

# LAD: Learning Advantage Distribution for Reasoning

Wendi Li      Sharon Li

Department of Computer Sciences  
University of the Wisconsin-Madison  
{wli679, sharonli}@cs.wisc.edu

## Abstract

Current reinforcement learning objectives for large-model reasoning primarily focus on maximizing expected rewards. This paradigm can lead to overfitting to dominant reward signals, while neglecting alternative yet valid reasoning trajectories, thereby limiting diversity and exploration. To address this issue, we introduce Learning Advantage Distributions (LAD), a distribution-matching framework that replaces advantage maximization with learning the advantage-induced distribution. By establishing the equivalence between the optimal policy update and an advantage-based target distribution, we derive a practical LAD objective formulated as minimizing an  $f$ -divergence between the policy-induced and advantage-induced distributions. This yields a gradient update that increases likelihood for high-advantage responses while suppressing over-confident probability growth, preventing collapse without requiring auxiliary entropy regularization. LAD incurs no extra training cost compared to GRPO and scales naturally to LLM post-training. In a controlled bandit setting, LAD faithfully recovers the multimodal advantage distribution, validating the theoretical formulation. Experiments on math and code reasoning tasks across several LLM backbones show that LAD reliably improves both accuracy and generative diversity.

## 1 Introduction

Large language models have recently shown impressive performance across a wide range of complex reasoning tasks (Jaech et al., 2024; Guo et al., 2025; Yang et al., 2025). A central approach driving this progress is reinforcement learning with verifiable rewards (RLVR) (Guo et al., 2025; Shao et al., 2024). In RLVR, response correctness is automatically verified using external tools such as compilers, symbolic solvers, or mathematical checkers. By replacing noisy and unreliable feedback with deterministic, verification-based rewards,

RLVR enables more stable optimization and mitigates reward hacking, leading to substantial gains in complex logical reasoning tasks (Shao et al., 2025; Wang et al., 2025c; Liu et al., 2025b; Luo et al., 2025b; Wang et al., 2025b; Jiang et al., 2025a; Dong et al., 2025; Da et al., 2025; Luo et al., 2025a).

However, current RLVR objectives for large-model reasoning are built around maximizing expected reward, a formulation that inherently overfits to dominant reward signals. This can drive policies toward mode collapses, suppressing alternative reasoning trajectories and reducing exploratory diversity (Cui et al., 2025b; Wu and Choi, 2025). As illustrated in Figure 1 (left), the policy distribution trained with the reward maximization objective (GRPO) concentrates on the most desirable action and neglects other meaningful modes. This reveals a fundamental limitation: the reward-maximization paradigm at the heart of RLVR imposes a structural bottleneck on the diversity and exploration (Liu et al., 2025a; Zhu et al., 2025). Although entropy-based constraints and advantage regularization methods have been proposed to mitigate this issue (Yu et al., 2025; Cheng et al., 2025; Wang et al., 2025a; Deng et al., 2025), these approaches largely retain reward expectation maximization as their core objective.

To address this limitation, we propose a framework that makes the policy Learn the Advantage Distribution (LAD). Our work takes a perspective shift: we formalize policy optimization as learning to match an advantage-induced distribution, rather than maximizing expected advantage. This shift turns RLVR from a scalar-expectation objective into a distribution-matching problem, allowing the policy to preserve multiple promising reasoning modes instead of collapsing onto a single high-reward trajectory (Figure 1, right). We develop both the theoretical LAD objective and a practical surrogate that avoids intractable computation while




Figure 1: Advantage distribution and policy distribution on 50-arm bandit trained via different training objectives.

preserving the same optima. The training cost of LAD is the same as standard GRPO, making LAD scalable to real-world LLM reasoning tasks. Concurrent work such as FlowRL (Zhu et al., 2025) also departs from pure reward maximization with a reward-matching objective, but it does not learn the full advantage-induced action distribution and can be proved as a more constrained instance of our broader distribution-matching framework.

Finally, we validate LAD through both controlled experiments and LLM reasoning evaluations. In a carefully designed bandit environment, LAD can learn the policy distribution that faithfully aligns with the underlying advantage-induced distribution. Additionally, we conduct experiments on six challenging mathematical reasoning and three coding reasoning benchmarks across multiple LLM backbones. LAD consistently improves both accuracy and diversity over GRPO, entropy-regularized baselines, and FlowRL. Further, we conduct comprehensive ablations to understand LAD’s gains across divergence choices and training conditions. Collectively, these findings validate LAD as a principled and effective training objective for RL-based post-training of LLMs, particularly in complex reasoning settings. We summarize our contributions as follows:

1. We introduce LAD, a new objective that re-frames RLVR as matching the advantage-induced action distribution instead of maximizing expected advantage.
2. We derive a practical surrogate with theoretical justification and improved optimization behavior, enabling scaling to real-world LLM reasoning tasks with comparable training cost as GRPO.
3. Through extensive experiments and systematic ablations, we provide empirical evidence that LAD achieves superior performance and improves diversity.

## 2 Preliminary

**Reward Maximization in RL for LLMs.** We consider a contextual bandit formulation of reinforcement learning (RL) for LLMs. Let  $x \in \rho$  denote a prompt sampled from a distribution  $\rho$ , and let the model (policy)  $\pi_\theta(y|x)$  generate a response  $y \in \mathcal{Y}$ . Upon observing the response, the environment returns a scalar reward  $r(x, y) \in \mathbb{R}$ . The objective of most existing RL algorithms is to learn policy parameters  $\theta$  that maximize the expected reward,  $J(\pi_\theta) = \mathbb{E}_{x \sim \rho, y \sim \pi_\theta(\cdot|x)}[r(x, y)]$ . Gradient-based methods (Williams, 1992) optimize this objective by exploiting the identity between  $\nabla_\theta J(\pi_\theta)$  and

$$\mathbb{E}_{x \sim \rho, y \sim \pi_\theta(\cdot|x)}[-A(x, y) \nabla_\theta \log \pi_\theta(y|x)], \quad (1)$$

where  $A(x, y)$  is an advantage function, commonly approximated by  $r(x, y) - b(x)$  with  $b(x)$  serving as a baseline to reduce gradient variance.

**Trust-Region Constrained Policy Optimization.** Direct gradient ascent on  $J(\pi_\theta)$  can induce large policy updates, often resulting in unstable training. Trust-region methods mitigate this issue by constraining each policy update to remain close to a behavior (sampling) policy  $\pi_{\text{old}}$ . A canonical formulation underlying Trust Region Policy Optimization (Schulman et al., 2015) and its practical variant Proximal Policy Optimization (Schulman et al., 2017) is given by:

$$\begin{aligned} \max_{\theta} \mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot|x)} \left[ \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} \cdot A(x, y) \right] \quad (2) \\ \text{s.t. } \mathbb{E}_{x \sim \rho} [D_{\text{KL}}(\pi_{\text{old}}(\cdot|x) \parallel \pi_\theta(\cdot|x))] \leq \epsilon, \end{aligned}$$

where  $\epsilon > 0$  controls the maximum allowable deviation from the behavior policy via a KL-divergence constraint.

**Reinforcement Learning with Verifiable Reward (RLVR)** is a post-training paradigm in which rewards are assigned by deterministic, rule-based

verifiers. Because such rewards are less susceptible to reward hacking, RLVR has proven particularly effective for reasoning-intensive tasks. Representative RLVR algorithms, including GRPO (Shao et al., 2024) and RLOO (Ahmadian et al., 2024), typically incorporate group-wise normalization or ranking-based estimators to reduce variance in advantage estimation. Despite these modifications, the core objective of existing RLVR methods remains the maximization of expected advantage, as in Eq. (1). This objective often leads to a rapid collapse of exploratory diversity, which tends to concentrate probability mass on a small set of high-reward responses and suppress alternative valid reasoning paths (Cui et al., 2025b; Wu and Choi, 2025). Motivated by this limitation, we propose to shift the learning objective from advantage maximization to learning the distribution of advantages.

### 3 Method

In this section, we propose **Learning Advantage Distribution (LAD)**, a framework that trains policies by matching the distribution induced by advantages rather than maximizing expected advantage. Building on the structure implied by trust-region policy optimization, we formulate LAD as minimizing an  $f$ -divergence between a policy-induced distribution and an advantage-induced distribution (Sec. 3.1). We further derive a practical surrogate loss that preserves the same optimal policy while enabling scalable training for LLMs (Sec. 3.2).

#### 3.1 Learn Advantage Distribution with $f$ -divergence

To align the policy with the distribution induced by advantage values, we first define two distributions: one derived from the policy  $\pi_\theta(y|x)$  and the other from the advantage function  $A(x, y)$ . At optimality, these two distributions coincide. We state this relationship through the following lemma, which directly motivates our training objective.

**Lemma 3.1 (Distribution equivalence).** *Let  $\pi_\theta$  denote the current policy parameterized by  $\theta$ , and let  $\pi_{\text{old}}$  be the behavior model. Suppose that  $\pi_\theta$  is optimized using a trust-region constrained reinforcement learning algorithm (e.g., PPO (Schulman et al., 2017)) as in Eq. 2. There exists a Lagrange multiplier  $\eta > 0$ , such that the optimal policy  $\pi_\theta^*$  satisfies*

$$\frac{\pi_\theta^*(y|x)}{\pi_{\text{old}}(y|x)} = \frac{e^{\frac{A(x,y)}{\eta}}}{Z_A(x)}, \quad (3)$$

where  $\eta$  is a Lagrangian multiplier,  $Z_\pi(x) = \sum_y \frac{\pi_\theta^*(y|x)}{\pi_{\text{old}}(y|x)}$  and  $Z_A(x) = \sum_y e^{\frac{A(x,y)}{\eta}}$  are two normalizations.

Lemma 3.1 implies that trust-region constrained optimization implicitly induces a policy-improvement probability distribution proportional to the exponentiated advantage. Rather than treating this form as a theoretical result, we make it explicit and treat it as a learning target. This leads to a shift in perspective: policy optimization can be viewed as learning to match an advantage-induced distribution, rather than maximizing expected advantage.

Formally, we define two distributions below:

$$\mathcal{P}_{\pi_\theta}(y|x) = \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)Z_\pi(x)}, \quad (4)$$

$$\mathcal{P}_A(y|x) = \frac{e^{\frac{A(x,y)}{\eta}}}{Z_A(x)}. \quad (5)$$

Intuitively,  $\mathcal{P}_A$  is the target advantage distribution, encouraging higher probability to higher-advantage responses;  $\mathcal{P}_{\pi_\theta}$  is the policy-induced distribution measuring how much the policy deviates from the behavior policy. The constructions of  $\mathcal{P}_A$  and  $\mathcal{P}_{\pi_\theta}$  are directly inspired by the optimal policy form arising from trust-region constrained policy optimization. This connection ensures that the proposed distributions are theoretically consistent with established policy improvement principles.

Given these two equivalent distributions, we formulate the learning objective by minimizing their divergence. To this end, we recall the definition of the  $f$ -divergence, which is a general class of discrepancy measures that subsumes many commonly used divergences (e.g., KL, Jensen-Shannon, Hellinger etc.).

**Definition 3.1 ( $f$ -divergence).** *Let  $\mathcal{P}$  and  $\mathcal{Q}$  be two distributions, then for a convex function  $f : [0, +\infty) \rightarrow (-\infty, +\infty)$  such that  $f(x)$  is finite for all  $x > 0$ , and  $f(1) = 0$ , the  $f$ -divergence of  $\mathcal{P}$  from  $\mathcal{Q}$  is defined as*

$$D_f(\mathcal{P}||\mathcal{Q}) = \int_{\Omega} f\left(\frac{d\mathcal{P}}{d\mathcal{Q}}\right)d\mathcal{Q}. \quad (6)$$

**LAD: Minimizing the divergence between  $\mathcal{P}_{\pi_\theta}$  and  $\mathcal{P}_A$ .** Using this measure of distributional discrepancy, we propose to train the policy by minimizing the  $f$ -divergence between  $\mathcal{P}_{\pi_\theta}$  and  $\mathcal{P}_A$ ,

thereby encouraging the policy to learn the advantage distribution. This divergence-minimization objective constitutes the central novelty of LAD:

$$\mathcal{L}_{\text{LAD}}^{\text{theorem}} = \mathbb{E}_{x \sim \rho} \sum_y \frac{e^{\frac{A(x,y)}{\eta}}}{Z_A(x)} f\left(\frac{\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}}{e^{\frac{A(x,y)}{\eta}}}\right) \cdot \frac{Z_A(x)}{Z_{\pi_\theta}(x)}. \quad (7)$$

Importantly, minimizing the LAD objective preserves the same optimal policy as standard trust-region policy optimization, namely  $\pi^* \propto \pi_{\text{old}} \cdot e^{\frac{A(x,y)}{\eta}}$ . While the resulting theoretical formulation is principled, it involves normalization terms  $Z_A(x)$  and  $Z_{\pi_\theta}(x)$  that are intractable for large action spaces such as those encountered in large language models. To address this challenge, we introduce a practical surrogate objective in the next section that is computationally tractable while preserving the same optimal policy.

### 3.2 Practical Implementations

Instead of approximating the intractable normalization constants directly, our key insight is that the optimal policy of an  $f$ -divergence objective depends only on the relative structure of its weighting functions. This allows us to replace the theoretical LAD objective with a surrogate loss that is tractable yet induces the same optimal policy. To facilitate this transition, we introduce the following lemma.

**Lemma 3.2** (Linking loss formulations to their induced optimal policies.). *Let  $\mathcal{L} = \mathbb{E}_{x \sim \rho} \sum_y g_2(x, y) f\left(\frac{p_\theta(y|x)}{g_1(x, y)}\right)$ , where  $p_\theta(\cdot)$  is a probability measure with  $\sum_y p(\pi) = 1$ .  $g_1(x, y)$ ,  $g_2(x, y)$  are functions independent of  $p_\theta$ . If the ratio  $\frac{g_2(x, y)}{g_1(x, y)} = Z_g(x)$  depends only on  $x$  and is independent of  $y$ , then any optimal solution  $p_\theta^*$  of  $\mathcal{L}$  satisfies  $p_\theta^*(y|x) \propto g_1(x, y)$ .*

**Interpretation.** Applying Lemma 3.2 to the theoretical LAD objective immediately yields the optimal policy  $\pi_\theta^* \propto \pi_{\text{old}} \cdot e^{\frac{A(x,y)}{\eta}}$ . Crucially, the lemma also implies that this optimal policy can be recovered by alternative loss formulations, provided that the induced ratio between the weighting functions remains invariant with respect to  $y$ . In particular, when setting  $p_\theta = \pi_\theta$  and choosing  $g_1(x, y) = g_2(x, y) = \pi_{\text{old}} \cdot e^{\frac{A(x,y)}{\eta}}$ , the resulting objective admits the same optimal policy,  $\pi_\theta^* \propto \pi_{\text{old}} \cdot e^{\frac{A(x,y)}{\eta}}$ , while avoiding the intractable normalization terms present in the theoretical LAD

loss. This observation enables transforming the theoretical LAD objective into a practical surrogate that is computationally tractable while preserving the same optimal policy.

#### LAD practical objective:

$$\mathcal{L}_{\text{LAD}} = \mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot|x)} e^{\frac{A(x,y)}{\eta}} f\left(\frac{\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}}{e^{\frac{A(x,y)}{\eta}}}\right). \quad (8)$$

This practical objective admits the same optimal policy as the theoretical loss  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$ , while exhibiting favorable optimization behavior and closely approximating it in practice.

### 3.3 Analyses and Discussions

$\mathcal{L}_{\text{LAD}}$  yields desirable loss behavior with implicit regularization. The gradient of the practical LAD loss with respect to the policy parameters  $\nabla_\theta \mathcal{L}_{\text{LAD}}$  can be written as,

$$\mathbb{E}_{x \sim \rho, y \sim \pi_\theta(\cdot|x)} f'\left(\frac{\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}}{e^{\frac{A(x,y)}{\eta}}}\right) \cdot \nabla_\theta \log \pi_\theta(y|x). \quad (9)$$

In contrast to the vanilla policy gradient in Eq. (1), which increases the log-likelihood of all actions with positive advantage, the weighting term in the LAD gradient,  $f'\left(\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} \cdot e^{-\frac{A(x,y)}{\eta}}\right)$  depends jointly on the advantage  $A(x, y)$  and the likelihood ratio  $\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}$ . As a result, even when a trajectory  $(x, y)$  has a large positive advantage, the LAD gradient naturally suppresses further increases in  $\log \pi_\theta(y|x)$  once the likelihood ratio becomes sufficiently large. This adaptive weighting prevents uncontrolled amplification of action probabilities and induces an *implicit regularization* effect, leading to more stable optimization and improved diversity preservation. Unlike PPO-style clipping as detailed in Appendix B.4, which enforces hard constraints on likelihood ratios, LAD induces a soft, advantage-aware suppression mechanism that emerges naturally from distribution matching.

$\mathcal{L}_{\text{LAD}}$  is a close surrogate to the theoretical LAD objective  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$ . Beyond inducing desirable gradient behavior, the practical LAD objective also closely approximates the theoretical LAD formulation, thereby serving as a principled and reliable surrogate. A formal bound is given in Theorem B.1, and additional discussion and empirical validation are provided in Appendix B.3. Finally, we summarize several representative choices of divergence functions within the LAD framework in Table 5.

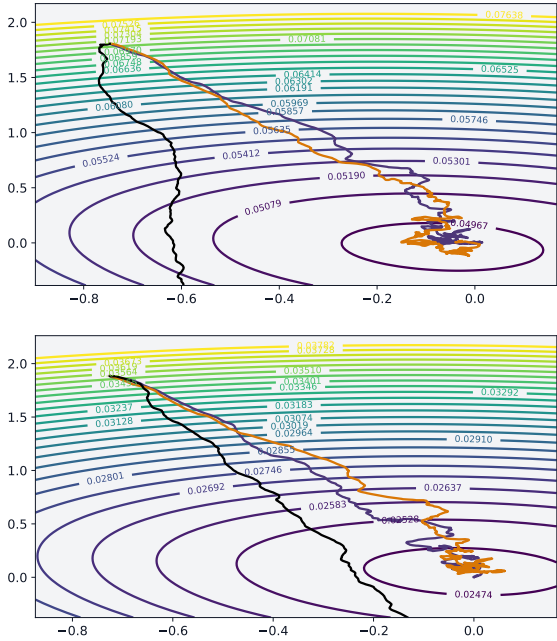


Figure 2: Learning dynamics of GRPO, the theoretical LAD loss, and the practical LAD loss. The loss landscape is constructed with the practical LAD loss. Top: LAD loss under Hellinger distance. Bottom: LAD loss under Jensen–Shannon divergence.

## 4 Experiments

We empirically evaluate LAD to answer three key questions: (i) Does the LAD objective successfully learn advantage distributions in controlled settings? (ii) Does this translate into improvements on large-scale reasoning tasks? (iii) How do different divergence choices and hyperparameters influence performance? To this end, we begin with a controlled distribution-matching task, followed by extensive LLM experiments, and conclude with a set of ablation studies.

### 4.1 Controlled Experiments

We begin with a controlled bandit environment to visualize the behavior of different objectives. This setting *allows us to directly compare the learned policy distribution against the ground-truth advantage distribution*, revealing whether an objective genuinely performs distribution matching.

**Settings.** We design a 50-bandit problem. The policy model is parameterized by a  $1 \times 50$  vector, where each entry represents the probability of selecting a corresponding arm. The policy model  $\pi_\theta$  and the behavior model  $\pi_{\text{old}}$  are initialized to the uniform distribution. The underlying advantage distribution is constructed as a mixture of three Gaussian distributions, resulting in three distinct

modes. Training is performed for 4000 steps. At each step, the behavior policy  $\pi_{\text{old}}$  is sampled 32 times with a sampling temperature of 1.5. The behavior policy is updated twice throughout training. We use a learning rate of  $5e-3$ . For LAD loss variants,  $\eta$  is set to 1. For GRPO, we additionally apply a KL-divergence regularization between  $\pi$  and  $\pi_{\text{old}}$  with weight 1. Under this setting, the theoretical optimal policy of GRPO coincides with that of the LAD objective, enabling a fair comparison.

### LAD loss successfully matches the two distributions.

Figure 1 visualizes the final learned policy distributions alongside the ground-truth advantage distribution. The advantage distribution is shown in orange, while the learned policy distribution is shown in blue. The LAD objective  $\mathcal{L}_{\text{LAD}}$  enables the policy-induced distribution  $\mathcal{P}_{\pi_\theta}$  to closely match the target advantage distribution  $\mathcal{P}_A$ , faithfully recovering all three modes. In contrast, the advantage-maximization objective of GRPO concentrates probability mass on a single arm. This leads to a collapse of the policy distribution and a substantial reduction in generative diversity during training, ultimately resulting in suboptimal performance on the less preferred arm.

### LAD loss $\mathcal{L}_{\text{LAD}}$ is an accurate surrogate of $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$ .

We further examine the learning dynamics induced by  $\mathcal{L}_{\text{LAD}}$ ,  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$  and GRPO. The corresponding trajectories are visualized in Figure 2. To construct the loss landscape, we use the practical LAD objective and project the model parameters from each training step onto this landscape. As shown in Figure 2, the optimization trajectories induced by the practical loss and the theoretical loss closely align, exhibiting highly similar learning dynamics. In contrast, GRPO follows a markedly different trajectory, deviating from the geometry implied by the theoretical LAD objective. These results empirically confirm that the practical LAD loss serves as a faithful and accurate surrogate for the theoretical formulation.

**Comparison with FlowRL.** Concurrent with our work, FlowRL (Zhu et al., 2025) proposes to move beyond reward maximization by adopting a reward-matching objective derived from the trajectory balance loss in GFlowNets (Bengio et al., 2023). While this formulation departs from conventional reward-maximization objectives, Figure 1 (middle) shows that FlowRL fails to fully align the advantage-induced distribution  $\mathcal{P}_A$  and the

Table 1: Performance of different losses on math reasoning benchmarks. The best two results under Avg@32 are highlighted in **bold** and underlined.

Methods	MATH	Olympiad Bench	Minerva	AIME 2024	AIME 2025	AMC	Average
GRPO(Shao et al., 2024)	75.57	39.47	36.32	14.03	4.86	53.44	37.28
EntAdv(Cheng et al., 2025)	76.60	<u>41.16</u>	35.68	15.00	7.23	52.97	38.11
KLCov(Cui et al., 2025b)	76.40	40.13	36.28	14.66	<b>9.81</b>	<u>53.45</u>	<u>38.46</u>
ClipCov(Cui et al., 2025b)	76.03	40.30	<b>37.58</b>	<u>16.29</u>	7.39	52.27	38.31
FlowRL(Zhu et al., 2025)	<u>77.17</u>	40.80	36.21	14.67	7.51	47.34	37.28
LAD	<b>77.66</b>	<b>41.25</b>	<u>36.68</u>	<b>19.60</b>	<u>8.84</u>	<b>56.45</b>	<b>40.08</b>

Table 2: Performance of different losses on code reasoning benchmarks. The best two results under a specific metric are highlighted in **bold** and underlined.

Method	LiveCodeBench		HumanEval+		CodeForces	
	Avg@16	Pass@16	Avg@16	Pass@16	Rating	Percentile
GRPO(Shao et al., 2024)	32.46	51.25	80.90	93.25	1355.35	70.30 %
EntAdv(Cheng et al., 2025)	32.75	50.90	81.05	95.09	1382.78	72.20%
KLCov(Cui et al., 2025b)	29.50	50.18	80.75	95.09	1373.07	71.60%
ClipCov(Cui et al., 2025b)	32.10	<b>51.97</b>	<u>81.48</u>	<u>94.49</u>	<u>1473.09</u>	<u>79.20%</u>
FLowRL(Zhu et al., 2025)	<u>33.24</u>	51.61	81.29	95.09	1364.70	71.00%
LAD	<b>33.51</b>	<b>51.97</b>	<b>82.29</b>	<b>95.71</b>	<b>1533.64</b>	<b>82.50%</b>

policy-induced distribution  $\mathcal{P}_{\pi_{\theta}}$ . As detailed in Appendix B.5, FlowRL can in fact be cast as a special case of the LAD framework, but only under substantially stricter conditions on policy updates and parameter scaling. In contrast, LAD provides a more general and theoretically grounded framework that directly learns the advantage distribution through principled  $f$ -divergence minimization, thus allowing LAD to more faithfully capture the full structure of the advantage distribution.

## 4.2 LLM Experiments

Having verified that LAD recovers the target distribution in a toy setting, we now test whether these benefits translate to real-world LLM tasks.

**Experimental settings.** We experiment on both the math and code domains. For the math domain, we use Qwen2.5-7B (Qwen et al., 2025) as our LLM backbone and DAPO-MATH (Yu et al., 2025) as our training dataset. The experimental protocol follows Cui et al. (2025b) with a dynamic sampling strategy and the clip-higher technique (Yu et al., 2025). Evaluation is conducted on a diverse set of mathematical reasoning benchmarks, including MATH500 (Lightman et al., 2023), AIME 2024, AIME 2025 (Li et al., 2024), AMC (Li et al., 2024), OlympiadBench (He et al., 2024), and Minerva (Lewkowycz et al., 2022). For the code domain, we use DeepSeek-R1-Distill-7B as our LLM backbone and DeepCoder (Luo et al., 2025b) as our training

dataset. We follow the experimental protocol of (Zhu et al., 2025), and evaluate the performance on LiveCodeBench (Jain et al., 2025), CodeForces (Penedo et al., 2025), and HumanEval+ (Chen et al., 2021). For both domains, we employ GRPO as the advantage estimator and use Jensen-Shannon divergence to implement LAD loss. The learning rate is set to 1e-6. More implementation details are in Appendix C.2.

**Baselines.** We compare against the latest methods designed to mitigate entropy collapse and improve generative diversity, including EntAdv, KLCov, and ClipCov (Cui et al., 2025b), as well as FlowRL (Zhu et al., 2025). EntAdv, KLCov, and ClipCov are regulation-based approaches that introduce additional regularization terms into the loss function; however, they retain expected advantage maximization as the primary optimization objective. In contrast, FlowRL represents the most recent departure from advantage maximization by explicitly learning a reward-matching objective.

**LAD leads to superior reasoning performance.** Table 1 reports the performance of different training objectives on a suite of mathematical reasoning benchmarks. Across nearly all benchmarks, LAD consistently outperforms prior methods, achieving the best average performance under both evaluation protocols. In particular, LAD attains the highest Avg@32 score overall (40.08), with substantial gains on challenging benchmarks such as

Table 3: Performance of LAD and FlowRL on different LLM backbones. The best result is highlighted in **bold**.

Methods	MATH	Olympiad Bench	Minerva	AIME 2024	AIME 2025	AMC	Average
OpenMath-Nemotron-1.5B: Avg@32							
FlowRL	80.43	45.81	21.02	21.36	22.66	55.47	41.13
LAD	<b>86.81</b>	<b>55.63</b>	<b>22.98</b>	<b>32.70</b>	<b>26.47</b>	<b>68.71</b>	<b>48.88</b>
OpenMath-Nemotron-1.5B: Pass@32							
FlowRL	94.60	65.93	38.24	56.35	37.60	80.65	66.98
LAD	<b>97.20</b>	<b>78.52</b>	<b>42.28</b>	<b>72.29</b>	<b>54.38</b>	<b>89.08</b>	<b>72.29</b>
Thinkless-DeepScaleR-1.5B: Avg@32							
FlowRL	80.90	42.71	<b>34.11</b>	19.41	16.90	54.80	41.47
LAD	<b>83.01</b>	<b>46.08</b>	33.71	<b>21.85</b>	<b>20.13</b>	<b>59.38</b>	<b>44.03</b>
Thinkless-DeepScaleR-1.5B: Pass@32							
FlowRL	96.40	69.19	56.62	51.56	<b>43.64</b>	84.79	67.03
LAD	<b>95.60</b>	<b>69.93</b>	<b>60.66</b>	<b>55.52</b>	42.60	<b>86.62</b>	<b>68.49</b>

Table 4: Generative diversity results of on rollouts of AIME 24/25 problems. The best two results under a specific metric are highlighted in **bold** and underlined.

Methods	dist-3( $\uparrow$ )	dist-4( $\uparrow$ )	GPT4-Judge( $\uparrow$ )
GRPO	0.2306	0.2902	2.04
EntAdv	0.2978	0.3782	2.07
KLCov	0.2471	0.3262	2.07
ClipCov	<u>0.3315</u>	<u>0.4171</u>	2.10
FLowRL	0.2878	0.3654	<u>2.38</u>
LAD	<b>0.3498</b>	<b>0.4442</b>	<b>2.58</b>

AIME 2024 (+3.31 over the strongest baseline) and AMC (+3). Performance under the Pass@32 metric is also reported in Appendix C.3, where LAD achieves a new state-of-the-art average score of 65.76, outperforming FlowRL by 2.63 points and regulation-based methods by an even larger margin. A similar performance advantage is observed on code reasoning tasks (Table 2), where LAD achieves the best results on LiveCodeBench and HumanEval+ and the highest CodeForces percentile (82.50%), indicating not only improved reasoning accuracy but also superior generalization to real-world programming problem difficulty. Taken together, these results demonstrate that LAD is broadly effective at improving both mathematical and code reasoning across diverse benchmarks and evaluation settings.

**LAD boosts generative diversity.** To evaluate the effectiveness of LAD in enhancing response diversity, we follow Zhu et al. (2025) and assess logical diversity using an LLM-as-a-Judge framework with GPT-4-Turbo on AIME 2024/2025 as detailed in Appendix C.2. In addition to logical diversity, we measure lexical diversity using standard automatic metrics, distinct- $n$  with  $n=3,4$ . As shown in Table 4, LAD attains the highest distinct-3 and distinct-4 scores, indicating substantially richer lex-




Figure 3: The average math reasoning performance of LAD loss variants under different  $f$ -divergence classes. The dashed lines represent the average performances of the **strict** and **non-strict** divergence classes.

ical variation than all baselines. Meanwhile, LAD achieves the highest GPT-4 Judge score among all methods, suggesting that a diverse logical pattern accompanies its increased diversity.

**LAD can be extended to different LLM backbones.** Besides the Qwen2.5-7B for math reasoning and DeepSeek-R1-Distill-7B for code reasoning, we further include two other LLM backbones, OpenMath-Nemotron-1.5B and Thinkless-DeepScaleR-1.5B. As shown in Table 3, LAD still outperforms FlowRL on both backbones.

### 4.3 Further Studies

**Effect of divergence classes: strict divergences outperform weaker ones.** The math reasoning performances of all LAD variants are shown in Figure 3. We comprehensively consider the KL divergence (KL), the reverse KL divergence (rKL), the Jeffreys divergence (JF), the total variation distance (TV), the Hellinger Distance (HD), and the Jensen-Shannon divergence (JS). Among all the divergence classes in our experiments, TV, HD, and JS are strict divergences that force exact matching at their minimum. On the contrary, KL, rKL, and JF




Figure 4: Average Avg@16 and Pass@16 scores over six mathematical benchmarks with different values of the hyperparameter  $\eta$ .




Figure 5: The distribution of GPU time per token throughout the training process.

are weaker divergences. Minimizing them does not necessarily imply distributional equality in practice. Overall, strict divergences generally lead to higher overall performance because they penalize discrepancies symmetrically, thereby helping LAD learn the advantage distribution more faithfully, which in turn leads to improved reasoning outcomes.

**Effect of  $\eta$ .** To evaluate the sensitivity of LAD to  $\eta$ , we train models with six values of  $\eta \in \{0.5, 1, 2, 4, 8, 16\}$ . Figure 4 reports the average Avg@16 and Pass@16 scores across six mathematical reasoning benchmarks. These results suggest that LAD consistently outperforms GRPO under a broad range of  $\eta$ , indicating that its effectiveness is not critically dependent on precise tuning of  $\eta$ .

**Computational cost of LAD is comparable to GRPO.** Figure 5 compares the training costs of LAD with those of the baseline methods. FlowRL and KLCov incur the highest computational overhead, as both require additional evaluations of the reference model probabilities; moreover, FlowRL introduces an extra auxiliary module to approximate an otherwise intractable term in its objective. In contrast, LAD does not require any additional model components or extra forward passes. Its implementation involves only a minimal modification to the loss function, resulting in a training cost that is comparable to that of the vanilla GRPO baseline.

## 5 Related Works

Reinforcement learning has emerged as a major approach for LLM post-training (Qwen et al., 2025; Grattafiori et al., 2024; Team et al., 2023), which have been widely applied to various domains, including instruction following (Stiennon et al., 2020; Ouyang et al., 2022) and human-value alignment (Rafailov et al., 2024; Azar et al., 2024; Meng et al., 2024; Dong et al., 2024). More recently, reinforcement learning has demonstrated substantial gains in enhancing the reasoning capabilities of LLMs (Li and Li, 2025; Yuan et al., 2025; Cui et al., 2025a). In this context, reinforcement learning with verifiable rewards (RLVR) has emerged as a particularly effective approach, owing to its robust and reliable reward signals (Zeng et al., 2025; Hu et al., 2025; Wang et al., 2024; Zhao et al., 2025; Yue et al., 2025; Zhang et al., 2025). Despite these advances, most existing RL methods for LLMs continue to rely on advantage maximization as the core training objective, which inherently biases policies toward a small set of high-reward responses and limits generative diversity (Cui et al., 2025b; Liu et al., 2025a). Concurrent with our work, FlowRL (Zhu et al., 2025) also moves beyond advantage maximization to learn reward-matching. However, as evidenced in Figure 1 and Appendix B.5, FlowRL exhibits limited distribution matching behavior and can be interpreted as a special case of our LAD framework under stricter conditions. Additional related work and detailed comparisons with traditional RL paradigms, including distributional RL, risk-sensitive RL and multi-objective RL, are provided in Appendix A.

## 6 Conclusion

Current RL objectives for large-model reasoning primarily maximize expected rewards, which encourages overfitting to dominant outcomes and suppresses viable alternative reasoning paths. We instead introduce Learning Advantage Distributions (LAD), a distribution-matching framework that optimizes policy updates by minimizing an  $f$ -divergence between the policy and the advantage-induced action distribution. This approach preserves multiple promising reasoning modes and avoids mode collapse without additional regularization. Experiments on multiple math and code reasoning benchmarks show that LAD consistently improves both accuracy and generative diversity over strong RLVR baselines.

## Limitations

While LAD demonstrates promising improvements in performance and diversity, several limitations remain. First, our work focuses on post-training scenarios where responses can be evaluated with verifiable signals (e.g., mathematical or programmatic correctness). Although this setting is common in recent RL research for LLMs, extending LAD to domains without objective verifiers, such as open-ended dialogue or creative writing, may require additional mechanisms for reward signal quality. This represents an orthogonal challenge rather than a failure of the method itself.

Second, our empirical evaluations focus on mathematical and code reasoning benchmarks. These benchmarks are representative of RLVR-style settings but do not encompass all possible LLM capabilities. Broader evaluation across multimodal tasks, multilingual reasoning, and agentic decision-making is left for future work. Meanwhile, our evaluation is conducted on 1.5B- and 7B-scale models. These sizes are standard for research-stage RL studies, but we do not presume automatic transfer to 30B–70B+ or frontier-scale models. Scaling studies are left for future work and are expected to be influenced by factors such as sampling budgets, hyperparameters, and verifier reliability, rather than limitations of LAD itself.

Finally, although LAD provides theoretical justification, we do not claim formal convergence guarantees for all divergence classes or optimization schedules. This is a standard limitation shared with most large-model RL methods, and strengthening theoretical guarantees remains an interesting direction for future research rather than a missing requirement for practical validity.

## Ethic Considerations

This work focuses on reinforcement learning for reasoning tasks, where correctness can be automatically evaluated. As such, our experiments are conducted exclusively in domains with objective ground-truth verification to reduce risks associated with generating misleading or harmful content.

Our approach of learning advantage distribution (LAD) introduced here operate purely at the optimization and distribution-learning level and do not involve new data sources, new modalities, or sensitive information. In deployment settings, LAD should be paired with standard safety practices for LLMs (e.g., alignment checks, content filtering,

and supervised oversight), which are orthogonal to the contributions of this work.

## Acknowledgement

We thank Changdae Oh and Yu Wang for their valuable suggestions on the paper. This work is supported in part by the AFOSR Young Investigator Program under award number FA9550-23-1-0184, National Science Foundation under awards IIS-2237037 and IIS-2331669, Office of Naval Research under grant number N00014-23-1-2643, Schmidt Sciences Foundation, Open Philanthropy, Alfred P. Sloan Fellowship, and gifts from Google and Amazon.

## References

- Axel Abels, Diederik Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher. 2019. Dynamic weights in multi-objective deep reinforcement learning. In *International conference on machine learning*, pages 11–20. PMLR.
- Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin, Ahmet Üstün, and Sara Hooker. 2024. Back to basics: Revisiting reinforce-style optimization for learning from human feedback in llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12248–12267.
- Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pages 4447–4455. PMLR.
- Marc G Bellemare, Will Dabney, and Rémi Munos. 2017. A distributional perspective on reinforcement learning. In *International conference on machine learning*, pages 449–458. PMLR.
- Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. 2023. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55.
- Shicong Cen, Jincheng Mei, Katayoon Goshvadi, Hanjun Dai, Tong Yang, Sherry Yang, Dale Schuurmans, Yuejie Chi, and Bo Dai. 2024. [Value-incentivized preference optimization: A unified approach to online and offline RLHF](#). *CoRR*, abs/2405.19320.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela

- Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *CoRR*, abs/2107.03374.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and Furu Wei. 2025. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*.
- Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu Yu, Qixin Xu, Weize Chen, and 1 others. 2025a. Process reinforcement through implicit rewards. *arXiv preprint arXiv:2502.01456*.
- Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen Fan, Huayu Chen, Weize Chen, and 1 others. 2025b. The entropy mechanism of reinforcement learning for reasoning language models. *arXiv preprint arXiv:2505.22617*.
- Jeff Da, Clinton Wang, Xiang Deng, Yuntao Ma, Nikhil Barhate, and Sean Hendryx. 2025. Agent-rlvr: Training software engineering agents via guidance and environment rewards. *arXiv preprint arXiv:2506.11425*.
- Will Dabney, Mark Rowland, Marc Bellemare, and Rémi Munos. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Nirjhar Das, Souradip Chakraborty, Aldo Pacchiano, and Sayak Ray Chowdhury. 2024. Active preference optimization for sample efficient rlhf. In *ICML 2024 Workshop on Theoretical Foundations of Foundation Models*. OpenReview.net.
- Jia Deng, Jie Chen, Zhipeng Chen, Wayne Xin Zhao, and Ji-Rong Wen. 2025. Decomposing the entropy-performance exchange: The missing keys to unlocking effective reinforcement learning. *arXiv preprint arXiv:2508.02260*.
- Guanting Dong, Hangyu Mao, Kai Ma, Licheng Bao, Yifei Chen, Zhongyuan Wang, Zhongxia Chen, Jiazhen Du, Huiyang Wang, Fuzheng Zhang, and 1 others. 2025. Agentic reinforced policy optimization. *arXiv preprint arXiv:2507.19849*.
- Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. 2024. RLHF workflow: From reward modeling to online RLHF. *CoRR*, abs/2405.07863.
- Yingjie Fei, Zhuoran Yang, and Zhaoran Wang. 2021. Risk-sensitive reinforcement learning with function approximation: A debiasing approach. In *International Conference on Machine Learning*, pages 3198–3207. PMLR.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, 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*.
- Conor F Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa M Zintgraf, Richard Dazeley, Fredrik Heintz, and 1 others. 2021. A practical guide to multi-objective reinforcement learning and planning. *arXiv preprint arXiv:2103.09568*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Jingcheng Hu, Yinmin Zhang, Qi Han, Daxin Jiang, Xiangyu Zhang, and Heung-Yeung Shum. 2025. Openreasoner-zero: An open source approach to scaling up reinforcement learning on the base model. *arXiv preprint arXiv:2503.24290*.
- Guanhua Huang, Tingqiang Xu, Mingze Wang, Qi Yi, Xue Gong, Siheng Li, Ruibin Xiong, Kejiao Li, Yuhao Jiang, and Bo Zhou. 2025. Low-probability tokens sustain exploration in reinforcement learning with verifiable reward. *arXiv preprint arXiv:2510.03222*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, and 1 others. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. Livecodebench: Holistic and contamination free evaluation of large language models for code. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24–28, 2025*. OpenReview.net.
- Kaixuan Ji, Jiafan He, and Quanquan Gu. 2024. [Reinforcement learning from human feedback with active queries](#). *CoRR*, abs/2402.09401.
- Xue Jiang, Yihong Dong, Mengyang Liu, Hongyi Deng, Tian Wang, Yongding Tao, Rongyu Cao, Binhua Li, Zhi Jin, Wenpin Jiao, and 1 others. 2025a.

- Coderl+: Improving code generation via reinforcement with execution semantics alignment. *arXiv preprint arXiv:2510.18471*.
- Yuxian Jiang, Yafu Li, Guanxu Chen, Dongrui Liu, Yu Cheng, and Jing Shao. 2025b. Rethinking entropy regularization in large reasoning models. *arXiv preprint arXiv:2509.25133*.
- Jean Kaddour. 2026. Target policy optimization. *arXiv preprint arXiv:2604.06159*.
- Panagiotis Kyriakis, Jyotirmoy Deshmukh, and Paul Bogdan. 2022. Pareto policy adaptation. In *International Conference on Learning Representations*.
- Prashanth La and Mohammad Ghavamzadeh. 2013. Actor-critic algorithms for risk-sensitive mdps. *Advances in neural information processing systems*, 26.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, and 1 others. 2022. Solving quantitative reasoning problems with language models. *Advances in neural information processing systems*, 35:3843–3857.
- Jia Li, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Huang, Kashif Rasul, Longhui Yu, Albert Q Jiang, Ziju Shen, and 1 others. 2024. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 13(9):9.
- Wendi Li and Yixuan Li. 2025. Process reward model with q-value rankings. In *The Thirteenth International Conference on Learning Representations*.
- Wendi Li, Changdae Oh, and Yixuan Li. 2025. General exploratory bonus for optimistic exploration in rlhf. *arXiv preprint arXiv:2510.03269*.
- 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*.
- Shiau Hong Lim and Ilyas Malik. 2022a. Distributional reinforcement learning for risk-sensitive policies. *Advances in Neural Information Processing Systems*, 35:30977–30989.
- Shiau Hong Lim and Ilyas Malik. 2022b. Distributional reinforcement learning for risk-sensitive policies. *Advances in Neural Information Processing Systems*, 35:30977–30989.
- Mingjie Liu, Shizhe Diao, Jian Hu, Ximing Lu, Xin Dong, Hao Zhang, Alexander Bukharin, Shaokun Zhang, Jiaqi Zeng, Makesh Narsimhan Sreedhar, and 1 others. 2025a. Scaling up rl: Unlocking diverse reasoning in llms via prolonged training. *arXiv preprint arXiv:2507.12507*.
- Zichen Liu, Changyu Chen, Chao Du, Wee Sun Lee, and Min Lin. 2024. Sample-efficient alignment for llms. *arXiv preprint arXiv:2411.01493*.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025b. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*.
- Kun Luo, Hongjin Qian, Zheng Liu, Ziyi Xia, Shitao Xiao, Siqi Bao, Jun Zhao, and Kang Liu. 2025a. Inflow: Reinforcing search agent via reward density optimization. *arXiv preprint arXiv:2510.26575*.
- Michael Luo, Sijun Tan, Roy Huang, Ameen Patel, Alpay Ariyak, Qingyang Wu, Xiaoxiang Shi, Rachel Xin, Colin Cai, Maurice Weber, and 1 others. 2025b. Deepcoder: A fully open-source 14b coder at o3-mini level. *Notion Blog*.
- Xiaoteng Ma, Junyao Chen, Li Xia, Jun Yang, Qianchuan Zhao, and Zhengyuan Zhou. 2025. Dsac: Distributional soft actor-critic for risk-sensitive reinforcement learning. *Journal of Artificial Intelligence Research*, 83.
- Borislav Mavrin, Hengshuai Yao, Linglong Kong, Kaiwen Wu, and Yaoliang Yu. 2019. Distributional reinforcement learning for efficient exploration. In *International conference on machine learning*, pages 4424–4434. PMLR.
- Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235.
- Oliver Mihatsch and Ralph Neuneier. 2002. Risk-sensitive reinforcement learning. *Machine learning*, 49(2):267–290.
- Mohammad Mirzanejad, Morteza Ebrahimi, Peter Vamplew, and Hadi Veisi. 2022. An online scalarization multi-objective reinforcement learning algorithm: Topsis q-learning. *The Knowledge Engineering Review*, 37:e7.
- Hossam Mossalam, Yannis M Assael, Diederik M Roijers, and Shimon Whiteson. 2016. Multi-objective deep reinforcement learning. *arXiv preprint arXiv:1610.02707*.
- 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.
- Guilherme Penedo, Anton Lozhkov, Hynek Krdlíček, Loubna Ben Allal, Edward Beeching, Agustín Piqueres Lajarín, Quentin Gallouédec, Nathan Habib, Lewis Tunstall, and Leandro von Werra. 2025. Codeforces. <https://huggingface.co/datasets/open-r1/codeforces>.

- Matteo Pirota, Simone Parisi, and Marcello Restelli. 2015. Multi-objective reinforcement learning with continuous pareto frontier approximation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Diederik M Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48:67–113.
- Mark Rowland, Robert Dadashi, Saurabh Kumar, Rémi Munos, Marc G Bellemare, and Will Dabney. 2019. Statistics and samples in distributional reinforcement learning. In *International Conference on Machine Learning*, pages 5528–5536. PMLR.
- Makoto Sato, Hajime Kimura, and Shibenobu Kobayashi. 2001. Td algorithm for the variance of return and mean-variance reinforcement learning. *Transactions of the Japanese Society for Artificial Intelligence*, 16(3):353–362.
- 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*.
- Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, and 1 others. 2025. Spurious rewards: Rethinking training signals in rlvr. *arXiv preprint arXiv:2506.10947*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Yun Shen, Michael J Tobia, Tobias Sommer, and Klaus Obermayer. 2014. Risk-sensitive reinforcement learning. *Neural computation*, 26(7):1298–1328.
- Ruizhe Shi, Runlong Zhou, and Simon S. Du. 2024. [The crucial role of samplers in online direct preference optimization](#). *CoRR*, abs/2409.19605.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Richard S Sutton. 1988. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44.
- Fahim Tajwar, Anikait Singh, Archit Sharma, Rafael Rafailov, Jeff Schneider, Tengyang Xie, Stefano Ermon, Chelsea Finn, and Aviral Kumar. 2024. [Preference fine-tuning of llms should leverage suboptimal, on-policy data](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Kristof Van Moffaert, Madalina M Drugan, and Ann Nowé. 2013. Scalarized multi-objective reinforcement learning: Novel design techniques. In *2013 IEEE symposium on adaptive dynamic programming and reinforcement learning (ADPRL)*, pages 191–199. IEEE.
- Kristof Van Moffaert and Ann Nowé. 2014. Multi-objective reinforcement learning using sets of pareto dominating policies. *The Journal of Machine Learning Research*, 15(1):3483–3512.
- Jun Wang, Meng Fang, Ziyu Wan, Muning Wen, Jiachen Zhu, Anjie Liu, Ziqin Gong, Yan Song, Lei Chen, Lionel M Ni, and 1 others. 2024. Openr: An open source framework for advanced reasoning with large language models. *arXiv preprint arXiv:2410.09671*.
- Kaiwen Wang, Nathan Kallus, and Wen Sun. 2023. Near-minimax-optimal risk-sensitive reinforcement learning with cvar. In *International Conference on Machine Learning*, pages 35864–35907. PMLR.
- Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, and 1 others. 2025a. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*.
- Yinjie Wang, Ling Yang, Ye Tian, Ke Shen, and Mengdi Wang. 2025b. Co-evolving llm coder and unit tester via reinforcement learning. *arXiv preprint arXiv:2506.03136*.
- Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, and 1 others. 2025c. Reinforcement learning for reasoning in large language models with one training example. *arXiv preprint arXiv:2504.20571*.

- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256.
- Fang Wu and Yejin Choi. 2025. The invisible leash: Why rlvr may not escape its origin. In *2nd AI for Math Workshop@ ICML 2025*.
- Tengyang Xie, Dylan J. Foster, Akshay Krishnamurthy, Corby Rosset, Ahmed Awadallah, and Alexander Rakhlin. 2024. Exploratory preference optimization: Harnessing implicit  $q^*$ -approximation for sample-efficient RLHF. *CoRR*, abs/2405.21046.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Rui Yang, Xiaoman Pan, Feng Luo, Shuang Qiu, Han Zhong, Dong Yu, and Jianshu Chen. 2024. Rewards-in-context: Multi-objective alignment of foundation models with dynamic preference adjustment. *arXiv preprint arXiv:2402.10207*.
- Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan. 2019. A generalized algorithm for multi-objective reinforcement learning and policy adaptation. *Advances in neural information processing systems*, 32.
- 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*.
- Lifan Yuan, Wendi Li, Huayu Chen, Ganqu Cui, Ning Ding, Kaiyan Zhang, Bowen Zhou, Zhiyuan Liu, and Hao Peng. 2025. Free process rewards without process labels. In *Forty-second International Conference on Machine Learning*.
- Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. 2025. Does reinforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv preprint arXiv:2504.13837*.
- Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. 2025. Simplerl-zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv preprint arXiv:2503.18892*.
- Qingyang Zhang, Haitao Wu, Changqing Zhang, Peilin Zhao, and Yatao Bian. 2025. Right question is already half the answer: Fully unsupervised llm reasoning incentivization. *arXiv preprint arXiv:2504.05812*.
- Shangdong Zhang, Bo Liu, and Shimon Whiteson. 2021. Mean-variance policy iteration for risk-averse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10905–10913.
- Shenao Zhang, Donghan Yu, Hiteshi Sharma, Ziyi Yang, Shuohang Wang, Hany Hassan, and Zhaoran Wang. 2024. Self-exploring language models: Active preference elicitation for online alignment. *CoRR*, abs/2405.19332.
- Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. 2025. Absolute zero: Reinforced self-play reasoning with zero data. *arXiv preprint arXiv:2505.03335*.
- Tianyu Zheng, Tianshun Xing, Qingshui Gu, Taoran Liang, Xingwei Qu, Xin Zhou, Yizhi Li, Zhoufutu Wen, Chenghua Lin, Wenhao Huang, and 1 others. 2025. First return, entropy-eliciting explore. *arXiv preprint arXiv:2507.07017*.
- Xuekai Zhu, Daixuan Cheng, Dinghui Zhang, Hengli Li, Kaiyan Zhang, Che Jiang, Youbang Sun, Ermo Hua, Yuxin Zuo, Xingtai Lv, and 1 others. 2025. Flowrl: Matching reward distributions for llm reasoning. *arXiv preprint arXiv:2509.15207*.

## A Related Works

### A.1 Traditional RL beyond Expected-Scalar Maximization

Beyond the standard formulation of reinforcement learning (RL) as maximizing expected cumulative scalar reward or advantage, a substantial body of work has explored alternative objectives that encode richer behavioral preferences.

**Distributional RL.** Distributional reinforcement learning (RL) (Bellemare et al., 2017; Dabney et al., 2018; Mavrin et al., 2019; Rowland et al., 2019; Lim and Malik, 2022a) departs from expectation-based value estimation by modeling the full distribution of returns and optimizing distributional objectives, thereby enabling finer-grained control over outcome profiles. Despite the similarity in terminology, distributional RL and learning the advantage distribution (LAD) address fundamentally different problems. In distributional RL, the learned distribution  $Z^\pi(x, y)$  captures the stochasticity of returns for a fixed state–action pair, reflecting randomness induced by the environment and transition dynamics. The return distribution thus serves as an intermediate representation to improve value estimation, while policy optimization remains expectation-based through  $\mathbb{E}_\pi[Z^\pi(x, y)]$ . In contrast, LAD models the distribution over the entire action space induced by advantage values. This distribution appears directly in the learning objective, which minimizes a divergence of the form  $\min_\pi \mathbb{D}_f(\mathcal{P}_\mathcal{A} || \mathcal{P}_\pi)$ , and therefore plays a central role in policy optimization rather than serving as an auxiliary estimator.

**Risk-sensitive RL.** Risk-sensitive RL (Mihatsch and Neuneier, 2002; Shen et al., 2014; Fei et al., 2021; Lim and Malik, 2022a; Ma et al., 2025) extends the standard RL objective by optimizing criteria that reflect attitudes toward uncertainty rather than mean performance alone. In many classical formulations (Sato et al., 2001; La and Ghavamzadeh, 2013; Zhang et al., 2021), risk sensitivity is introduced by augmenting the expected return with an additional term that penalizes variability or tail risk. More recent approaches include variance-constrained actor–critic methods and distributional techniques that explicitly model the return distribution and optimize risk-aware criteria such as conditional value-at-risk (CVaR) within that framework (Wang et al., 2023; Lim and Malik, 2022b). Nevertheless, similar to distributional RL,

the distributions considered in risk-sensitive RL are induced by environmental stochasticity, and policy optimization typically retains an expectation-based form, albeit with a modified objective reflecting risk preferences.

**Multi-objective RL.** Multi-objective reinforcement learning (MORL) (Van Moffaert et al., 2013; Mossalam et al., 2016; Abels et al., 2019; Yang et al., 2019; Hayes et al., 2021; Yang et al., 2024) generalizes the standard RL framework by considering vector-valued reward functions and learning policies that optimize trade-offs among multiple, potentially conflicting objectives. Common solution concepts include learning Pareto-optimal policies (Van Moffaert and Nowé, 2014; Kyriakis et al., 2022), approximating Pareto fronts (Pirotta et al., 2015), or optimizing scalarized objectives defined by linear or nonlinear combinations of the individual reward components (Mirzanejad et al., 2022; Roijers et al., 2013). In these approaches, the multiple objectives are typically specified exogenously and correspond to distinct semantic criteria (e.g., performance, cost, safety), and policy optimization proceeds by reducing the vector-valued returns to a scalar objective through scalarization or preference weights. In contrast, learning the advantage distribution (LAD) does not assume multiple predefined reward dimensions. Instead, LAD induces a distribution over actions from scalar advantage values and directly matches this distribution in the learning objective. As a result, LAD modifies the policy objective through distributional alignment over actions rather than through explicit multi-criteria optimization or scalarization of vector-valued rewards.

### A.2 Exploration-exploitation trade-off in RL for LLMs

The exploration–exploitation trade-off (Sutton, 1988) is a foundational challenge in reinforcement learning. In the context of large language models (LLMs), most existing work has focused on improving exploitation through increasingly sophisticated objectives. However, a growing body of research argues that insufficient exploration substantially limits the effectiveness of current RL-based fine-tuning methods for LLMs, motivating explicit mechanisms to promote exploration during training.

**Optimistic exploration in RLHF.** Several recent works (Xie et al., 2024; Tajwar et al., 2024) in rein-

forcement learning from human feedback (RLHF) emphasize the importance of deliberate exploration during policy optimization. These approaches introduce structured exploration mechanisms to mitigate mode collapse and premature convergence, including modifying sampling strategies during rollouts (Xie et al., 2024; Shi et al., 2024) and selecting training samples based on uncertainty estimates (Dong et al., 2024; Das et al., 2024; Ji et al., 2024). Other methods explicitly encourage exploration by optimizing objectives beyond immediate reward maximization, such as incorporating uncertainty-aware rewards (Liu et al., 2024) or adding exploration bonuses derived from disagreement or novelty signals (Zhang et al., 2024; Cen et al., 2024; Li et al., 2025). Collectively, these works demonstrate that, in the LLM setting, exploration is not merely a training stability issue but a critical factor for discovering high-quality and diverse generations under sparse, noisy, or biased reward signals. While effective, many of these methods rely on computationally expensive uncertainty estimation or auxiliary models. In contrast, LAD promotes sustained exploration by reformulating the advantage maximization objective as a distribution-matching objective, naturally maintaining higher policy entropy throughout training without additional uncertainty modeling.

**Entropy collapse and regulation in RLVR.** A related line of work (Cui et al., 2025b; Deng et al., 2025; Huang et al., 2025) identifies rapid entropy collapse in the early stages of reinforcement learning with verifiable rewards (RLVR), motivating approaches that encourage exploration by augmenting advantage maximization with entropy-based regularization (Cheng et al., 2025; Wang et al., 2025a; Zheng et al., 2025; Jiang et al., 2025b). However, these methods retain advantage maximization as the primary objective, limiting the extent to which entropy regularization can counteract exploitation-dominated updates. Moreover, entropy regularization generally alters the optimization landscape and may shift the optimum away from the original policy improvement objective. In contrast, LAD directly reformulates the learning objective by shifting from advantage maximization to learning the advantage distribution, while preserving the same optimal policy implied by trust-region policy optimization. This reformulation enables principled exploration without introducing auxiliary regularizers or modifying the target optimum. Notably, a

recent work, TPO (Kaddour, 2026), shares a similar idea to bridge the advantage distribution and policy distribution with a different loss objective.

## B Analyses & Proofs

### B.1 Proofs of Lemma 3.1

**Lemma 3.1.** *Let  $\pi_\theta$  denote the current policy parameterized by  $\theta$ , and let  $\pi_{\text{old}}$  be the behavior model. Suppose that  $\pi_\theta$  is optimized using a trust-region constrained reinforcement learning algorithm (e.g., PPO (Schulman et al., 2017)) as in Eq. 2. There exists a Lagrange multiplier  $\eta > 0$ , such that the policy update satisfies*

$$\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} = \frac{e^{\frac{A(x,y)}{\eta}}}{Z_A(x)}, \quad (10)$$

where  $\eta$  is a Lagrangian multiplier,  $Z_\pi(x) = \sum_y \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}$  and  $Z_A(x) = \sum_y e^{\frac{A(x,y)}{\eta}}$  are two normalizations.

*Proof.* Consider one step of policy improvement that maximizes the expected advantage under the new policy while staying close to the old policy  $\pi_{\text{old}}$  (trust-region). For each prompt  $x$ , we can formulate a constrained maximization problem as

$$\begin{aligned} \max_{\pi_\theta} \sum_y A(x,y)\pi_\theta(y|x) \quad \text{s.t.} \quad \sum_y \pi_\theta(y|x) = 1 \\ \text{and} \quad D_{\text{KL}}(\pi_\theta(\cdot|x) \parallel \pi_{\text{old}}(\cdot|x)) \leq \epsilon. \end{aligned} \quad (11)$$

We can form the per-state Lagrangian with multipliers  $\eta \geq 0$  and  $\lambda$  as follows,

$$\begin{aligned} L = \sum_y A(x,y)\pi_\theta(y|x) + \lambda(\sum_y \pi_\theta(y|x) - 1) \\ - \eta \sum_y (\pi_\theta(y|x) \log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}) - \epsilon. \end{aligned} \quad (12)$$

Stationarity w.r.t.  $\pi_\theta$  gives

$$A(x,y) - \eta(\log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} + 1) + \lambda = 0, \quad (13)$$

which equals  $\log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} = \frac{A(x,y) + \lambda - \eta}{\eta}$ . With some algebra, we have

$$\frac{\pi(y|x)}{\pi_{\text{old}}(y|x)} = \frac{e^{\frac{A(x,y)}{\eta}}}{Z_A(x)}, \quad (14)$$

where  $Z_\pi(x) = \sum_y \frac{\pi(y|x)}{\pi_{\text{old}}(y|x)}$  and  $Z_A(x) = \sum_y e^{\frac{A(x,y)}{\eta}}$  are two normalizations.  $\square$

## B.2 Proofs of Lemma 3.2

**Lemma 2.** Let  $\mathcal{L} = \mathbb{E}_{x \sim \rho} \sum_y g_2(x, y) f\left(\frac{p_\theta(y|x)}{g_1(x, y)}\right)$ , where  $p_\theta(\cdot)$  is a probability measure with  $\sum_y p(\pi) = 1$ .  $g_1(x, y), g_2(x, y)$  are two functions independent to  $p_\theta$ .  $p_\theta(y|x)$  always satisfy  $p_\theta(y|x) \propto g_1(x, y)$  as long as the fraction  $\frac{g_2(x, y)}{g_1(x, y)} = Z_g(x)$  is irrespective to  $y$ .

*Proof.* Similar to Lemma 3.1, we adopt the same constrained maximization problem with Lagrange multipliers, and the stationarity condition can be formulated as

$$\frac{g_2(x, y)}{g_1(x, y)} f'\left(\frac{p(\pi(y|x))}{g_1(x, y)}\right) = \lambda, \quad (15)$$

where  $\lambda$  is a Lagrange multiplier of  $\sum_y p(\pi(y|x)) = 1$ . Since  $f'$  is an increasing function according to the definition in Definition 3.1, we have for a fixed state  $x, \forall y, \frac{p(\pi(y|x))}{g_1(x, y)}$  is a constant, thus  $p_\theta(y|x) \propto g_1(x, y)$ .  $\square$

## B.3 LAD is an $\mathcal{O}(\delta)$ -accurate surrogate of the theoretical LAD objective

First, we made some standard assumptions for the following proofs.

**Assumption B.1** (finite actions with uniform lower bound  $p_{\min}$ ). For any policy  $\pi$  and any state-action pair  $x, y, 1 \geq p_{\max} \geq \pi(y|x) \geq p_{\min} > 0$ .

**Assumption B.2** (bounded advantage). For any state-action pair  $(x, y)$ , the advantage value is bounded by a small value  $|A(x, y)| < \gamma$

**Assumption B.3** (small update per step). Under a small learning rate, we assume that for any  $(x, y)$ ,  $\left| \frac{\pi(y|x)}{\pi_{\text{old}}(y|x)} - 1 \right| < \delta$

The first assumption is generally established in the LLM setting due to a finite vocabulary size. For the second assumption, the advantage scale is controllable by constraining the reward scale and normalization, which are generally adopted in existing implementations like GRPO (Shao et al., 2024). For the third assumption, since the sampling distribution  $\pi_{\text{old}}$  is frequently updated (generally within ten steps), the overall distribution of  $\pi$  would not have a sharp change compared to  $\pi_{\text{old}}$ . Hence, Assumption B.3 also holds empirically, especially considering that standard RL implementations for LLMs generally have small learning rates and gradient clipping mechanisms.

With the above assumptions, we obtain the following conditions,

$$\left| \frac{Z_A(x)}{Z_\pi(x)} - 1 \right| < \frac{\mathcal{O}(\gamma + \delta)}{1 + \mathcal{O}(\delta)} \quad (16)$$

Holding this inequality, we can then prove that the practical LAD loss in Eq. 8 is an  $\mathcal{O}(\delta)$ -accurate surrogate of the theoretical LAD objective in Eq. 7.

**Theorem B.1** (LAD is an  $\mathcal{O}(\delta)$ -accurate surrogate of the theoretical LAD objective). Under Assumptions B.1, B.2, B.3, the gap between the theoretical LAD objective  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$  and the practical LAD objective  $\mathcal{L}_{\text{LAD}}$  is bounded by

$$\left| \mathcal{L}_{\text{LAD}}^{\text{theorem}} - \mathcal{L}_{\text{LAD}} \right| < \mathcal{O}(\gamma + \delta) \quad (17)$$

*Proof.* In the proof, we denote  $c_\theta = \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} \cdot e^{\frac{-A(x, y)}{\eta}}$  for simplicity. First, we re-express the expectations in  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$  and  $\mathcal{L}_{\text{LAD}}$  with respect to a common base distribution  $\mu_x$ , and absorb the differences into weights.

$$\begin{aligned} \mathcal{L}_{\text{LAD}}^{\text{theorem}} &= \mathbb{E}_{x \sim \rho, y \sim \mu_x} \frac{e^{\frac{A(x, y)}{\eta}}}{\mu_x(y) Z_A(x)} f\left(c_\theta \frac{Z_A(x)}{Z_{\pi_\theta}(x)}\right), \\ \mathcal{L}_{\text{LAD}} &= \mathbb{E}_{x \sim \rho, y \sim \mu_x} \frac{e^{\frac{A(x, y)}{\eta}}}{\mu_x(y)} \pi_{\text{old}}(y|x) f(c_\theta). \end{aligned}$$

Now, both objectives are written as expectations over the same uniform distribution  $\mu_x = 1/N$ , where  $N$  is the number of all possible actions. Then we can have the following subtraction,

$$\begin{aligned} \left| \mathcal{L}_{\text{LAD}}^{\text{theorem}} - \mathcal{L}_{\text{LAD}} \right| &\leq \mathbb{E}_{x \sim \rho, y \sim \mu_x} \left| \frac{e^{\frac{A(x, y)}{\eta}}}{\mu_x(y)} F(x, y) \right| \\ F(x, y) &= \frac{1}{Z_A(x)} f\left(\frac{c_\theta Z_A(x)}{Z_{\pi_\theta}(x)}\right) - \pi_{\text{old}}(y|x) f(c_\theta) \end{aligned}$$

where we use the triangle inequality here. With Assumption B.1 and Assumption B.2, for any  $(x, y)$ , we have

$$\left| \frac{e^{\frac{A(x, y)}{\eta}}}{\mu_x(y)} F(x, y) \right| \leq C \left| f\left(c_\theta \cdot \frac{Z_A(x)}{Z_{\pi_\theta}(x)}\right) - f(c_\theta) \right|$$

where  $C = e^{\frac{\gamma}{\eta}} \cdot N \cdot \max(p_{\max}, \frac{e^{\frac{\gamma}{\eta}}}{N})$  is a constant. Next, we write the Taylor expansion of  $f\left(c_\theta \cdot \frac{Z_A(x)}{Z_{\pi_\theta}(x)}\right)$  to bound the right term.

$$f\left(c_\theta \cdot \frac{Z_A(x)}{Z_{\pi_\theta}(x)}\right) = f(c_\theta) + f'(\zeta) c_\theta \left(\frac{Z_A(x)}{Z_{\pi_\theta}(x)} - 1\right),$$

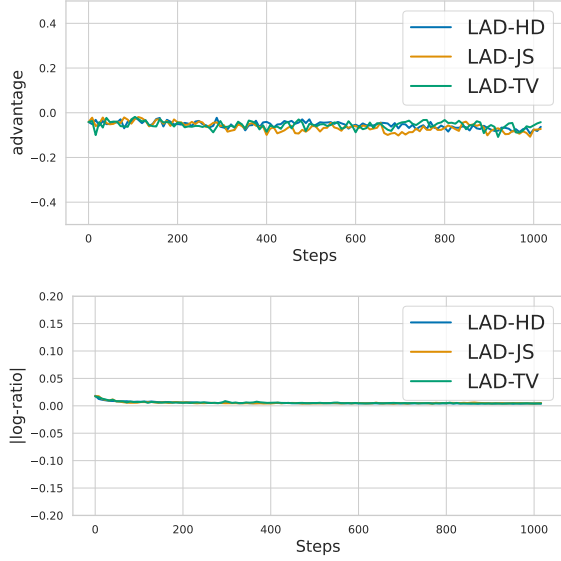


Figure 6: The average advantage value and the log-ratio of  $|\log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}|$  throughout the training.

where  $\zeta$  is a value between  $c_\theta$  and  $c_\theta \frac{Z_A(x)}{Z_{\pi_\theta}(x)}$ . Hence, for any  $(x, y)$ , we can obtain

$$\begin{aligned} & \left| \mathcal{L}_{\text{LAD}}^{\text{theorem}} - \mathcal{L}_{\text{LAD}} \right| \\ & \leq \frac{C}{N} \mathbb{E}_{x \sim \rho} \sum_y \left| f'(\zeta) c_\theta \left( \frac{Z_A(x)}{Z_{\pi_\theta}(x)} - 1 \right) \right| \end{aligned}$$

Thus, we have

$$\left| \mathcal{L}_{\text{LAD}}^{\text{theorem}} - \mathcal{L}_{\text{LAD}} \right| < \mathcal{O}(\gamma + \delta)$$

□

**The empirical validation for assumptions** To empirically validate our assumption hold true, we show the average advantage value and the log-ratio of  $|\log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}|$  during the training process of three LAD variants under strict divergence classes. As shown in Figure 6, advantage values are empirically bounded by a small value, i.e.,  $\gamma$  in Assumption B.2 is small. Meanwhile, the  $\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}$  is indeed near 1, i.e.,  $\epsilon$  in Assumption B.3 is close to 0. Therefore, the bound of  $|\mathcal{L}_{\text{LAD}}^{\text{theorem}} - \mathcal{L}_{\text{LAD}}|$  as proved in Theorem B.1 is small, which indicates that  $\mathcal{L}_{\text{LAD}}$  is indeed a close surrogate of  $\mathcal{L}_{\text{LAD}}^{\text{theorem}}$ .

#### B.4 Gradient Comparisons

Beyond the main paper’s comparison between the gradients of the LAD loss and the vanilla policy gradient, we further analyze how the LAD gradient

differs from both the PPO gradient and the theoretical LAD formulation.

In practice, PPO enforces trust-region constraints through a clipped surrogate objective. The resulting gradient can be written as

$$\mathbb{E}_{x \sim \rho, y \sim \pi_\theta(\cdot|x)} \left[ -\text{sg} \left( \frac{\pi_\theta}{\pi_{\text{old}}}, \epsilon \right) A(x, y) \nabla_\theta \log \pi_\theta(y|x) \right],$$

where  $\text{sg}(\frac{\pi_\theta}{\pi_{\text{old}}}, \epsilon)$  is a binary signal function that equals 1 when  $(A(x, y) > 0 \wedge \frac{\pi_\theta}{\pi_{\text{old}}} \leq 1 + \epsilon) \vee (A(x, y) < 0 \wedge \frac{\pi_\theta}{\pi_{\text{old}}} \geq 1 - \epsilon)$ , and 0 otherwise. Compared with the vanilla policy gradient in Eq. (1), PPO introduces an additional gating mechanism that suppresses excessively large updates. However, once the gate is active, the update direction remains proportional to the advantage. As a result, PPO still predominantly increases the likelihood of responses with larger advantages, reinforcing already high-probability actions. Although this clipping improves stability, it does not fundamentally alter the advantage-maximization nature of the objective, and therefore still tends to reduce generative diversity in a manner similar to the vanilla policy gradient. For the theoretical LAD objective, the gradient takes the following form  $\nabla_\theta \mathcal{L}_{\text{LAD}}^{\text{theorem}}$  is

$$\begin{aligned} & \mathbb{E}_{x \sim \rho, y \sim \mathcal{P}_{\pi_\theta}(\cdot|x)} f' \left( \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} \right) \cdot e^{-\frac{A(x, y)}{\eta}} \cdot \frac{Z_A(x)}{Z_{\pi_\theta}(x)} \cdot \\ & (\nabla_\theta \log \pi_\theta(y|x) - \mathbb{E}_{y' \sim \mathcal{P}_{\pi_\theta}(\cdot|x)} \nabla_\theta \log \pi_\theta(y'|x)), \end{aligned}$$

where  $\mathcal{P}_{\pi_\theta}$  is the policy-induced distribution defined in Eq. 4. Compared to the gradient of the practical LAD loss, this gradient includes an explicit centering term that subtracts the expected score function under the policy-induced distribution. This structure enforces competition among actions and prevents any single response from dominating the update solely due to a large advantage. Nevertheless, the practical LAD loss retains the core behavior of the theoretical gradient while removing intractable normalization terms. Compared with the theoretical LAD gradient, the practical LAD gradient omits the expectation-centering term but preserves the crucial dependence on both the advantage and the likelihood ratio. As a result, practical LAD continues to suppress updates for actions whose probabilities are already sufficiently large, even when their advantages remain high. This mechanism fundamentally differs from advantage maximization and enables LAD to maintain stable optimization while preserving generative diversity.

## B.5 FlowRL is a special case in our framework, but with stricter conditions

FlowRL (Zhu et al., 2025) proposes to learn a reward-matching objective by directly regressing the log policy ratio onto advantage values using a mean-squared-error loss inspired by GFlowNets (Bengio et al., 2023). In contrast, LAD provides a more general and theoretically grounded framework, within which FlowRL arises as a special case under restrictive assumptions on policy updates and parameter scaling. Specifically, under a small-step learning regime where the updated policy  $\pi_\theta$  remains close to the behavior policy  $\pi_{\text{old}}$ , the likelihood ratio  $\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)}$  stays near 1. If we further assume that this ratio satisfies  $\frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} > \frac{1}{e}$ , then the function  $f(x) = x(\log x)^2$  convex over the relevant domain  $x > \frac{1}{e}$ . The convexity ensures that the corresponding  $f$ -divergence objective is well-defined. Further, when  $\eta \rightarrow \infty$ ,  $e^{\frac{A}{\eta}} \rightarrow 1 + \frac{A}{\eta}$ . Therefore, for each  $x \sim \rho$ , the LAD loss can be formulated as

$$\mathbb{E}_{y \sim \pi_\theta(\cdot|x)} \left[ \log Z(x) + \log \frac{\pi_\theta(y|x)}{\pi_{\text{old}}(y|x)} - \frac{A(x, y)}{\eta} \right]^2, \quad (18)$$

where  $Z(x) = \frac{Z_A(x)}{e^{Z_{\pi_\theta}(x)}}$ . Eq. 18 is exactly the objective optimized by FlowRL, showing that FlowRL can be viewed as a special instantiation of LAD.

## C Experiments

### C.1 Loss Instances

All the instances involved in our experiments are listed in Table 5. It reveals that LAD loss is highly extensible to different divergence classes. As the star sign indicated, Jensen–Shannon divergence, Hellinger Distance, and Total Variance Distance are strict divergences, while KL divergence, the reverse KL divergence, and Jeffreys divergence are not strict divergences. As analyzed in §4.3, the strict divergence classes are recommended, which generally yield better overall performance.

### C.2 Implementation Details

**Math reasoning.** The experiments on math reasoning tasks are all conducted on two NVIDIA H200 140GB. We train our models on the DAPO-MATH dataset (Yu et al., 2025), following the experimental protocol of Cui et al. (2025b). We use Qwen2.5-7B (Qwen et al., 2025) as our LLM backbone, and Jensen-Shannon divergence to implement LAD loss in the main experiments. At each

rollout step, we sample 8 responses with a max length of 8192 tokens per prompt for a batch of 256 prompts using a sampling temperature of 1, and subsequently perform 8 policy update steps using the collected trajectories with the learning rate 1e-6. We apply the dynamic sampling strategy and the clip-higher technique from DAPO (Yu et al., 2025) uniformly across all baselines to ensure a fair comparison. Evaluation is conducted on a diverse set of mathematical reasoning benchmarks, including MATH500 (Lightman et al., 2023), AIME 2024, AIME 2025 (Li et al., 2024), AMC (Li et al., 2024), OlympiadBench (He et al., 2024), and Minerva (Lewkowycz et al., 2022). These settings ensure direct comparability with prior work. We follow Zhu et al. (2025) to assess logical diversity using an LLM-as-a-Judge framework with GPT-4-Turbo. Specifically, we generate 16 responses per question on AIME 2024/2025 and use the same evaluation prompt as shown in Box C.1. For the two 1.5B models, we change the max lengths to 4096 and keep other settings the same. For the ablation studies of the hyperparameter value, we kept all settings the same, except for the  $\eta$ . For the ablation studies of the divergence class, we conduct a hyperparameter search among  $1/\eta \in \{0.25, 1, 4, 16\}$ . For all baseline methods, we adopt the default hyperparameters identified as optimal through the hyperparameter searches described in the original papers.

**Code reasoning.** The experiments on math reasoning tasks are all conducted on eight NVIDIA A100 80GB. We train our models on the dataset of DeepCoder (Luo et al., 2025b), following the experimental protocol of (Zhu et al., 2025). We use DeepSeek-R1-Distill-Qwen-7B as our LLM backbone, and Jensen-Shannon divergence to implement LAD loss in the main experiments. The batch size is set to 64, and the learning rate is set to 1e-6. For all baseline methods, we adopt the default hyperparameters identified as optimal through the hyperparameter searches described in the original papers. We evaluate the code reasoning performance on LiveCodeBench (Jain et al., 2025), CodeForces (Penedo et al., 2025), and HumanEval+ (Chen et al., 2021).

**Evaluation metrics.** Pass@k is a common metric in reasoning tasks. It measures the probability that at least one correct solution appears among the top k generated outputs. Higher Pass@k means the model is more likely to produce a correct answer within its top candidates. Avg@k represents

Table 5: LAD losses under different divergences. The \* represents the corresponding divergence class is a strict divergence that forces exact matching at its minimum.

Divergence	$f(x)$	$\mathcal{L}_{\text{LAD}}$
KL divergence	$x \log x$	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} \left( \log \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} - \frac{A(x,y)}{\eta} \right)$
reverse KL divergence	$-\log x$	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} - e^{\frac{A(x,y)}{\eta}} \log \pi_{\theta}(y x)$
Jeffreys divergence	$(x-1) \log x$	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} \left( \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} - e^{\frac{A(x,y)}{\eta}} \right) \log \left[ \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} e^{-\frac{A(x,y)}{\eta}} \right]$
Total Variance Distance*	$ x-1 $	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} e^{\frac{A(x,y)}{\eta}} \left  \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} \cdot e^{-\frac{A(x,y)}{\eta}} - 1 \right $
Hellinger Distance*	$\frac{1}{2}(\sqrt{t}-1)^2$	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} e^{\frac{A(x,y)}{\eta}} \left( \sqrt{\frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)}} \cdot e^{-\frac{A(x,y)}{2\eta}} - 1 \right)^2$
Jensen–Shannon divergence*	$\frac{1}{2}(x \log x - (x+1) \log \frac{x+1}{2})$	$\mathbb{E}_{x \sim \rho, y \sim \pi_{\text{old}}(\cdot x)} e^{\frac{A(x,y)}{\eta}} \left( \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} e^{-\frac{A(x,y)}{\eta}} \log \left( \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} \right) - e^{-\frac{A(x,y)}{\eta}} \right) - \left( \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} e^{\frac{A(x,y)}{\eta}} + 1 \right) \log \left( \frac{\pi_{\theta}(y x)}{\pi_{\text{old}}(y x)} e^{\frac{A(x,y)}{\eta}} \right)$

Table 6: The detailed math reasoning performances of LAD with different hyperparameters.

Methods	$1/\eta$	MATH	Olympiad Bench	Minerva	AIME 2024	AIME 2025	AMC	Average
GRPO	-	75.71	39.24	36.63	Avg@16 13.89		53.39	37.29
LAD-JS	1/4	75.24	38.83	35.59	15.68	5.70	50.34	36.90
	1/2	76.20	38.02	38.07	18.15	6.73	53.08	38.38
	1	74.84	37.16	39.84	14.95	11.28	55.98	39.00
	2	76.90	36.50	34.93	18.67	9.13	55.12	38.54
	4	77.60	41.29	36.81	19.59	8.71	56.37	40.06
	8	77.15	42.06	35.64	14.54	12.11	49.34	38.47
	16	74.89	39.30	36.01	18.54	12.40	51.60	38.79
32	76.47	40.02	34.74	14.15	9.33	51.08	37.63	
GRPO	-	89.80	57.78	53.31	Pass@16 27.60		74.92	54.44
LAD-JS	1/4	92.20	64.44	55.88	34.38	25.83	79.82	58.76
	1/2	91.60	61.93	58.46	35.63	26.67	80.16	59.08
	1	87.4	53.78	55.15	35.42	32.92	78.65	57.22
	2	91.40	57.33	57.35	40.42	30.42	80.91	59.64
	4	92.40	64.15	56.62	40.63	30.73	84.30	61.47
	8	93.00	68.15	52.57	34.58	31.04	79.78	59.85
	16	91.60	65.33	56.99	37.50	29.79	80.08	60.22
32	94.40	65.63	57.72	38.13	29.79	81.36	61.17	

Table 7: Pass@32 performance of different losses on math reasoning benchmarks. The best two results under a specific metric are highlighted in **bold** and underlined.

Methods	MATH	Olympiad Bench	Minerva	AIME 2024	AIME 2025	AMC	Average
GRPO(Shao et al., 2024)	90.60	60.15	55.88	30.52	26.77	80.50	57.40
EntAdv(Cheng et al., 2025)	90.60	62.51	55.51	35.83	34.27	81.93	60.10
KLCov(Cui et al., 2025b)	92.80	64.15	56.99	36.25	<b>41.25</b>	<u>85.99</u>	63.04
ClipCov(Cui et al., 2025b)	91.40	61.48	<b>60.29</b>	37.40	32.08	80.61	60.54
FlowRL(Zhu et al., 2025)	<b>94.20</b>	<b>69.04</b>	59.56	<u>37.50</u>	35.52	82.94	<u>63.13</u>
LAD	<u>93.80</u>	<u>67.70</u>	<u>59.93</u>	<b>47.19</b>	<u>37.08</u>	<b>88.86</b>	<b>65.76</b>

the average accuracy over the top k generated outputs. Unlike Pass@k, which is binary (pass/fail), Avg@k captures overall quality across multiple candidates. Lexical diversity measures the variety of vocabulary used in a text. Higher lexical diver-

sity indicates less repetition and a broader range of words, often used to assess richness or creativity in language generation. Logical Diversity evaluates the variety of reasoning patterns, argument structures, or logical paths in generated outputs. Higher

logical diversity suggests the model can approach a problem from multiple valid reasoning perspectives rather than repeating the same logic.

### C.3 Experimental Results

**Supplemented results of the main experiments on math reasoning.** As shown in Table 7, for the evaluation metric Pass@32. Across six benchmarks, LAD achieves the best average performance (65.76), outperforming the second-best method (63.13).

**Supplemented results of the ablation studies on divergence classes.** As shown in Table 6, the performances of LAD variants with stricter divergence classes are generally better than ones with non-stricter divergence classes.

**Supplemented results of the ablation studies on the hyperparameter  $\eta$ .** As shown in Table 6, LAD can perform well in a large range of  $\eta$  values. Notably, across nearly all settings, LAD consistently outperforms GRPO, indicating that its effectiveness is not critically dependent on precise tuning of  $\eta$ .

Table 8: Performance of different losses on math reasoning benchmarks. The best two results under a specific metric are highlighted in **bold** and underlined.

Methods	$k$	MATH	Olympiad Bench	Minerva	AIME 2024	AIME 2025	AMC	Average
		Avg@ $k$						
LAD-KL	@1	76.80	37.33	38.97	20.00	5.83	51.13	38.01
	@16	76.35	36.93	39.18	20.03	4.64	51.74	38.15
	@32	76.20	36.83	38.73	19.91	4.68	51.63	37.99
LAD-JF	@1	72.60	32.74	37.13	15.52	7.50	53.69	36.20
	@16	72.65	32.54	34.92	15.96	7.23	53.47	36.13
	@32	72.83	32.67	34.60	15.79	7.24	53.48	36.10
LAD-rKL	@1	73.20	37.62	37.50	11.67	8.33	46.15	35.41
	@16	73.86	36.63	36.28	12.21	8.03	45.93	35.49
	@32	73.82	36.76	36.12	12.25	8.10	46.08	35.52
LAD-TV	@1	<u>76.80</u>	<u>40.59</u>	35.66	15.31	6.25	55.61	38.37
	@16	<u>77.00</u>	<u>40.92</u>	<b>38.14</b>	14.98	7.58	55.27	38.98
	@32	76.83	40.66	<b>37.74</b>	14.96	<u>7.61</u>	55.19	38.83
LAD-HD	@1	<b>77.60</b>	39.70	35.29	<u>15.63</u>	<u>7.92</u>	<b>57.34</b>	<u>38.91</u>
	@16	76.88	39.06	<u>36.83</u>	<u>16.27</u>	<u>8.08</u>	<b>57.40</b>	<u>39.08</u>
	@32	76.83	38.98	<u>36.74</u>	<u>16.18</u>	7.29	<b>57.40</b>	<u>38.90</u>
LAD-JS	@1	<u>76.80</u>	40.30	<u>36.40</u>	<b>19.38</b>	<b>8.65</b>	<u>57.00</u>	<b>39.76</b>
	@16	<b>77.60</b>	<b>41.29</b>	36.81	<b>19.59</b>	<b>8.71</b>	<u>56.37</u>	<b>40.06</b>
	@32	<b>77.66</b>	<b>41.25</b>	36.68	<b>19.60</b>	<b>8.84</b>	<u>56.45</u>	<b>40.08</b>
		Pass@ $k$						
LAD-KL	@1	86.40	51.11	56.25	34.48	17.92	66.19	52.39
	@16	87.80	55.26	60.29	39.69	22.50	71.16	56.12
	@32	88.60	57.19	61.40	43.85	26.46	76.02	58.92
LAD-JF	@8	88.2	51.56	54.04	32.50	21.25	77.18	54.12
	@16	90.4	56.74	58.09	38.85	26.98	82.38	58.90
	@32	92.2	60.59	59.93	46.04	33.85	86.71	63.22
LAD-rKL	@8	89.60	58.67	53.68	29.06	25.41	72.29	54.78
	@16	91.20	64.30	59.93	34.90	30.83	79.97	60.19
	@32	93.60	68.89	62.87	43.75	37.50	86.07	65.45
LAD-TV	@8	87.60	55.70	<b>55.51</b>	26.35	<b>26.04</b>	73.53	54.12
	@16	90.20	58.96	<b>60.29</b>	29.79	<u>31.15</u>	77.67	58.01
	@32	91.60	62.07	<b>62.50</b>	34.06	<u>35.00</u>	81.36	61.10
LAD-HD	@8	88.20	56.59	<u>54.41</u>	31.04	<u>25.42</u>	75.00	<u>55.11</u>
	@16	91.00	60.89	<u>59.93</u>	<u>36.04</u>	<b>31.67</b>	<u>79.71</u>	<u>59.87</u>
	@32	93.00	64.74	<b>62.50</b>	<u>42.71</u>	<b>38.23</b>	<u>84.22</u>	<u>64.23</u>
LAD-JS	@8	<u>89.00</u>	<u>59.56</u>	54.04	<b>35.00</b>	<u>25.42</u>	<b>78.58</b>	<b>56.93</b>
	@16	<u>92.40</u>	<u>64.15</u>	56.62	<b>40.63</b>	30.73	<b>84.30</b>	<b>61.47</b>
	@32	<u>93.80</u>	<u>67.70</u>	59.93	<b>47.19</b>	<u>37.08</u>	<b>88.86</b>	<b>65.76</b>

### Box C.0: GPT-4 prompt for diversity judgement

<System Prompt>

You are evaluating the DIVERSITY of solution approaches for a mathematics competition problem. Focus on detecting even SUBTLE differences in methodology that indicate different problem-solving strategies.

<User Prompt>

PROBLEM:

{problem}

16 SOLUTION ATTEMPTS:

{formatted responses}

EVALUATION CRITERIA - Rate diversity from 1 to 5:

Score 1 - Minimal Diversity:

- 14+ responses use essentially identical approaches
- Same mathematical setup, same variable choices, same solution path
- Only trivial differences (arithmetic, notation, wording)
- Indicates very low exploration/diversity in the generation process

Score 2 - Low Diversity:

- 11-13 responses use the same main approach
- 1-2 alternative approaches appear but are rare
- Minor variations within the dominant method (different substitutions, orderings)
- Some exploration but heavily biased toward one strategy

Score 3 - Moderate Diversity:

- 7-10 responses use the most common approach
- 2-3 distinct alternative approaches present
- Noticeable variation in problem setup or mathematical techniques
- Balanced mix showing reasonable exploration

Score 4 - High Diversity:

- 4-6 responses use the most common approach
- 3-4 distinct solution strategies well-represented
- Multiple mathematical techniques and problem framings
- Strong evidence of diverse exploration strategies

Score 5 - Maximum Diversity:

- No single approach dominates ( $\leq 3$  responses use same method)
- 4+ distinctly different solution strategies
- Wide variety of mathematical techniques and creative approaches
- Excellent exploration and generation diversity

IMPORTANT: Focusing on the DIVERSITY of the attempted approaches. Return ONLY a number from 1 to 5.