 \end{cases} \quad (37)$$

A.3.6 Reverse Left-Boundary Bias

Based on the definition of ASPO, if the reverse left-boundary is treated as a clipping region $1 - \varepsilon_{\text{low}}$, the gradient estimates are listed in Eq. (39). The bias magnitude is shown in Table 9.

The reverse left-boundary bias magnitude relationship is directly derived as:

$$\begin{aligned}
 0 &= \text{Bias}_{\text{DGPO}}^{\text{LP}} = \text{Bias}_{\text{GRPO}}^{\text{LP}} = \text{Bias}_{\text{GPPO}}^{\text{LP}} \\
 &= \text{Bias}_{\text{CE}}^{\text{LP}} < \text{Bias}_{\text{CISPO}}^{\text{LP}} = \text{Bias}_{\text{ASPO}}^{\text{LP}}.
 \end{aligned}$$

A.3.7 Reverse Right-Boundary Bias

Assuming the reverse right-boundary of ASPO is $1 + \varepsilon_{\text{high}}$, the gradient estimates are listed in Eq. (40). The bias magnitude is shown in Table 10.

The reverse right-boundary bias magnitude relationship is directly derived as:

$$\begin{aligned}
 0 &= \text{Bias}_{\text{DGPO}}^{\text{HN}} = \text{Bias}_{\text{GRPO}}^{\text{HN}} = \text{Bias}_{\text{GPPO}}^{\text{HN}} \\
 &= \text{Bias}_{\text{CE}}^{\text{HN}} < \text{Bias}_{\text{CISPO}}^{\text{HN}} = \text{Bias}_{\text{ASPO}}^{\text{HN}}.
 \end{aligned}$$

A.4 Derivation of Learning Rate Scaling

To ensure consistent training dynamics across models of varying sizes (1.5B, 7B, and 14B), we employ a learning rate scaling rule based on the principle of *Constant Total Gradient Variance*. Below, we provide the formal derivation of this scaling law.

Assumptions. Consider two models with identical architecture but different parameter counts, denoted by N . We make the following standard assumptions for large-scale model training:

- The training data, batch size, and optimizer configuration remain constant.
- The gradient variance of a single parameter, denoted as σ^2 , is constant and independent of the total model size (assuming high redundancy in LLM parameters).
- The gradients of individual parameters are approximately independent.

Table 8: Right-Boundary bias among various policy optimization algorithms

Algorithm	Right-Boundary Bias Magnitude
GRPO	$\text{Bias}_{\text{GRPO}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
CISPO	$\text{Bias}_{\text{CISPO}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left\{ v_t^{\text{HP}} \cdot [r_t - (1 + \varepsilon_{\text{high}})] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right\} \right\ $
GPPO	$\text{Bias}_{\text{GPPO}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left\{ v_t^{\text{HP}} \cdot [r_t - (1 + \varepsilon_{\text{high}})] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right\} \right\ $
CE-GPPO	$\text{Bias}_{\text{CE}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left\{ v_t^{\text{HP}} \cdot [r_t - \beta_2(1 + \varepsilon_{\text{high}})] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right\} \right\ $
ASPO	$\text{Bias}_{\text{ASPO}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $
DGPO	$\text{Bias}_{\text{DGPO}}^{\text{HP}} = \left\ \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{HP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) \right\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot [r_t - (1 + \varepsilon_{\text{high}})^{\frac{1}{m}} r_t^{1 - \frac{1}{m}}] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t}) \right] \right\ $

$$\left\{ \begin{array}{l} \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot 0 \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot (1 + \varepsilon_{\text{high}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot (1 + \varepsilon_{\text{high}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot \beta_2(1 + \varepsilon_{\text{high}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot 0 \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \\ \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{HP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[v_t^{\text{HP}} \cdot (1 + \varepsilon_{\text{high}})^{\frac{1}{m}} r_t^{1 - \frac{1}{m}} \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t | q, o_{<t}) \right] \end{array} \right. \quad (38)$$

Total Gradient Variance. Let $\nabla_{\theta} \mathcal{L}$ represent the gradient vector of the loss function with respect to the model parameters $\theta \in \mathbb{R}^N$. Based on the independence assumption, the variance of the total gradient norm is the sum of the variances of individual parameter gradients:

$$\text{Var}(\nabla_{\theta} \mathcal{L}) = \sum_{i=1}^N \text{Var}(g_i) = N \cdot \sigma^2 \quad (41)$$

Stability Condition. To maintain a consistent convergence speed across different scales, we require the variance of the parameter update step (the “step size” in the parameter space) to remain constant. Let η be the learning rate. The update step is $\Delta\theta = \eta \cdot \nabla_{\theta} \mathcal{L}$. The variance of this update step is:

$$\begin{aligned} \text{Var}(\Delta\theta) &= \text{Var}(\eta \cdot \nabla_{\theta} \mathcal{L}) \\ &= \eta^2 \cdot \text{Var}(\nabla_{\theta} \mathcal{L}) = \eta^2 N \sigma^2 \end{aligned} \quad (42)$$

We define a stability constant C such that:

$$\eta^2 N \sigma^2 = C \quad (43)$$

Scaling Law Derivation. Let (η_1, N_1) be the configuration for the base model (1.5B) and (η_2, N_2) be the configuration for the target model (e.g., 7B or 14B). From Eq. (43), we have:

$$\eta_1^2 N_1 \sigma^2 = \eta_2^2 N_2 \sigma^2 = C \quad (44)$$

Eliminating constant terms σ^2 and C , we obtain the relationship:

$$\eta_1^2 N_1 = \eta_2^2 N_2 \implies \eta_2 = \eta_1 \sqrt{\frac{N_1}{N_2}} \quad (45)$$

This establishes the **Inverse Square Root Scaling Law** relative to the number of parameters.

Numerical Verification. Applying this rule to our experimental settings with the base learning rate $\eta_{1.5B} = 1.0 \times 10^{-6}$:

• **For 7B Model:**

$$\begin{aligned} \eta_{7B} &= 1.0 \times 10^{-6} \times \sqrt{\frac{1.5}{7}} \\ &\approx 1.0 \times 10^{-6} \times 0.4629 \approx 4.63 \times 10^{-7} \end{aligned}$$

Table 9: Reverse Left-Boundary bias among various policy optimization algorithms

Algorithm	Reverse Left-Boundary Bias Magnitude
GRPO, GPPO CE-GPPO, DGPO	$\text{Bias}_X^{\text{LP}} = \ \nabla_{\theta} \mathcal{J}_X^{\text{LP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LP}}(\theta)\ = 0$
CISPO, ASPO	$\text{Bias}_{\text{CISPO}}^{\text{LP}} = \ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{LP}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LP}}(\theta)\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot (1 - \varepsilon_{\text{low}} - r_t) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t})] \right\ \neq 0^*$

Table 10: Reverse Right-Boundary bias among various policy optimization algorithms

Algorithm	Reverse Left-Boundary Bias Magnitude
GRPO, GPPO CE-GPPO, DGPO	$\text{Bias}_X^{\text{HN}} = \ \nabla_{\theta} \mathcal{J}_X^{\text{HN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HN}}(\theta)\ = 0$
CISPO, ASPO	$\text{Bias}_{\text{CISPO}}^{\text{HN}} = \ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{HN}}(\theta) - \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HN}}(\theta)\ $ $= \left\ \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot [r_t - (1 + \varepsilon_{\text{high}})] \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t q, o_{<t})] \right\ \neq 0^*$

$$\left\{ \begin{array}{l} \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot (1 - \varepsilon_{\text{low}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot (1 - \varepsilon_{\text{low}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{LP}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{LP}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \end{array} \right. \quad (39)$$

$$\left\{ \begin{array}{l} \nabla_{\theta} \mathcal{J}_{\text{PG}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{GRPO}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{CISPO}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot (1 + \varepsilon_{\text{high}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{GPPO}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{CE}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{ASPO}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot (1 + \varepsilon_{\text{high}}) \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \\ \nabla_{\theta} \mathcal{J}_{\text{DGPO}}^{\text{HN}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} [v_t^{\text{HN}} \cdot r_t \cdot A_t \nabla_{\theta} \log \pi_{\theta}(o_t|q, o_{<t})] \end{array} \right. \quad (40)$$

- **For 14B Model:**

$$\begin{aligned} \eta_{14B} &= 1.0 \times 10^{-6} \times \sqrt{\frac{1.5}{14}} \\ &\approx 1.0 \times 10^{-6} \times 0.3273 \approx 3.27 \times 10^{-7} \end{aligned}$$

These calculated values correspond exactly to the learning rates reported in Table 12.

B Implementation Details

B.1 Infrastructure and Environment

Computational Resources. Our experiments are conducted on a high-performance computing cluster consisting of 30 nodes. Each node is equipped with $8 \times$ NVIDIA A100 GPUs (80GB VRAM),

interconnected via NVLink for high-bandwidth intra-node communication. The multi-node training relies on a robust Ethernet/InfiniBand fabric to ensure synchronization efficiency. The total effective training time for the largest model (14B) was approximately 500 hours.

Software Stack. We build our reinforcement learning pipeline upon VeRL (Sheng et al., 2025), a flexible framework designed for post-training.

- **Training Backend:** We utilize Fully Sharded Data Parallel (FSDP) for distributed training. This setup manages memory efficiency through parameter sharding and offloading, allowing us to train 14B models with full-parameter updates without memory overflow.

- **Inference Engine:** To maximize rollout throughput, we integrate vLLM as the inference backend. We leverage its **PagedAttention** mechanism to efficiently manage Key-Value (KV) cache memory, significantly reducing fragmentation during the generation of long reasoning chains. We align the environment configurations (CUDA version 12.4, PyTorch version 2.6.0) across all nodes to prevent numerical discrepancies.

B.2 Data and Evaluation Benchmarks

Training Dataset. We utilize the DAPO-Math-17k dataset for training. We employ the official tokenizer corresponding to the Qwen2.5-Math series to ensure consistent token mapping between the pre-trained backbone and the RL fine-tuning stage. The maximum response length is set to 8192 tokens. This length is empirically chosen to sufficiently accommodate the chain-of-thought reasoning steps required by benchmarks like AIME and MATH, while maintaining high training throughput.

Benchmark Details To rigorously evaluate the mathematical reasoning capabilities of our models, we selected a suite of diverse benchmarks. These datasets cover a spectrum of difficulty levels, from foundational high school mathematics to expert-level olympiad problems. Table 11 provides a comprehensive overview of their characteristics.

The datasets used for evaluation can be accessed via the following repositories:

- **AIME 2024:** <https://huggingface.co/datasets/math-ai/aime24>
- **AIME 2025:** <https://huggingface.co/datasets/math-ai/aime25>
- **AMC 2023:** <https://huggingface.co/datasets/math-ai/amc23>
- **MATH-500:** <https://huggingface.co/datasets/HuggingFaceH4/MATH-500>
- **Minerva:** <https://huggingface.co/datasets/math-ai/minervamath>
- **OlympiadBench:** <https://huggingface.co/datasets/math-ai/olympiadbench>

Reward Function. We employ a strict rule-based reward mechanism to rigorously verify the correctness of the generated solutions.

- **Format Verification:** We first check if the output follows the required format (e.g., enclosing the answer in `\boxed{\}`).
- **Correctness Check:** Using the `math_verify` toolkit, we compare the extracted answer against the ground truth. A reward of $r = 1$ is assigned for a correct match, and $r = -1$ otherwise. This binary reward setting poses a significant challenge for exploration, as no partial rewards are applied.

B.3 Training Protocols

Optimization Strategy. We employ the AdamW optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.95$, and a small weight decay. Unlike standard pre-training schedules, we adopt a **constant learning rate** strategy without warm-up or cosine decay. This design choice eliminates the confounding factors of learning rate scheduling, allowing us to attribute performance gains solely to the algorithmic improvements. To ensure consistent convergence dynamics across different model scales (1.5B, 7B, 14B), we calibrate the learning rate based on the *Constant Total Gradient Variance* principle.

Numerical Precision and Stability. All models are trained using `bf16` precision. We apply gradient clipping with a norm threshold of 1.0 to mitigate gradient explosion. Random seeds for model initialization, data shuffling, and sampling are fixed to 42 to ensure reproducibility.

B.4 Hyperparameter Specifications

Table 12 summarizes the key hyperparameters used across all experiments.

Controlled Variables. To ensure a fair comparison and strictly isolate the impact of different gradient weighting strategies, we unify the clipping thresholds across all experiments. Specifically, for all methods involving trust region clipping or boundary definitions (including GRPO, CISPO, GPPO, CE-GPPO, ASPO and our DGPO), we fix:

$$\varepsilon_{\text{low}} = \varepsilon_{\text{high}} = 0.2$$

This rigorous control ensures that any observed performance differences are primarily attributable to the gradient dynamics rather than minor variations in the trust region size.

Dataset	Core Description	Key Characteristics
AIME 2024	The 2024 edition of the American Invitational Mathematics Examination, serving as a bridge between the AMC and the USAMO.	<ul style="list-style-type: none"> • Focuses on arithmetic precision and number theory. • Requires answers in a strict integer format (000–999). • Tests robustness against recent contamination.
AIME 2025	The most recent iteration of the AIME competition, representing a strictly “held-out” set for evaluating generalization to unseen problems.	<ul style="list-style-type: none"> • Zero data contamination risk due to recency. • High complexity requiring multi-step logical chains. • Validates the model’s potential for future reasoning.
AMC 2023	Selected problems from the 2023 American Mathematics Competitions (AMC 10/12), representing the entry-level olympiad difficulty.	<ul style="list-style-type: none"> • Covers broad topics: Algebra, Geometry, Counting. • Functions as a baseline for competitive math ability. • Requires mapping multiple-choice logic to open-ended generation.
MATH-500	A curated subset of 500 representative problems from the widely used MATH dataset, designed by OpenAI for efficient evaluation.	<ul style="list-style-type: none"> • Spans 7 categories including Calculus and Probability. • Reduces evaluation costs while maintaining distribution fidelity. • Heavily relies on LaTeX understanding.
Minerva	A collection of technical mathematics problems derived from scientific papers and undergraduate-level coursework.	<ul style="list-style-type: none"> • Involves higher-order symbolic reasoning. • Contains domain-specific vocabulary and notation. • Tests capabilities beyond standard competition math.
OlympiadBench	A comprehensive aggregate of international mathematics competitions (e.g., IMO, CMO) spanning multiple languages and formats.	<ul style="list-style-type: none"> • Represents the upper bound of mathematical reasoning. • Includes theorem proving and fill-in-the-blank types. • Challenges the model’s cross-lingual mathematical logic.

Table 11: Detailed comparison of the mathematical reasoning benchmarks used in this study.

Table 12: Detailed hyperparameter configurations for all experiments. To ensure a fair comparison, we maintain identical training settings across all algorithms, varying only the learning rate according to the model scale.

Model Scale	Algorithm	Learning Rate	Specific Hyperparameters
<i>Common Settings: Mini-batch Size = 32, Rollout Batch Size = 512, Max Length = 8192</i>			
DeepSeek-R1-Distill-Qwen-1.5B	GRPO	1.0×10^{-6}	$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CISPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CE-GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \beta_1 = 0.75, \beta_2 = 1$
	ASPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \epsilon'_{\text{low}} = 0.33, \epsilon'_{\text{high}} = 3$
	DGPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, n = 2, m = 2$
DeepSeek-R1-Distill-Qwen-7B	GRPO	4.63×10^{-7}	$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CISPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CE-GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \beta_1 = 0.75, \beta_2 = 1$
	ASPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \epsilon'_{\text{low}} = 0.33, \epsilon'_{\text{high}} = 3$
	DGPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, n = 1, m = 2$
DeepSeek-R1-Distill-Qwen-14B	GRPO	3.27×10^{-7}	$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CISPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2$
	CE-GPPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \beta_1 = 0.75, \beta_2 = 1$
	ASPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, \epsilon'_{\text{low}} = 0.33, \epsilon'_{\text{high}} = 3$
	DGPO		$\epsilon_{\text{low}} = \epsilon_{\text{high}} = 0.2, n = 1, m = 2$

Common Configurations. We set the KL coefficient $\beta_{\text{KL}} = 0$. By removing the explicit KL penalty, we rely implicitly on the trust region con-

straints imposed by the clipping mechanisms to prevent policy collapse. The rollout generation uses a temperature of 1.0 and a top- p of 1.0.

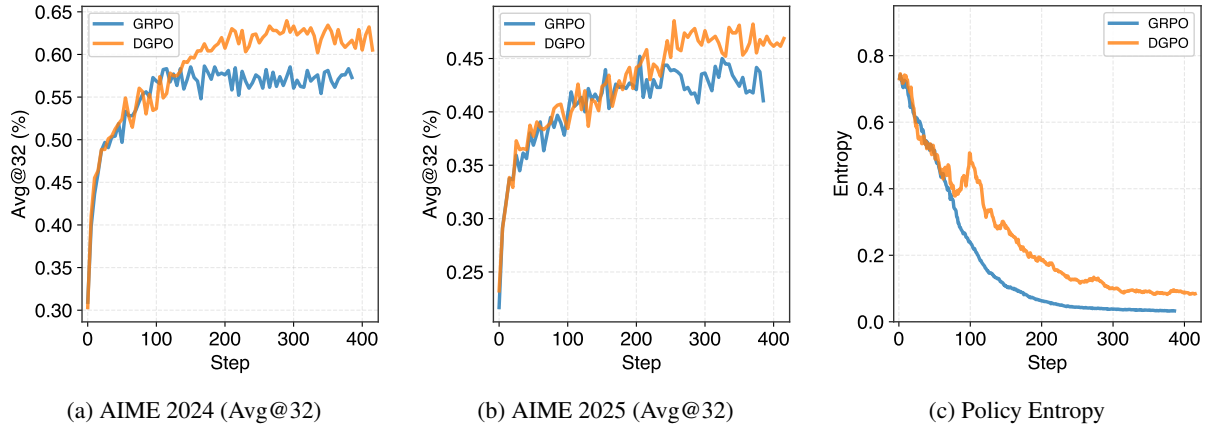


Figure 5: Training dynamics of DeepSeek-R1-Distill-Qwen-14B comparing GRPO and DGPO.

Method-Specific Settings. The method-specific configurations are as follows:

- **CE-GPPO:** Adopts scaled soft clipping with $\beta_1 = 0.75$ and $\beta_2 = 1.0$, as recommended in the original paper.
- **ASPO:** Use $\varepsilon'_{\text{low}} = 0.33$ and $\varepsilon'_{\text{high}} = 3$ as soft dual clip threshold.
- **DGPO:** We introduce the decoupled decay parameters n and m :
 - **1.5B:** $n = 2, m = 2$.
 - **7B & 14B:** $n = 1, m = 2$.

The continuity constants C_{left} and C_{right} are automatically calculated based on $\pi_{\theta_{\text{old}}}$.

C Additional Experimental Results

C.1 14B Model Training Dynamics

We visualize the training trajectories of the DeepSeek-R1-Distill-Qwen-14B model in Figure 5. Similar to the 1.5B and 7B models, DGPO on 14B scale demonstrates faster convergence and higher asymptotic performance on AIME benchmarks while maintaining a stable entropy reduction curve, avoiding the collapse issues seen in RLVR.

C.2 Detailed 14B Performance Comparison

Table 13 provides comprehensive evaluation results for 14B model, including both Avg@32 (expected performance) and Pass@32 (potential capability).

C.3 Pass@K Statistics on AIME Benchmarks

We report Pass@K ($k \in \{1, 2, 4, 8, 16, 32\}$) metrics specifically for the **AIME 2024** and **AIME 2025** benchmarks across 1.5B, 7B, and 14B scales

(Table 14). The results demonstrate that DGPO consistently achieves higher coverage of the solution space (higher Pass@K) compared to baselines, particularly as k increases.

Table 13: Full performance comparison on DeepSeek-R1-Distill-Qwen-14B.

Method	AIME24		AIME25		AMC23		MATH500		Minerva		Olympiad		Avg.	
	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32	A@32	P@32
DEEPSEEK-R1-DISTILL-QWEN-14B														
GRPO	56.6	82.9	40.5	67.8	92.2	97.4	66.5	70.4	22.6	35.3	43.2	50.7	53.6	67.4
DGPO	63.3	86.2	47.6	66.4	93.9	99.1	67.0	75.1	23.0	37.9	45.1	57.4	56.7	70.4

Table 14: Detailed Pass@K performance on AIME 2024 and AIME 2025 across all model scales.

Method	P@1	P@2	P@4	P@8	P@16	P@32
DeepSeek-R1-Distill-Qwen-1.5B						
<i>Dataset: AIME 2024</i>						
GRPO	33.2	43.2	53.2	61.6	67.9	71.8
CISPO	34.8	44.1	52.0	58.7	64.4	69.1
GPPO	29.6	38.4	47.1	53.6	57.9	60.5
CE-GPPO	35.1	44.4	52.6	60.0	66.2	70.2
ASPO	36.4	44.4	53.1	61.7	68.5	73.2
DGPO	43.3	53.4	63.1	70.6	75.6	79.3
<i>Dataset: AIME 2025</i>						
GRPO	27.7	33.0	38.1	42.6	46.5	49.9
CISPO	25.8	31.8	37.6	43.6	49.2	53.3
GPPO	23.5	29.3	34.3	40.0	46.4	51.9
CE-GPPO	27.7	33.1	38.3	44.1	50.1	55.1
ASPO	28.3	33.0	37.1	41.2	45.9	51.5
DGPO	32.8	38.3	43.6	48.6	52.7	56.1
DeepSeek-R1-Distill-Qwen-7B						
<i>Dataset: AIME 2024</i>						
GRPO	48.2	58.0	66.6	73.4	78.6	82.5
CISPO	51.6	60.9	68.5	73.5	75.8	76.6
GPPO	43.1	52.1	59.3	65.0	69.7	72.5
CE-GPPO	48.7	55.7	62.3	68.1	73.0	76.9
ASPO	51.8	61.2	68.5	73.4	76.7	79.6
DGPO	55.5	64.1	71.6	77.0	80.2	81.9
<i>Dataset: AIME 2025</i>						
GRPO	37.4	44.4	50.8	55.2	58.1	60.5
CISPO	38.2	45.2	51.7	56.5	60.8	65.4
GPPO	31.7	38.5	45.8	53.0	58.6	62.5
CE-GPPO	36.4	43.1	49.1	53.9	57.4	60.4
ASPO	37.1	43.3	48.4	51.2	52.6	54.1
DGPO	43.1	49.7	55.2	60.1	64.6	68.0
DeepSeek-R1-Distill-Qwen-14B						
<i>Dataset: AIME 2024</i>						
GRPO	56.6	69.3	76.0	79.8	81.9	82.9
DGPO	63.3	70.4	76.4	81.1	84.6	86.2
<i>Dataset: AIME 2025</i>						
GRPO	40.5	48.4	54.2	60.1	64.7	67.8
DGPO	47.6	53.7	58.7	61.7	63.9	66.4