

BASIS: Batchwise Advantage Estimation from Single-Rollout Information Sharing for LLM Reasoning

Shijin Gong^{1,*} Erhan Xu^{2,*} Kai Ye^{2,*}
Giulia Livieri^{2,†} Francesco Quinzan^{3,†} Chengchun Shi^{2,†}

Abstract

Reinforcement learning with verifiable rewards has become a standard recipe for improving the reasoning abilities of large language models. Existing algorithms face a tradeoff between computational efficiency and sample efficiency in value estimation and policy learning. We introduce BASIS, a critic-free post-training algorithm designed to address this tradeoff. At each online training step, BASIS samples only one rollout per prompt, but leverages rich information across prompts in the entire batch to improve value function estimation. Our experiments demonstrate that BASIS reduces MSE in value function estimation by **69%** compared to REINFORCE++, a representative single-rollout baseline, and achieves lower MSE with one rollout than GRPO estimators with **8** rollouts. This improvement in value estimation translates to better policy optimization: using substantially less online training time, BASIS achieves performance close to multi-rollout GRPO-type baselines and often outperforms single-rollout REINFORCE-type baselines.

1 Introduction

Recent progress in large language model (LLM) reasoning has increasingly relied on large-scale reinforcement learning (RL). In particular, RL with verifiable rewards (RLVR, [Lambert et al., 2024](#)) scores LLM responses using automated task verifiers and uses these scores as rewards to update the model’s policy. Popularized by DeepSeekMath and DeepSeek-R1 ([Shao et al., 2024](#); [DeepSeek-AI](#)

[et al., 2025](#)), RLVR has become a standard post-training framework for open-source large reasoning models (see e.g., [Liu et al., 2025a](#); [Yu et al., 2025](#); [Zheng et al., 2025](#)), with numerous follow-up works appearing shortly thereafter (Section 2).

Existing RLVR algorithms face a fundamental tradeoff between computational efficiency and sample efficiency. More accurate value or advantage estimates can improve the sample efficiency of the resulting policy optimization algorithm, but they are often expensive to obtain. PPO-type algorithms train an auxiliary critic network for value estimation ([Ouyang et al., 2022](#)), while GRPO-type algorithms sample multiple rollouts per prompt and use their average reward for the estimation ([Shao et al., 2024](#); [Liu et al., 2025b](#); [Chu et al., 2025](#); [Zheng et al., 2025](#)). These algorithms improve sample efficiency at the cost of additional computation. In contrast, single-rollout critic-free algorithms are computationally cheaper, but their value estimates are less accurate, leading to noisier policy updates.

In response to this tradeoff, we propose BASIS, short for **B**atchwise **A**dvantage estimation from **S**ingle-rollout **I**nformation **S**haring, a critic-free RLVR algorithm that samples only one rollout per prompt at each online training step.

- **Methodologically**, BASIS introduces a novel batchwise advantage estimator that borrows rich information across prompts in the whole batch to improve value and advantage estimation (see [Figure 1](#) for an illustration). Together with offline estimation and online calibration ([Section 4](#)), this enables BASIS to retain much of the sample efficiency of multi-rollout algorithms while using only a single rollout during online training.
- **For value estimation**, BASIS is highly sample efficient: it substantially reduces the MSE of single-rollout baselines and achieves lower MSE than multi-rollout baselines using **8** rollouts. Moreover, unlike existing baselines, BA-

*These authors contributed equally and are listed in alphabetical order.

†Joint senior authors and corresponding authors. Emails: g.livieri@lse.ac.uk, francesco.quinzan@eng.ox.ac.uk, c.shi7@lse.ac.uk

¹School of Management, University of Science and Technology of China.

²Department of Statistics, London School of Economics and Political Science.

³Department of Engineering Science, University of Oxford.

SIS is robust to reward heterogeneity and prompt difficulty while producing informative advantage estimates (Section 5.1).

- **For policy optimization**, the BASIS advantage estimator can serve as a plug-in component for a range of RLVR algorithms. It stabilizes training and mitigates collapse during optimization. The resulting fine-tuned model often outperforms single-rollout baselines and remains competitive with multi-rollout baselines while requiring considerably less training time (Section 5.2).

2 Related Work

The algorithmic foundations of RLVR trace back to a long line of policy gradient algorithms in the RL literature (Sutton et al., 1999; Williams, 1992). One central theme in this literature is variance reduction: more accurate value or advantage estimates can reduce the variance of policy gradient estimation and improve the quality of policy optimization (Greensmith et al., 2004). Motivated by this observation, existing approaches can be divided into four categories:

1. The first category, represented by PPO and SAO (Hou et al., 2026), employs a deep neural network to learn the value function and applies generalized advantage estimation for variance reduction (Schulman et al., 2016, 2017). While effective, this approach introduces an additional value network, making training substantially more costly.
2. The second category, represented by GRPO, removes the value network and uses the average reward over multiple rollouts of the same prompt as a proxy for the value function (Shao et al., 2024; Ahmadian et al., 2024).
3. The third category follows the critic-free spirit of the second category, but aims to reduce the number of rollouts required per prompt. One approach uses the reward of a greedy, deterministic rollout, which only requires one rollout per prompt (Li et al., 2023). Another approach reduces the required number of rollouts while maintaining sample efficiency by borrowing reward information either across training iterations for the same prompt (Wang et al., 2025; Xu and Ding, 2025; Gong et al., 2026; Ruan et al., 2026), or across different prompts in the same batch (Hu et al., 2025; Zeng et al., 2025; Han et al., 2026).
4. The last category studies variance reduction baselines beyond standard value functions (Hao et al., 2025).

BASIS is most closely related to the third category, especially those methods that exploit batch-level reward information. However, as our experiments show, BASIS achieves substantially smaller MSE in value estimation than REINFORCE++ (Hu et al., 2025), a representative batch-level baseline.

Beyond value and advantage estimation, existing work also studies exploration, clipping, rejection sampling, experience reuse, rollout down-sampling and pruning, uncertainty and difficulty-aware updates (Chen et al., 2025; Dai et al., 2025; Lin et al., 2025; Shrivastava et al., 2025; Su et al., 2025; Xiong et al., 2025; Xu et al., 2025; Zhang and Zuo, 2025; Zhan et al., 2025; Cheng et al., 2026). Since BASIS only modifies the value function baseline, it can serve as a plug-in component for many of these RLVR pipelines.

Finally, recent work has begun to study the theoretical properties of RLVR (Pang et al., 2025; Davis and Recht, 2025; Vojnovic and Yun, 2025; Yang et al., 2026; Huang et al., 2026; Zhou et al., 2026).

3 Preliminaries

We first introduce notation that will be used throughout the paper. We then describe how the classical REINFORCE algorithm (Williams, 1992) applies to LLM post-training. Finally, we use PPO, GRPO, and REINFORCE++ as representative variance reduction baselines to illustrate the first three categories of algorithms reviewed in Section 2.

Notation. We use $x \sim \mathcal{D}$ to denote a prompt drawn from a training dataset $\mathcal{D}$, and $y \sim \pi_\theta(\bullet|x)$ to denote a rollout sampled from the LLM policy π_θ that we wish to fine-tune, with parameters θ . For reasoning tasks, y contains both a reasoning trace and a final answer. An automated task verifier returns a scalar reward $r(x, y) \in \mathbb{R}$, where the reward function r usually measures the correctness of the final answer. Let $\mathcal{B} = \{x_i\}_{i=1}^B$ denote a batch of prompts. To simplify the notation, we use r_i to denote $r(x_i, y_i)$ for a given prompt-rollout pair (x_i, y_i) . We use θ_t to denote the model parameters at training iteration t , and define $V_{i,t} = \mathbb{E}_{y \sim \pi_{\theta_t}(\bullet|x_i)}[r(x_i, y)]$, $A_{i,t} = r_i - V_{i,t}$ as the policy’s value and advantage, respectively.




Figure 1: Overview of BASIS for constructing advantage estimates. BASIS samples one completion per prompt and uses reward information from the whole batch to estimate value baselines. The baselines are weighted averages of batch rewards, with the weight matrix W shown in the center. The diagonal entries of W are zero, shown in white, enforcing a prompt’s own reward to be excluded when estimating its baseline. Unlike single-rollout baselines such as REINFORCE++, which use a simple global average, BASIS chooses W by minimizing the MSE of the value baselines. Advantages are then computed by subtracting the baselines from the observed rewards.

REINFORCE. RLVR algorithms fine-tune the parameters θ by maximizing the expected reward

$$\mathcal{J}(\theta) = \mathbb{E}_{x \sim \mathcal{D}} \left\{ \mathbb{E}_{y \sim \pi_{\theta}(\bullet|x)} [r(x, y)] \right\}, \quad (1)$$

which can be optimized using stochastic gradient methods (Robbins and Monro, 1951). Specifically, its gradient $\nabla_{\theta} \mathcal{J}(\theta)$ can be shown to equal

$$\mathbb{E}_{x \sim \mathcal{D}} \left\{ \mathbb{E}_{y \sim \pi_{\theta}(\bullet|x)} [r(x, y) \nabla_{\theta} \log \pi_{\theta}(y|x)] \right\}.$$

This identity motivates REINFORCE-type algorithms: starting from an initial parameter θ_1 , at each training step t , the algorithm samples a batch of prompts $\mathcal{B} = \{x_i\}_{i=1}^B$ and one rollout $y_i \sim \pi_{\theta_t}(\bullet|x_i)$ per prompt. It then estimates $\nabla_{\theta} \mathcal{J}(\theta_t)$ by averaging the product of the reward r_i and policy score $\nabla_{\theta} \log \pi_{\theta_t}(y_i|x_i)$ over the batch, and applies stochastic gradient ascent to update θ_t to θ_{t+1} .

Variance reduction baselines. The REINFORCE policy gradient estimator is unbiased, but often suffers from high variance, especially when the model is uncertain and sampled rollouts can vary substantially in quality. This variability leads to noisy rewards and gradient estimates, making policy optimization less stable. A standard remedy is to construct an advantage estimate $r_i - b_i$ by subtracting a baseline b_i from the reward r_i before multiplying by the policy score. Under suitable conditions, the optimal baseline that minimizes the variance of the policy gradient estimator is the value function (Greensmith et al., 2004).

PPO, GRPO, and REINFORCE++ all adopt this variance reduction principle, but differ in how they construct b_i , and hence the resulting advantage estimate. PPO sets the baseline to an estimated value function learned by an auxiliary neural network,

which improves sample efficiency but introduces substantial computational cost. GRPO sets b_i to the average reward over multiple rollouts from the same prompt, which eliminates the need for a value network but also increases computation through repeated rollout sampling. REINFORCE++ instead uses a global batch baseline, given by a simple average of $\{r_i\}_i$, which avoids both a value network and multiple rollouts per prompt, but the resulting baseline is shared across all prompts and may poorly approximate each prompt’s individual value.

These choices highlight the central tradeoff addressed in this paper: accurate baseline estimation can be obtained by learning a value model or by repeating rollouts for each prompt, but both approaches increase computation.

4 BASIS

We now introduce BASIS to address the tradeoff between accurate baseline estimation and computational efficiency. At each online training step, BASIS samples only one rollout per prompt while leveraging the rich information available across the entire training batch to improve each prompt’s value and advantage estimation.

BASIS proceeds in three steps: offline value estimation, batchwise refinement, and online calibration. It first constructs an initial prompt-level value estimate through offline estimation. It then refines this estimate online using reward information from the entire training batch. Finally, it calibrates the refined estimator at each training step. Algorithm 1 summarizes the procedure for each training step. Below, we first describe the second step, which is the key ingredient of BASIS. We then explain how the offline value estimates are computed, and how the refined estimators are calibrated.

Batchwise refinement. At each training step t , given a prompt-rollout batch $\{(x_i, y_i)\}_{i=1}^B$, BASIS estimates the value $V_{i,t}$ for each x_i . Its main idea is simple: each $V_{i,t}$ is estimated by a weighted average of rewards from other prompts in the batch:

$$\tilde{V}_{i,t} = \sum_{j \neq i} w_{ij} r_j, \quad (2)$$

where the weights $\{w_{ij}\}_{j \neq i}$ are prompt-dependent. Such a linear combination allows BASIS to borrow information across prompts in the same batch, yielding a low-MSE estimate from only one rollout per prompt, as we show in the experiments.

Before specifying the weights w_{ij} , we highlight two features of this estimator. (i) It is prompt-dependent: the weights vary with i , unlike the global batch baseline used by REINFORCE++. With a proper choice of weights, this allows BASIS to better approximate the prompt-level value. (ii) It follows the leave-one-out principle of RLOO (Ahmadian et al., 2024): the reward r_i of the target prompt is excluded from its own baseline.

It remains to specify the weights w_{ij} . The following proposition motivates our choice.

Proposition 1 (Best linear unbiased estimator (BLUE)). *For each prompt i , among all linear estimators of the form (2) that are unbiased for $V_{i,t}$, i.e., $\mathbb{E}(\tilde{V}_{i,t}|\mathcal{B}) = V_{i,t}$, the weights minimizing $\mathbb{E}[(V_{i,t} - \tilde{V}_{i,t})^2|\mathcal{B}]$ are*

$$w_{ij} = \frac{V_{i,t} V_{j,t} / \sigma_j^2}{\sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}, \quad \forall j \in \{1, \dots, B\} \setminus \{i\}, \quad (3)$$

where $\sigma_j^2 = \text{Var}(r_j|x_j)$.

Proposition 1 characterizes the ideal weights one would use to minimize the MSE of the resulting value estimator while preserving its unbiasedness. Intuitively, when estimating the value of the i th prompt, the weight increases with both the target value $V_{i,t}$ and the source value $V_{j,t}$, and decreases with the source reward variance σ_j^2 .

Of course, the BLUE weights cannot be used directly during online training, because the true values $V_{i,t}$ and variances σ_i^2 are unknown. BASIS therefore uses this proposition as a guiding principle. For binary rewards, given an initial value estimate $\hat{V}_{i,t}$, we estimate the reward variance by $\hat{\sigma}_i^2 = \hat{V}_{i,t}(1 - \hat{V}_{i,t})$. We then plug these estimates into (3) to construct the refined estimator $\tilde{V}_{i,t}$.

When some initial values are close to zero or one, their Bernoulli variance estimates are close to zero,

making the weights in (3) numerically unstable. To address this, we define the following active set

$$\mathcal{A}_t = \{i : \epsilon < \hat{V}_{i,t} < 1 - \epsilon\},$$

for some small threshold $\epsilon > 0$. For prompts in the active set, the sums in (2) and the denominator of (3) are taken over $\mathcal{A}_t \setminus \{i\}$ rather than all prompts. For prompts outside the active set, BASIS falls back to the conservative zero baseline $b_i = 0$.

We will describe how the initial value estimates are obtained below. To conclude this batchwise refinement step, we note that these initial estimates could also be used directly for advantage estimation without refinement. However, the batchwise refinement yields a substantially more accurate estimator by leveraging information shared across prompts. Figure 2 illustrates this effect by comparing the MSEs of the initial and refined estimators at three checkpoints. It can be seen that the refined estimator achieves considerably lower MSE and is much less sensitive to a calibration hyperparameter.

Offline value estimation. We next describe how BASIS computes the initial value estimates $\hat{V}_{i,t}$. The idea is motivated by direct preference optimization (Rafailov et al., 2023) and its extension to LLM reasoning (Brantley et al., 2025). To begin with, consider the following Kullback–Leibler (KL)-regularized objective function,

$$\mathbb{E}_{x \sim \mathcal{D}} \mathbb{E}_{y \sim \pi(\bullet|x)} [r(x, y) - \beta \text{KL}(\pi(\bullet|x) \parallel \pi_{\text{ref}}(\bullet|x))],$$

where $\pi_{\text{ref}} = \pi_{\theta_1}$ denotes the reference policy, taken to be the initial policy before fine-tuning, and $\beta > 0$ controls the strength of the KL penalty.

This objective is particularly appealing because its maximizer π_β^* admits the following closed-form expression (Rafailov et al., 2023):

$$\pi_\beta^*(y|x) = \frac{\pi_{\text{ref}}(y|x)}{Z_\beta(x)} \exp\left(\frac{r(x, y)}{\beta}\right),$$

where $Z_\beta(x)$ is a normalizing constant. This expression connects π_β^* directly to π_{ref} . As a result, estimating the value of π_β^* does not require training or sampling from π_β^* itself. Instead, it can be done offline by sampling rollouts from π_{ref} (Brantley et al., 2025). We formalize this observation in the following proposition.

Proposition 2 (Closed-form value under π_β^*). *For any prompt x and $\beta > 0$, the value function under*




Figure 2: MSEs of the initial and refined estimators at three model checkpoints. The left panel reports MSE at an early training stage, while the middle and right panels report MSE at later stages. Both estimators depend on a hyperparameter β , and the curves visualize their MSE across different values of β .

Algorithm 1 One online training step of BASIS

Require: Policy π_{θ_t} , batch $\mathcal{B} = \{x_i\}_{i=1}^B$, search interval $\mathbb{B}$, cached reward statistics for evaluating $\hat{V}_\beta(x_i)$, threshold ϵ

- 1: Set $\theta_{\text{old}} \leftarrow \theta_t$
- 2: **for** $i = 1, \dots, B$ **do**
- 3: Sample $y_i \sim \pi_{\theta_{\text{old}}}(\bullet|x_i)$ and compute r_i
- 4: Set $\hat{V}_{i,t,\beta} \leftarrow \hat{V}_\beta(x_i)$, $\mathcal{A}_{t,\beta} \leftarrow \{i : \epsilon < \hat{V}_{i,t,\beta} < 1 - \epsilon\}$, compute $\tilde{V}_{i,t,\beta}$ by Eqs. (2) and (3) for $\beta \in \mathbb{B}$
- 5: **end for**
- 6: Select β_t by Eq. (5)
- 7: **for** $i = 1, \dots, B$ **do**
- 8: Set $b_i \leftarrow \tilde{V}_{i,t,\beta_t} \mathbb{1}_{\{i \in \mathcal{A}_{t,\beta_t}\}}$ and $\hat{A}_i \leftarrow r_i - b_i$
- 9: **end for**
- 10: Perform the policy update using $\{\hat{A}_i\}_{i=1}^B$

π_β^* equals

$$V_\beta^*(x) := \mathbb{E}_{y \sim \pi_\beta^*(\bullet|x)}[r(x, y)] \\ = \frac{\mathbb{E}_{y \sim \pi_{\text{ref}}(\bullet|x)}[r(x, y) \exp(r(x, y)/\beta)]}{\mathbb{E}_{y \sim \pi_{\text{ref}}(\bullet|x)}[\exp(r(x, y)/\beta)]}. \quad (4)$$

The expectations in both the numerator and denominator of (4) are taken with respect to the reference policy π_{ref} . This makes it feasible to estimate $V_\beta^*(x)$ offline, without training or sampling from π_β^* itself. Specifically, for each prompt, we sample a set of reference rollouts from π_{ref} and score them using the same automated verifier. We then compute $\hat{V}_\beta(x)$, a plug-in estimate of (4) by replacing the expectations in the numerator and denominator with their empirical averages over these reference rollouts. These offline estimates are then used as the initial value estimates in BASIS. Below, we describe how the hyperparameter β is adaptively selected across training steps.

Online calibration. To illustrate how the KL-regularized optimal values V_β^* can be used to ap-

proximate the values $V_{i,t}$ across training steps, consider two extreme cases:

1. At the first training step, the value $V_{i,t}$ is computed under the initial policy, which is essentially the reference policy. This corresponds to setting β to infinity: the KL penalty becomes dominant, and the induced optimal policy π_β^* coincides with the reference policy.
2. When the algorithm converges, for sufficiently large t , $V_{i,t}$ approaches the value under the reward-optimal policy that maximizes (1). This corresponds to setting β close to zero: the KL penalty vanishes, and the induced optimal policy π_β^* approaches the reward-optimal policy.

More generally, the KL-regularized optimal values provide a continuum of value estimates that approximate how the policy value evolves from the reference policy toward a reward-optimal policy during training. Motivated by this observation, for an arbitrary training step t and prompt x , we use the family of values $\{V_\beta^*(x) : \beta \in \mathbb{B}\}$ to approximate the value of the current learning policy $V_{i,t}$.

In our implementation, we set $\mathbb{B} = [0.01, 5]$. The calibration parameter β is selected adaptively at each training step, allowing its value to vary as training progresses. Specifically, at each training step t , we select β by minimizing

$$\beta_t \in \arg \min_{\beta \in \mathbb{B}} \frac{1}{|\mathcal{A}_{t,\beta}|} \sum_{i \in \mathcal{A}_{t,\beta}} (r_i^{(t)} - \tilde{V}_{i,t,\beta})^2, \quad (5)$$

where $\tilde{V}_{i,t,\beta}$ and $\mathcal{A}_{t,\beta}$ denote, respectively, the batchwise refined estimator and the active set obtained by using $\hat{V}_\beta(x_i)$ as the initial value estimate.

Figure 8 in Appendix D visualizes the selected β_t . As expected, it generally decreases over training. Finally, we set the baseline $b_i = \tilde{V}_{i,t,\beta_t}$ and compute the advantage as $\hat{A}_{i,t} = r_i - b_i$.

5 Experiments

We conduct numerical experiments to evaluate BASIS along two dimensions: value estimation (Section 5.1), where we measure the accuracy of the estimated values, and policy optimization (Section 5.2), where we assess the downstream performance of the fine-tuned models.

5.1 Value estimation

We begin with a high-level summary of our findings and provide details in the following paragraphs. BASIS is highly sample efficient and robust for value estimation:

1. It reduces the REINFORCE++ value estimator’s MSE by **69%**, and with only a single rollout per prompt, achieves smaller MSE than GRPO-type estimators with **8** rollouts (Figure 3(a)).
2. Its estimator is also robust to reward heterogeneity within a batch, unlike the global baseline used by REINFORCE++ (Figure 3(b)), and remains robust across prompt difficulty levels, unlike GRPO-type algorithms (Figure 3(c)).
3. Moreover, unlike GRPO, BASIS often produces an informative nonzero advantage (Figure 3(d)).

Setup. We post-train Qwen2.5-Math-7B with GRPO on a subset of the MATH dataset consisting of Level 3–5 problems (Hendrycks et al., 2021). We then use an intermediate training checkpoint to evaluate the accuracy of different value estimators.

Specifically, for each prompt, we first approximate its oracle value by averaging rewards over 256 Monte Carlo rollouts from the checkpoint. We then repeat the following procedure: In each repeat, we independently sample a batch of prompts from the same checkpoint, compute the BASIS, GRPO, RLOO and REINFORCE++ value estimators, and compute their squared errors against the Monte Carlo oracle values. We then average these squared errors over different repeats to obtain the MSE. More details are provided in Appendix A.

Finding 1: BASIS is sample efficient. Figure 3(a) reports the MSE of different value estimators averaged over all prompts. Compared with single-rollout baselines such as REINFORCE++, BASIS reduces the MSE by **69%**. More strikingly, even with one rollout per prompt, BASIS achieves lower MSE than GRPO and RLOO with **8** rollouts. This demonstrates the utility of batchwise learning: by borrowing information across the entire

batch, BASIS substantially improves value estimation without requiring repeated rollouts from the same prompt. Additionally, RLOO has larger MSE than GRPO, since its leave-one-out construction uses one fewer rollout for value estimation.

Finding 2: BASIS is robust. Figure 3(b) reports the MSE of different value estimators across data batches grouped by the heterogeneity of their prompt-level values, measured by the standard deviation of $\{V_{i,t}\}_{i=1}^B$. Figure 3(c) reports the MSE across subsets of prompts grouped by their values.

We make three observations. (i) First, BASIS is nearly flat across all five groups in both figures, demonstrating its robustness to both within-batch reward heterogeneity and prompt difficulty. (ii) Second, REINFORCE++ is sensitive to both factors. As batch-level heterogeneity increases from **0.21** to **0.31**, its MSE nearly doubles. This is expected because the variance of the global batch baseline used by REINFORCE++ increases naturally with the within-batch reward heterogeneity. Similarly, across prompt difficulty levels, REINFORCE++ performs best on medium-difficulty prompts but incurs much larger MSEs on easy and hard prompts, since the global batch mean pulls all baseline estimates toward the middle. (iii) Finally, GRPO and RLOO perform worst on medium-difficulty prompts. This is also expected for binary rewards: prompts with values near the middle have the largest reward variance.

Finding 3: BASIS is informative. Figure 3(d) reports how often the estimated baseline is exactly 0 or 1, leading to a zero advantage and therefore a zero contribution to the policy gradient estimator. BASIS never returns a 0 or 1 baseline in this setting, and therefore consistently provides an informative signal. In contrast, GRPO often produces such extreme baselines, particularly for easy and hard prompts.

5.2 Policy optimization

We now show that improved value and advantage estimation leads to more effective policy optimization. In particular, our experiments demonstrate three major findings:

1. The BASIS advantage estimator can be used as a versatile plug-in component for a number of multi-rollout GRPO-type algorithms, roughly **halving** online learning time while preserving most of their downstream performance (Table 1).




Figure 3: MSE of various value estimates. (a) MSE aggregated over all prompts. (b) MSE across batches grouped by the heterogeneity of within-group values, measured by the standard deviation of $\{V_{i,t}\}_{i=1}^B$. (c) MSE across prompts grouped by prompt-level value. (d) Frequency with which the estimated baseline is exactly 0 or 1. G in the legend denotes the number of rollouts per prompt. In (b), we report GRPO and RLOO only with 8 rollouts, since their MSEs are much larger with fewer rollouts. In (c) and (d), only GRPO results are reported to keep the figure readable. MSEs of RLOO value estimators are deferred to Figure 4 in Appendix A. In (d), the yellow line (BASIS) and the red line (REINFORCE++) largely overlap.

2. Compared with single-rollout REINFORCE-type baselines, BASIS often achieves better downstream accuracy, with absolute improvements of up to **44.8** and **9.5** percentage points over REINFORCE and REINFORCE++, respectively, while using roughly **one-half** of their computation (Table 1).
3. REINFORCE can **collapse** during training: its downstream accuracy may drop considerably as the number of training iterations increases. BASIS shows a much more stable learning curve, suggesting that more accurate value and advantage estimation stabilizes single-rollout policy optimization and helps avoid such collapse (see Figure 6 in Appendix C).

Setup. We compare different algorithms by using them to post-train Qwen3-4B on the DAPO-Math-17K training split (Yu et al., 2025) and evaluating the resulting models on seven benchmarks: AIME 2024, AIME 2025, AMC 2023, MATH-500 (Hendrycks et al., 2021), Minerva Math (Lewkowycz et al., 2022), OlympiadBench (He et al., 2024), and HMMT 2025 (Balunovic et al., 2025). For AIME, AMC, and HMMT, we report avg@32, the average accuracy of the resulting model over 32 sampled rollouts per problem. For MATH-500, Minerva Math, and OlympiadBench, we report accuracy from a single greedy deterministic rollout per problem.

We also evaluate the models on five out-of-domain (OOD) benchmarks: HumanEval+ and MBPP+ for code generation (Liu et al., 2023), MMLU-Pro for broad-domain knowledge and reasoning (Wang et al., 2024), MGSM for multilingual mathematical reasoning (Shi et al., 2023), and IFE-

val for instruction following (Zhou et al., 2023). We report each benchmark’s standard task metric as well as their average.

To demonstrate the versatility of BASIS as a plug-in advantage estimation method, we combine it with the original GRPO and two representative follow-up variants: GPG (Chu et al., 2025), which removes the PPO step that requires importance sampling, and GSPO (Zheng et al., 2025), which replaces the per-token importance ratio with a sequence-level importance ratio. In each of these combinations, BASIS modifies only the advantage estimator: we use the BASIS baseline with a single rollout per prompt, and compare it against three baselines: the original GRPO-type algorithm, which uses eight rollouts per prompt and the group mean as the baseline; its vanilla single-rollout version, which uses a zero baseline; and REINFORCE++, which uses a global batch baseline. To ensure a fair comparison, all algorithms use the same sampling budget of 512 rollouts per training step. Thus, the 8-rollout methods sample 64 prompts per training step, while the single-rollout methods sample 512 prompts per step. We train all baseline algorithms for 300 steps, corresponding to approximately 1.1 passes over the training set with 17,398 prompts for 8-rollout methods and approximately 8.8 passes for single-rollout methods. For BASIS, we train for only 150 steps, requiring roughly half the online learning time. As shown below, despite this reduced training budget, BASIS achieves performance comparable to the multi-rollout methods and often better performance than the single-rollout baselines.

Objective	Setting	Time	AIME 2024	AIME 2025	AMC 2023	MATH -500	Minerva Math	Olympiad Bench	HMMT 2025	Avg
GRPO	$G = 8$, step 300	15.5h	0.312	0.298	0.716	0.880	0.467	0.537	0.088	0.471
GRPO Vanilla	$G = 1$, step 300	20.7h	0.023	0.021	0.331	0.444	0.279	0.234	0.015	0.192
+REINFORCE++	$G = 1$, step 300	17.7h	0.277	0.271	0.662	0.826	0.408	0.487	0.088	0.431
+BASIS	$G = 1$, step 150	8.3h	0.303	0.281	0.758	0.892	0.426	0.559	0.093	0.473
	Δ vs $G = 8$	-7.2h	-0.8	-1.7	+4.2	+1.2	-4.0	+2.2	+0.4	+0.2
	Δ vs Vanilla	-12.4h	+28.0	+26.0	+42.7	+44.8	+14.7	+32.5	+7.8	+28.1
	Δ vs REINFORCE++	-9.4h	+2.6	+1.0	+9.5	+6.6	+1.8	+7.3	+0.4	+4.2
GPG	$G = 8$, step 300	14.0h	0.369	0.286	0.788	0.880	0.426	0.570	0.123	0.492
GPG Vanilla	$G = 1$, step 300	18.6h	0.056	0.093	0.558	0.488	0.184	0.182	0.025	0.227
+REINFORCE++	$G = 1$, step 300	12.9h	0.272	0.281	0.748	0.874	0.426	0.558	0.082	0.463
+BASIS	$G = 1$, step 150	6.4h	0.338	0.284	0.762	0.884	0.430	0.574	0.107	0.483
	Δ vs $G = 8$	-7.6h	-3.0	-0.2	-2.7	+0.4	+0.4	+0.5	-1.6	-0.9
	Δ vs Vanilla	-12.2h	+28.2	+19.2	+20.4	+39.6	+24.6	+39.2	+8.2	+25.6
	Δ vs REINFORCE++	-6.5h	+6.7	+0.3	+1.4	+1.0	+0.4	+1.6	+2.5	+2.0
GSPO	$G = 8$, step 300	14.4h	0.365	0.296	0.794	0.896	0.463	0.588	0.117	0.502
GSPO Vanilla	$G = 1$, step 300	13.0h	0.234	0.241	0.751	0.868	0.423	0.586	0.105	0.458
+REINFORCE++	$G = 1$, step 300	13.0h	0.350	0.265	0.794	0.908	0.441	0.599	0.108	0.495
+BASIS	$G = 1$, step 150	7.2h	0.320	0.269	0.754	0.888	0.460	0.565	0.114	0.481
	Δ vs $G = 8$	-7.2h	-4.5	-2.7	-4.0	-0.8	-0.4	-2.2	-0.3	-2.1
	Δ vs Vanilla	-5.8h	+8.5	+2.8	+0.3	+2.0	+3.7	-2.1	+0.8	+2.3
	Δ vs REINFORCE++	-5.8h	-3.0	+0.4	-4.0	-2.0	+1.8	-3.4	+0.5	-1.4

Table 1: Online policy optimization experiments on Qwen3-4B. G denotes the number of rollouts per prompt. BASIS with $G=1$ is trained for 150 steps and compared against three baselines trained for 300 steps: (i) the 8-rollout ($G=8$) GRPO-type algorithm, (ii) its vanilla single-rollout ($G=1$) variant, and (iii) another single-rollout variant combined with the REINFORCE++ global batch baseline. **Avg** is the unweighted average accuracy across the seven math benchmarks shown. Each BASIS row is followed by three Δ rows comparing it against the three baselines, respectively. Accuracy differences are reported in percentage points. The Δ entry under *Time* gives the reduction in online learning wall-clock time achieved by BASIS. Green indicates that BASIS outperforms the baseline; gray indicates that BASIS is within 5 percentage points below the baseline.

For the OOD evaluation, we apply BASIS only to GRPO. Additional implementation details as well as the computational cost of BASIS for baseline estimation are provided in Appendix D. In Appendix C, we report an additional study in which we post-train Qwen2.5-Math-7B and compare BASIS against GRPO under the same protocol.

Finding 1. BASIS is competitive with 8-rollout methods while halving the online learning budget. Table 1 reports the results on the seven in-domain mathematical reasoning benchmarks. Compared with 8-rollout GRPO, BASIS achieves a slightly higher average accuracy across the seven benchmarks, with a gain of **0.2** percentage points, and outperforms GRPO on **4** of the seven benchmarks. Compared with 8-rollout GPG and GSPO, BASIS is slightly lower on average, by **0.9** and **2.1** percentage points, respectively. Importantly, these results are obtained with half the online sampling budget: BASIS uses **76.8K** sampled responses, compared with **153.6K** for the 8-rollout baselines, reducing online learning wall-clock time by **7.2** to **7.6** hours across the three GRPO-type algorithms.

Finding 2. BASIS often outperforms single-rollout methods using roughly one-half of their online learning time. Table 1 further shows that BASIS improves the average accuracy across the seven benchmarks by **28.1** and **25.6** percentage points over single-rollout vanilla GRPO and GPG, respectively, and by **4.2** and **2.0** percentage points over their REINFORCE++ variants. These gains are positive on every individual benchmark, ranging from **0.3** to **44.8** percentage points. Importantly, BASIS achieves these improvements with substantially less training time, ranging from about **half** the compute used by the REINFORCE++ variants to about one-third of that used by single-rollout GPG.

Compared with single-rollout GSPO, the improvement is more modest. Using roughly **55%** of the training time of the two GSPO single-rollout variants, BASIS improves the average accuracy over vanilla GSPO by **2.3** percentage points, but is **1.4** percentage points lower than its REINFORCE++ variant with half of the compute budget.

The OOD results reported in Table 2 further confirm this finding: BASIS improves the aver-

Objective	Setting	HumanEval+	MBPP+	MMLU-Pro	MGSM	IFEval	Avg
Base model	no RL	67.1	64.0	51.2	78.9	79.7	68.2
GRPO	$G = 8$, step 300	82.3	71.2	61.5	82.0	83.4	76.1
GRPO Vanilla	$G = 1$, step 300	57.9	53.4	47.8	69.1	76.3	60.9
+REINFORCE++	$G = 1$, step 300	73.8	66.9	56.3	80.5	78.9	71.3
+BASIS	$G = 1$, step 150	82.3	69.8	61.8	82.6	80.2	75.3
	Δ vs $G = 8$	0.0	-1.4	+0.3	+0.6	-3.2	-0.7
	Δ vs Vanilla	+24.4	+16.4	+14.0	+13.5	+3.9	+14.4
	Δ vs REINFORCE++	+8.5	+2.9	+5.5	+2.1	+1.3	+4.1

Table 2: OOD evaluation of GRPO variants trained on DAPO-Math-17K.

age score by **14.4** percentage points over vanilla single-rollout GRPO and by **4.1** points over REINFORCE++, with consistent gains across all five OOD benchmarks.

Finding 3. BASIS prevents collapse in single-rollout policy optimization. A closer look into the learning curves in Figure 6 (Appendix C) reveals that vanilla GRPO and GPG collapse during training: vanilla GRPO barely learns, and its accuracy measure generally decreases over training. Vanilla GPG improves at first, but after peaking before the middle of training, its performance steadily declines. In contrast, BASIS does not suffer from this collapse. Together with the results in Section 5.1, these findings suggest that sample efficient value and advantage estimation stabilizes single-rollout policy optimization and helps prevent collapse.

6 Discussion

This paper introduces BASIS for rollout-efficient RLVR. Methodologically, BASIS samples only one rollout per prompt during online training, while leveraging batchwise information sharing to improve value and advantage estimation (Figure 2). Empirically, BASIS is highly sample efficient and robust for value function estimation, and often produces informative advantage estimates (Figure 3). For policy optimization, BASIS often improves over single-rollout baselines, helps prevent collapse during training and achieves similar performance to multi-rollout baselines with substantially less computation (Table 1). These gains also persist in out-of-domain evaluations covering code generation, general knowledge, multilingual reasoning, and instruction following (Table 2).

Although our experiments in Section 5 primarily use binary verifiable rewards, BASIS can also be extended to settings with continuous rewards. Appendix E discusses this extension and evaluates it

in a reinforcement learning from human feedback (RLHF) setting on the Helpful and Harmless (HH) dataset (Bai et al., 2022). It can be seen from Figure 9 that BASIS attains higher training rewards than the single-rollout REINFORCE and REINFORCE++ baselines in the later stage of training, providing evidence that its benefits extend beyond binary reward RLVR settings.

Limitations

This work has a few limitations. First, the BLUE and unbiasedness guarantees in Proposition 1 require the initial value estimates to be independent of the data batch used for batchwise refinement. These guarantees therefore do not directly apply when the initial estimates are adaptively calibrated using the same batch. While Appendix D outlines a simple way to preserve the required independence by calibrating the initial value estimates using the preceding batch instead, a more thorough investigation of this issue is left for future work. Second, BASIS improves value and advantage estimation by borrowing information across prompts within the same training batch. A complementary line of work borrows information for the same prompt across different training steps (Wang et al., 2025; Xu and Ding, 2025; Gong et al., 2026; Ruan et al., 2026). These approaches could potentially be combined with BASIS to share information both across prompts within a batch and across training iterations, further improving the post-training algorithm’s sample efficiency.

Acknowledgments

This work was supported by the UK Government’s AI Research Resource (AIRR) programme and undertaken using the Isambard-AI system. This paper used generative AI tools for language editing and code development. All AI-assisted content and

code were reviewed and verified by the authors, who take full responsibility for the final manuscript and implementation.

References

- 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. *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova Dassarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Thomas Henighan, Nicholas Joseph, Saurav Kadavath, John Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, and 12 others. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *ArXiv*, abs/2204.05862.
- Mislav Balunovic, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. 2025. Matharena: Evaluating llms on uncontaminated math competitions. *Advances in Neural Information Processing Systems*, 38.
- Kianté Brantley, Mingyu Chen, Zhaolin Gao, Jason Lee, Wen Sun, Wenhao Zhan, and Xuezhou Zhang. 2025. Accelerating rl for llm reasoning with optimal advantage regression. *Advances in Neural Information Processing Systems*, 38:151492–151531.
- Minghan Chen, Guikun Chen, Wenguan Wang, and Yi Yang. 2025. [Seed-grpo: Semantic entropy enhanced grpo for uncertainty-aware policy optimization](#). *arXiv preprint arXiv:2505.12346v1*.
- Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Xin Zhao, Zhenliang Zhang, and Furu Wei. 2026. Reasoning with exploration: An entropy perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 30377–30385.
- Xiangxiang Chu, Hailang Huang, Xiao Zhang, Fei Wei, and Yong Wang. 2025. [GPG: A simple and strong reinforcement learning baseline for model reasoning](#). *Preprint*, arXiv:2504.02546.
- Muzhi Dai, Shixuan Liu, and Qingyi Si. 2025. Stable reinforcement learning for efficient reasoning. *arXiv preprint arXiv:2505.18086*.
- Damek Davis and Benjamin Recht. 2025. [What is the objective of reasoning with reinforcement learning?](#) *arXiv preprint arXiv:2510.13651v1*.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, and 1 others. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Shijin Gong, Kai Ye, Jin Zhu, Xinyu Zhang, Hongyi Zhou, and Chengchun Shi. 2026. Kernelized advantage estimation: From nonparametric statistics to llm reasoning. *arXiv preprint arXiv:2604.28005*.
- Evan Greensmith, Peter L Bartlett, and Jonathan Baxter. 2004. Variance reduction techniques for gradient estimates in reinforcement learning. *Journal of Machine Learning Research*, 5(Nov):1471–1530.
- Kevin Han, Yuhang Zhou, Mingze Gao, Gedi Zhou, Serena Li, Abhishek Kumar, Xiangjun Fan, Weiwei Li, and Lizhu Zhang. 2026. Ebpo: Empirical bayes shrinkage for stabilizing group-relative policy optimization. *arXiv preprint arXiv:2602.05165*.
- Yaru Hao, Li Dong, Xun Wu, Shaohan Huang, Zewen Chi, and Furu Wei. 2025. [On-policy rl with optimal reward baseline](#). *arXiv preprint arXiv:2505.23585v2*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 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. Association for Computational Linguistics.
- 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. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks)*.
- Zhenyu Hou, Yujiang Li, Jie Tang, and Yuxiao Dong. 2026. Single-rollout asynchronous optimization for agentic reinforcement learning. *arXiv preprint arXiv:2607.07508*.
- Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. 2025. Reinforce++: Stabilizing critic-free policy optimization with global advantage normalization. *arXiv preprint arXiv:2501.03262*.
- Yu Huang, Zixin Wen, Yuejie Chi, Yuting Wei, Aarti Singh, Yingbin Liang, and Yuxin Chen. 2026. [The implicit curriculum: Learning dynamics in rl with verifiable rewards](#). *arXiv preprint arXiv:2602.14872*.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Øyvind Tafjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh,

- Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. 2022. Solving quantitative reasoning problems with language models. In *Advances in Neural Information Processing Systems*, volume 35.
- Ziniu Li, Tian Xu, Yushun Zhang, Zhihang Lin, Yang Yu, Ruoyu Sun, and Zhi-Quan Luo. 2023. ReMax: A simple, effective, and efficient reinforcement learning method for aligning large language models. *arXiv preprint arXiv:2310.10505*.
- Zhihang Lin, Mingbao Lin, Yuan Xie, and Rongrong Ji. 2025. Cppo: Accelerating the training of group relative policy optimization-based reasoning models. *arXiv preprint arXiv:2503.22342*.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is your code generated by ChatGPT really correct? rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems*, volume 36.
- Zhaowei Liu, Xin Guo, Zhi Yang, Fangqi Lou, Lingfeng Zeng, Jinyi Niu, Mengping Li, Qi Qi, Zhiqiang Liu, Yiyang Han, and 1 others. 2025a. Fin-r1: A large language model for financial reasoning through reinforcement learning. *arXiv preprint arXiv:2503.16252*.
- 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*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Lei Pang, Jun Luo, and Ruinan Jin. 2025. **Tic-grpo: Provable and efficient optimization for reinforcement learning from human feedback**. *arXiv preprint arXiv:2508.02833*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Herbert Robbins and Sutton Monro. 1951. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407.
- Kai Ruan, Jinghao Lin, Qianshan Wei, Ziqi Zhou, and Zihe Huang. 2026. Spo++: Stream-aligned policy optimization for asynchronous agentic rl. *arXiv preprint arXiv:2608.24870*.
- John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. 2016. High-dimensional continuous control using generalized advantage estimation. *International Conference on Learning Representations (ICLR)*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 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. 2024. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint arXiv:2409.19256*.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2023. Language models are multi-lingual chain-of-thought reasoners. In *International Conference on Learning Representations*.
- Vaishnavi Shrivastava, Ahmed Awadallah, Vidhisha Balachandran, Shivam Garg, Harkirat Behl, and Dimitris Papailiopoulos. 2025. Sample more to think less: Group filtered policy optimization for concise reasoning. *arXiv preprint arXiv:2508.09726*.
- Zhenpeng Su, Leiyu Pan, Xue Bai, Dening Liu, Guanting Dong, Jiaming Huang, Minxuan Lv, Wenping Hu, Fuzheng Zhang, Kun Gai, and 1 others. 2025. Klear-reasoner: Advancing reasoning capability via gradient-preserving clipping policy optimization. *arXiv preprint arXiv:2508.07629*.
- Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, pages 1057–1063.
- Milan Vojnovic and Se-Young Yun. 2025. **What is the alignment objective of grpo?** *arXiv preprint arXiv:2502.18548v3*.
- Hu Wang, Congbo Ma, Ian Reid, and Mohammad Yaqub. 2025. Kalman filter enhanced grpo for reinforcement learning-based language model reasoning. *arXiv preprint arXiv:2505.07527*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024. MMLU-Pro: A more robust and challenging multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*.

- Ronald J. Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3–4):229–256.
- Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong Zhang, Caiming Xiong, and 1 others. 2025. A minimalist approach to llm reasoning: from rejection sampling to reinforce. *arXiv preprint arXiv:2504.11343*.
- Yixuan Even Xu, Yash Savani, Fei Fang, and J Zico Kolter. 2025. Not all rollouts are useful: Down-sampling rollouts in llm reinforcement learning. *arXiv preprint arXiv:2504.13818*.
- Zhongwen Xu and Zihan Ding. 2025. Single-stream policy optimization. *arXiv preprint arXiv:2509.13232*.
- Fengkai Yang, Zherui Chen, Xiaohan Wang, Xiaodong Lu, Jiajun Chai, Guojun Yin, Wei Lin, Shuai Ma, Fuzhen Zhuang, Deqing Wang, Yaodong Yang, Jianxin Li, and Yikun Ban. 2026. [Your group-relative advantage is biased](#). *arXiv preprint arXiv:2601.08521v2*.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, and 1 others. 2025. DAPO: An open-source LLM reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Guanning Zeng, Zhaoyi Zhou, Daman Arora, and Andrea Zanette. 2025. Shrinking the variance: Shrinkage baselines for reinforcement learning with verifiable rewards. *arXiv preprint arXiv:2511.03710*.
- Runzhe Zhan, Yafu Li, Zhi Wang, Xiaoye Qu, Dongrui Liu, Jing Shao, Derek F Wong, and Yu Cheng. 2025. Exgrp: Learning to reason from experience. *arXiv preprint arXiv:2510.02245*.
- Jixiao Zhang and Chunsheng Zuo. 2025. Grpo-lead: A difficulty-aware reinforcement learning approach for concise mathematical reasoning in language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 5642–5665.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, and 1 others. 2025. Group sequence policy optimization. *arXiv preprint arXiv:2507.18071*.
- Hongyi Zhou, Kai Ye, Erhan Xu, Jin Zhu, Ying Yang, Shijin Gong, and Chengchun Shi. 2026. Demystifying group relative policy optimization: Its policy gradient is a u-statistic. *arXiv preprint arXiv:2603.01162*.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

Appendix

A Experimental Details for Value Estimation

This section details the value estimation experiments in Section 5.1. All experiments are conducted at a fixed GRPO-trained checkpoint π_{θ_t} , initialized from Qwen2.5-Math-7B and optimized on the Level 3–5 subset of the MATH dataset (Hendrycks et al., 2021). For each repeat, we sample a batch of prompts and estimate each prompt’s oracle value $V_{i,t}$ using 256 Monte Carlo rollouts from the checkpoint. All value estimators are then constructed from a separate set of reward samples from the same checkpoint.

More specifically, we fix the batch size at $B = 64$ and vary the number of rollouts per prompt $G \in \{1, 2, 4, 8\}$ for GRPO and $G \in \{2, 4, 8\}$ for RLOO (which requires at least two rollouts to leave one out), while BASIS is evaluated in the single-rollout setting with $G = 1$. For each repeat, we draw a batch of prompts without replacement and sample up to eight rewards per prompt. To ensure a fair comparison across different values of G , the value estimator for each G is computed using the first G rewards from the same set of rollouts. The reported metric is the mean squared error $(\hat{V}_{i,t} - V_{i,t})^2$, averaged over prompts and then over 10 repeats.

Baseline algorithms. We compare BASIS with three baseline algorithms:

- **GRPO** uses the within-prompt group reward mean as its baseline:

$$b_{i,g}^{\text{GRPO}} = \bar{r}_i := \frac{1}{G} \sum_{\ell=1}^G r_{i,\ell}$$

- **RLOO** uses a leave-one-out group baseline, where the baseline for each sample is computed from the other completions generated from the same prompt:

$$b_{i,g}^{\text{RLOO}} = \bar{r}_{i,-g} := \frac{1}{G-1} \sum_{\ell \neq g} r_{i,\ell},$$

- **REINFORCE++** uses the global reward mean within the current batch as its baseline:

$$b_i^{\text{R++}} = \bar{r}_{\text{batch}} := \frac{1}{B} \sum_{j=1}^B r_j,$$




Figure 4: MSE of RLOO value estimator (left) and frequency with which its estimated baseline is exactly 0 or 1 (right). Results for BASIS and REINFORCE++ are also reported for completeness.

Effect of batch heterogeneity. We examine how the accuracy of the value estimator varies with the heterogeneity of the sampled prompt batch. We fix $B = 64$ and sample 500 batches to ensure sufficient diversity across the sampled prompt batches. For each batch b , we define its heterogeneity score as $s_b := \text{std}_{i \in b}(V_{i,t})$, which is the standard deviation of the 64 oracle prompt values in that batch. The minimum and maximum observed heterogeneity scores across the 500 batches are 0.206 and 0.307, and we divide this interval into five bins with uniformly spaced boundaries: $[0.206, 0.226)$, $[0.226, 0.246)$, $[0.246, 0.266)$, $[0.266, 0.287)$, and $[0.287, 0.307]$. Within each bin, we report the MSE of each value estimator, first averaged over prompts within each batch and then averaged across batches in the bin.

Effect of prompt difficulty. We further evaluate how the accuracy of the value estimator varies across different prompt-difficulty levels. For each repeat, we first sample a batch of $B = 64$ prompts and compute the oracle value estimate per prompt. We next stratify all prompts by their oracle values into five difficulty bins: $[0, 0.2)$, $[0.2, 0.4)$, $[0.4, 0.6)$, $[0.6, 0.8)$, and $[0.8, 1]$. Within each bin, we similarly report the MSE of each value estimator, averaged over prompts within that bin and across 10 repeats.

Sensitivity analysis. We evaluate the sensitivity of BASIS to the active-set threshold ϵ , the number of offline reference samples n , the online estimation batch size B , and the number of candidate values in the β grid. Following the value estimation experiments in Section 5.1, we use Qwen2.5-Math-7B checkpoints from the early, middle, and late stages of training and compute oracle values using 256 held-out rollouts. We vary one hyperparameter at a time while holding all others at their default values ($\epsilon = 10^{-6}$, $n = 64$, and $B = 64$, together

with the default β grid).

Figure 5 shows that BASIS is stable over a broad range of hyperparameter choices. The MSE is nearly unchanged for $\epsilon \in [10^{-8}, 10^{-2}]$, and using between 8 and 64 candidate β values produces similar results. Increasing the offline cache from $n = 8$ to $n = 64$ reduces MSE by approximately 55–59% across the three training stages. Increasing B improves estimation up to about $B = 64$ –128, beyond which the improvement is marginal.

B Alternative Information-Sharing Approaches

In Section 4, we construct the BASIS baseline as a weighted average of rewards from the same batch, with data-adaptive weights chosen to form a best linear unbiased baseline; we refer to this main estimator as UNB for unbiasedness. More broadly, batchwise information can be shared in other ways, such as using weighted averages without the unbiasedness constraint, or using importance weights based on ratios of initial value estimates. In this section, we study two such alternatives, namely, the Variance-Optimized shrinkage baseline (VOP) and the Ratio-average Value-Guided baseline (RVG). Using the Qwen2.5-Math-7B checkpoints, single-rollout VOP and RVG perform comparably to UNB and match or outperform GRPO with $G=8$ rollouts on most math benchmarks, while using roughly half the online training time (see Appendix C, Table 4).

Before defining each variant, we note that both approaches reuse the same offline initial value estimator and the same calibration method for selecting β_t as UNB does. We use notation similar to that in Section 4: $\hat{V}_{i,t} = \hat{V}_{\beta_t}(x_i)$, $\hat{\sigma}_i^2 = \hat{V}_{i,t}(1 - \hat{V}_{i,t})$, and $\mathcal{A}_t$ is the active set in Algorithm 1. The formulas below apply only to active prompts; prompts outside $\mathcal{A}_t$ use the zero baseline. For reference, the UNB proposed in the main text can be written as

$$\tilde{V}_{i,t}^{\text{UNB}} = \hat{V}_{i,t} \frac{\sum_{j \in \mathcal{A}_t \setminus \{i\}} \hat{V}_{j,t} r_j / \hat{\sigma}_j^2}{\sum_{j \in \mathcal{A}_t \setminus \{i\}} \hat{V}_{j,t} / \hat{\sigma}_j^2}, \quad i \in \mathcal{A}_t, \quad (6)$$

by plugging $\hat{V}_{i,t}$ and $\hat{\sigma}_i^2$ into Eq. (3).

Both variants detailed below share the same idea of weighting the batch rewards by the offline values $\hat{V}_{i,t}$, and they differ in the weights.

Variance-Optimized shrinkage baseline (VOP).

VOP drops the unbiasedness constraint adopted by UNB. The following proposition derives VOP’s weights; its proof is given in Appendix F.

Proposition 3. *Fix a training step t , a target prompt i , and condition on the prompt batch $\mathcal{X}_B := \{x_k\}_{k=1}^B$. Among all leave-one-out linear baselines $b_i = \sum_{j \neq i} w_{ij} r_j$, without imposing the unbiasedness constraint, the weights minimizing $\mathbb{E}[(V_{i,t} - b_i)^2 \mid \mathcal{X}_B]$ are*

$$w_{ij}^{\text{VOP}} = \frac{V_{i,t} V_{j,t} / \sigma_j^2}{1 + \sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}, \quad j \neq i.$$

By Proposition 3, for each active prompt i , we define the VOP estimator by

$$\tilde{V}_{i,t}^{\text{VOP}} = \hat{V}_{i,t} \frac{\sum_{j \in \mathcal{A}_t \setminus \{i\}} \hat{V}_{j,t} r_j / \hat{\sigma}_j^2}{1 + \sum_{j \in \mathcal{A}_t \setminus \{i\}} \hat{V}_{j,t}^2 / \hat{\sigma}_j^2}, \quad i \in \mathcal{A}_t.$$

Compared to the UNB estimator in Eq. (6), VOP includes an additional 1 in the denominator, which shrinks the estimator toward zero, trading unbiasedness for lower MSE.

Ratio-average Value-Guided baseline (RVG).

RVG computes the value estimator by

$$\tilde{V}_{i,t}^{\text{RVG}} = \frac{1}{|\mathcal{A}_t| - 1} \sum_{j \in \mathcal{A}_t \setminus \{i\}} \frac{\hat{V}_{i,t}}{\hat{V}_{j,t}} r_j, \quad i \in \mathcal{A}_t.$$

It can be viewed as an importance sampling estimator, where the ratio $\hat{V}_{i,t} / \hat{V}_{j,t}$ reweights rewards from other prompts to borrow information.

C Additional Online Learning Results

Section 5 reports BASIS’s performance at step 150. In this section, we evaluate its performance over steps 150–300 and compare against the vanilla single-rollout algorithms. The results are visualized in Figure 6. We observe that vanilla GRPO and GPG collapse after approximately step 150, with their performance deteriorating substantially as training progresses. Only vanilla GSPO survives beyond this point. BASIS lifts every objective onto a monotonically improving trajectory. Figure 7 compares BASIS against a stronger single-rollout baseline, REINFORCE++. It can be seen that BASIS dominates uniformly across all seven benchmarks, with a visible gap from Step 100 onward.




Figure 5: MSE of BASIS value estimator at early, middle, and late training stages in the sensitivity analysis. Shaded regions are 95% confidence intervals. The batch-size panel also reports the performance of REINFORCE++ for reference.




Figure 6: Accuracy of Qwen3-4B models fine-tuned by *Vanilla* (with a zero baseline) and *BASIS* coupled with GRPO (orange), GPG (red), GSPO (teal) across seven math benchmarks. Solid lines are exponentially smoothed with a window of 3; lighter curves in the same colors show the raw scores.

In addition to the experiments on Qwen3-4B, we apply the same online RL protocol to Qwen2.5-Math-7B, compare the three BASIS weighting variants described in Appendix B, and report the results in Table 4. We observe that the three BASIS variants match or outperform GRPO on five of the seven benchmarks, despite being trained for only 150 steps, compared with 300 steps for GRPO: UNB performs best on AMC 2023 and MATH-500, RVG on AIME 2024 and OlympiadBench, and VOP on AIME 2025. These results demonstrate that BASIS generalizes across models (Qwen2.5-Math-7B and Qwen3-4B) and across different weighting variants, suggesting that the benefits of information sharing are not tied to a particular choice of weights and models.

D Implementation Details for Policy Optimization

Before online RL, we sample n completions from the reference policy on each training prompt and score them with the same verifier used for online rewards. For each prompt x_i , the cache stores the offline sample size $n_i = n$ and the empirical mean reward $\hat{p}_i$. We set $n = 64$ for the Qwen3-4B experiments and $n = 32$ for the Qwen2.5-Math-7B experiments. In general, we recommend using a comparable offline sample size, subject to the available offline sampling budget.

For binary rewards, our estimated value takes the following form:

$$\hat{V}_\beta(x_i) = \frac{\hat{p}_i \exp(1/\beta)}{1 - \hat{p}_i + \hat{p}_i \exp(1/\beta)},$$

given a regularization parameter β . We evaluate $\hat{V}_\beta(x_i)$ on a fixed grid of 230 values of β spanning




Figure 7: Accuracy of Qwen3-4B models fine-tuned by *REINFORCE++* (with a global batch baseline) and *BASIS* coupled with GRPO (orange), GPG (red), GSPO (teal) across seven math benchmarks. The rest is the same as in Figure 6.

$[0.01, 5.0]$: 200 values with a step size of 0.01 over $[0.01, 2.00]$, followed by 30 values with a step size of 0.1 over $[2.1, 5.0]$. This yields a $17,398 \times 230$ lookup table, occupying approximately 16 MB on disk. Figure 8 visualizes the resulting calibrated β values over the course of online training.

The offline rollout process takes a few hours on the same four-GPU node used for training, while the subsequent closed-form evaluation of $\tilde{V}_\beta$ takes only seconds on a single CPU. Once computed, the initial estimator can be reused across all three policy objectives and all three weighting variants (UNB / RVG / VOP). Since this precomputation is shared across the BASIS settings, its effective per-experiment overhead is well below 5% of the training wall-clock time.

For online policy learning, we adopt the implementation from ver1 (Sheng et al., 2024), replacing only its advantage estimator with BASIS’s estimator. At each calibration step, BASIS chooses β_t by minimizing Eq. (5) over the precomputed grid. This amounts to minimizing over the 230 precomputed grid values for each of the 512 active prompts, costing less than a millisecond per step on a single GPU. As an alternative to using the same data batch for both batchwise refinement and calibration, one can select the calibration parameter using the preceding online batch to preserve the




Figure 8: Calibrated values of β_t during BASIS training on Qwen3-4B when coupled with GRPO, GPG, and GSPO. At each training step, the calibrator selects β_t by minimizing Eq. (5) over the precomputed 230-point grid $\beta \in [0.01, 5.0]$. All three runs start near the upper end of the grid and settle into a narrow band around $\beta \approx 0.4$ by approximately step 80, where they remain for the rest of training. Solid lines show the smoothed calibrated β_t values; lighter curves in the same colors show the corresponding raw selected β_t values.

unbiasedness of the value estimate:

$$\mathcal{L}_s(\beta) = \frac{1}{|\mathcal{A}_{s,\beta}|} \sum_{i \in \mathcal{A}_{s,\beta}} (r_i^{(s)} - \tilde{V}_{i,s,\beta})^2,$$

$$\beta_t^{\text{PB}} \in \arg \min_{\beta \in \mathbb{B}} \mathcal{L}_{t-1}(\beta), \quad b_i^{(t)} = \tilde{V}_{i,t,\beta_t^{\text{PB}}}.$$

Prompts whose selected offline value falls outside the active set use the zero baseline; all other prompts use the cross-prompt BASIS baseline. We set $\epsilon = 10^{-6}$ for the active set threshold throughout.

Objective	Setting	Step	Responses	Time
GRPO	$G = 8$	300	153.6K	15.5h
GRPO Vanilla	$G = 1$	300	153.6K	20.7h
+REINFORCE++	$G = 1$	300	153.6K	17.7h
+BASIS	$G = 1$	150	76.8K	8.3h
+BASIS	$G = 1$	300	153.6K	16.6h
GPG	$G = 8$	300	153.6K	14.0h
GPG Vanilla	$G = 1$	300	153.6K	18.6h
+REINFORCE++	$G = 1$	300	153.6K	12.9h
+BASIS	$G = 1$	150	76.8K	6.4h
+BASIS	$G = 1$	300	153.6K	12.7h
GSPO	$G = 8$	300	153.6K	14.4h
GSPO Vanilla	$G = 1$	300	153.6K	13.0h
+REINFORCE++	$G = 1$	300	153.6K	13.0h
+BASIS	$G = 1$	150	76.8K	7.2h
+BASIS	$G = 1$	300	153.6K	14.3h

Table 3: Computation time for fine-tuning Qwen3-4B models. Both single-rollout and multi-rollout baseline algorithms are trained for 300 steps, corresponding to approximately 1.1 passes over the DAPO-Math-17K training set for 8-rollout methods and approximately 8.8 passes for single-rollout methods. BASIS is trained for 150 steps.

Method	AIME 2024	AIME 2025	AMC 2023	MATH -500	Minerva Math	Olympiad Bench	HMMT 2025
GRPO ($G = 8$)	0.244	0.102	0.597	0.762	0.324	0.406	0.019
+BASIS-UNB ($G = 1$)	<u>0.245</u>	0.107	0.655	0.780	<u>0.309</u>	0.409	0.019
+BASIS-RVG ($G = 1$)	0.256	<u>0.117</u>	0.622	<u>0.773</u>	0.298	0.417	0.015
+BASIS-VOP ($G = 1$)	0.240	0.118	<u>0.634</u>	0.768	0.303	<u>0.416</u>	<u>0.018</u>

Table 4: Online policy optimization experiments on Qwen2.5-Math-7B. Similar to the experiments on Qwen3-4B, BASIS is trained for 150 steps, whereas GRPO is trained for 300 steps. All algorithms sample 512 rollouts per training step, so BASIS methods train on eight times as many different prompts per step as 8-rollout GRPO. UNB is the main algorithm presented in the main paper (Table 1); RVG and VOP variants are defined in Appendix B. The best result in each column is bolded, and the second best result is underlined.

Finally, the Qwen2.5-Math-7B experiments use max prompt length 1024, max response length 2048, learning rate 10^{-6} , rollout temperature 1.0, top- $p=1.0$, top- $k=-1$, DAPO-Math-17K training data (Yu et al., 2025), on four GH200 GPUs. The Qwen3-4B experiments use the same learning rate and temperature with max response length 4096, on four GH200 GPUs. In both cases, the rollout budget is 512 per online training step. The PPO mini-batch size is 256 for single-rollout algorithms and 32 for 8-rollout algorithms, respectively, so that each training step performs 2 PPO inner updates. We did not implement KL regularization, following the recent consensus (Yu et al., 2025).

E Extensions to Continuous Rewards

BASIS requires an estimate of the reward variance to implement batchwise refinement. For binary rewards, we can use the plug-in estimate $\hat{\sigma}_i^2 = \hat{V}_{i,t}(1 - \hat{V}_{i,t})$. For a continuous scalar reward, however, the reward variance cannot be recovered from its value function alone. We instead estimate the conditional second moment $M_{2,j,t} := \mathbb{E}[r_j^2 | x_j]$ to derive the reward variance

$$\sigma_j^2 = M_{2,j,t} - V_{j,t}^2.$$

This additional quantity can be obtained from the same offline reference-policy rollouts used to estimate the value. Indeed, applying the exponential-tilting identity in Proposition 2 to the first and second reward moments gives

$$V_{\beta}^*(x) = \frac{\mathbb{E}_{\pi_{\text{ref}}} [r(x, y) e^{r(x, y)/\beta} | x]}{\mathbb{E}_{\pi_{\text{ref}}} [e^{r(x, y)/\beta} | x]},$$

$$M_{2,\beta}^*(x) = \frac{\mathbb{E}_{\pi_{\text{ref}}} [r(x, y)^2 e^{r(x, y)/\beta} | x]}{\mathbb{E}_{\pi_{\text{ref}}} [e^{r(x, y)/\beta} | x]}.$$

Suppose the offline cache contains M reference-policy rewards $r_{j1}, \dots, r_{jM}$ for prompt x_j and let $q_{jm}(\beta) = \exp(r_{jm}/\beta)$. We then compute

$$\hat{V}_{\beta}(x_j) = \frac{\sum_{m=1}^M q_{jm}(\beta) r_{jm}}{\sum_{m=1}^M q_{jm}(\beta)},$$

$$\hat{M}_{2,\beta}(x_j) = \frac{\sum_{m=1}^M q_{jm}(\beta) r_{jm}^2}{\sum_{m=1}^M q_{jm}(\beta)},$$

$$\hat{\sigma}_{\beta}^2(x_j) = \max\{\hat{M}_{2,\beta}(x_j) - \hat{V}_{\beta}(x_j)^2, \delta\},$$

where $\delta > 0$ is a small constant.

At each online step t , we set $\hat{V}_{j,t,\beta} = \hat{V}_{\beta}(x_j)$ and $\hat{\sigma}_{j,t,\beta}^2 = \hat{\sigma}_{\beta}^2(x_j)$ and substitute these estimates into the formula from Proposition 1, which yields:

$$\hat{w}_{ij,t,\beta} = \frac{\hat{V}_{i,t,\beta} \hat{V}_{j,t,\beta} / \hat{\sigma}_{j,t,\beta}^2}{\sum_{k \neq i} \hat{V}_{k,t,\beta}^2 / \hat{\sigma}_{k,t,\beta}^2},$$

$$\tilde{V}_{i,t,\beta} = \sum_{j \neq i} \hat{w}_{ij,t,\beta} r_j. \quad (7)$$

If the denominator in Eq. (7) is zero, we use the same zero baseline as in the binary reward setting. Finally, β_t is selected from the current online batch by Eq. (5).

We further implement this extension using the HH dataset, which contains preferred and rejected




Figure 9: Training rewards on the HH dataset for BASIS and the single-rollout REINFORCE and REINFORCE++ baselines.

responses. We partition the data into three disjoint subsets, containing 5,000 preference pairs for training a reward model, 2,000 preferred responses for supervised fine-tuning, and 2,048 prompts for RL policy optimization. For each subset, we reserve 5% of the data for validation.

We initialize the policy and reward model from Qwen2.5-1.5B-Instruct, fine-tune the policy on preferred responses for 200 steps, and train a Bradley-Terry reward model on paired preferences for 1,000 steps. For RL policy optimization, all methods use maximum prompt and response lengths of 512 tokens, sampling temperature 1.0, learning rate 5×10^{-7} . BASIS constructs its offline prompt values from 32 responses per prompt sampled from the supervised policy and scored by the reward model. We standardize the reward-model scores using the mean and standard deviation computed from these offline scores, subtracting the offline mean and dividing by the standard deviation of these scores. Figure 9 reports the resulting standardized rewards. We observe that BASIS achieves substantially higher rewards than both REINFORCE and REINFORCE++ at later stages of training.

F Proofs

F.1 Proof of Proposition 1

Proof. Fix a training step t , a target prompt i , and condition on the prompt batch $\mathcal{X}_B := \{x_k\}_{k=1}^B$. Under the independent rollout sampling used by BASIS, the rewards from different prompts are conditionally uncorrelated given the prompt batch. Hence, for any leave-one-out linear baseline $b_i =$

$$\sum_{j \neq i} w_{ij} r_j,$$

$$\mathbb{E}[b_i | \mathcal{X}_B] = \sum_{j \neq i} w_{ij} V_{j,t}.$$

The unbiasedness constraint is therefore

$$\sum_{j \neq i} w_{ij} V_{j,t} = V_{i,t}.$$

On this constrained set, the conditional MSE equals the conditional variance:

$$\begin{aligned} \mathbb{E}[(V_{i,t} - b_i)^2 | \mathcal{X}_B] &= \text{Var}(b_i | \mathcal{X}_B) \\ &= \sum_{j \neq i} w_{ij}^2 \sigma_j^2. \end{aligned}$$

Thus the optimal weights solve the quadratic program

$$\begin{aligned} \min_{(w_{ij})_{j \neq i}} \quad & \sum_{j \neq i} w_{ij}^2 \sigma_j^2 \\ \text{s.t.} \quad & \sum_{j \neq i} w_{ij} V_{j,t} = V_{i,t}. \end{aligned}$$

In the non-degenerate case $\sum_{k \neq i} V_{k,t}^2 / \sigma_k^2 > 0$, the Lagrangian

$$\mathcal{L}(w_i, \lambda) = \sum_{j \neq i} w_{ij}^2 \sigma_j^2 + \lambda \left(\sum_{j \neq i} w_{ij} V_{j,t} - V_{i,t} \right)$$

has first-order condition

$$2w_{ij}\sigma_j^2 + \lambda V_{j,t} = 0, \quad j \neq i.$$

Therefore $w_{ij} = -\lambda V_{j,t} / (2\sigma_j^2)$. Substituting this expression into the unbiasedness constraint gives

$$-\frac{\lambda}{2} \sum_{k \neq i} \frac{V_{k,t}^2}{\sigma_k^2} = V_{i,t}, \quad \lambda = -\frac{2V_{i,t}}{\sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}.$$

Hence

$$w_{ij} = \frac{V_{i,t} V_{j,t} / \sigma_j^2}{\sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}, \quad j \neq i,$$

which is Eq. (3). Since the objective is strictly convex whenever $\sigma_j^2 > 0$, these weights are the unique minimizer. $\square$

F.2 Proof of Proposition 2

Proof. Fix a prompt x and $\beta > 0$. We first derive the optimizer of the KL-regularized objective using the same variational calculation as in DPO-style derivations (Rafailov et al., 2023). Let

$$p_y = \pi(y | x), \quad p_y^{\text{ref}} = \pi_{\text{ref}}(y | x),$$

where the sums below range over responses in the support of $\pi_{\text{ref}}(\bullet \mid x)$. For this fixed prompt, the objective is

$$J_x(p) = \sum_y p_y r(x, y) - \beta \sum_y p_y \log \frac{p_y}{p_y^{\text{ref}}}.$$

The maximization is over distributions p satisfying $\sum_y p_y = 1$. Its Lagrangian is

$$\begin{aligned} \mathcal{L}(p, \lambda) &= \sum_y p_y r(x, y) - \beta \sum_y p_y \log \frac{p_y}{p_y^{\text{ref}}} \\ &\quad + \lambda \left(\sum_y p_y - 1 \right). \end{aligned}$$

For any response y with $p_y^{\text{ref}} > 0$, the first-order condition is

$$\begin{aligned} 0 &= \frac{\partial \mathcal{L}}{\partial p_y} \\ &= r(x, y) - \beta \left(\log \frac{p_y}{p_y^{\text{ref}}} + 1 \right) + \lambda. \end{aligned}$$

Solving this equation gives

$$p_y = C p_y^{\text{ref}} \exp(r(x, y)/\beta),$$

where $C = \exp(\lambda/\beta - 1)$ does not depend on y . Enforcing $\sum_y p_y = 1$ yields

$$\begin{aligned} C^{-1} &= \sum_y p_y^{\text{ref}} \exp(r(x, y)/\beta) \\ &= \mathbb{E}_{y \sim \pi_{\text{ref}}(\bullet \mid x)}[\exp(r(x, y)/\beta)] =: Z_\beta(x). \end{aligned}$$

Therefore the unique maximizer is

$$\pi_\beta^*(y \mid x) = \frac{\pi_{\text{ref}}(y \mid x) \exp(r(x, y)/\beta)}{Z_\beta(x)}.$$

The objective is concave in p because it is a linear reward term plus $-\beta$ times a convex KL term, so the stationary point is globally optimal.

Now substitute this optimizer into the reward value:

$$\begin{aligned} V_\beta^*(x) &= \sum_y \pi_\beta^*(y \mid x) r(x, y) \\ &= \frac{N_\beta(x)}{Z_\beta(x)}, \end{aligned}$$

where

$$N_\beta(x) := \sum_y p_y^{\text{ref}} r(x, y) \exp(r(x, y)/\beta).$$

Writing the sums in $N_\beta(x)$ and $Z_\beta(x)$ as expectations under $\pi_{\text{ref}}(\bullet \mid x)$ gives Eq. (4). $\square$

F.3 Proof of Proposition 3

Proof. Fix a training step t , a target prompt i , and condition on the prompt batch $\mathcal{X}_B$. For a leave-one-out linear baseline $b_i = \sum_{j \neq i} w_{ij} r_j$, the conditional MSE decomposes as

$$\begin{aligned} \mathbb{E}[(V_{i,t} - b_i)^2 \mid \mathcal{X}_B] \\ = \left(V_{i,t} - \sum_{j \neq i} w_{ij} V_{j,t} \right)^2 + \sum_{j \neq i} w_{ij}^2 \sigma_j^2. \end{aligned}$$

This is the same quadratic objective as in the proof of Proposition 1, but without the linear unbiasedness constraint. Its first-order condition is

$$\begin{aligned} -2V_{j,t} \left(V_{i,t} - \sum_{k \neq i} w_{ik} V_{k,t} \right) + 2w_{ij} \sigma_j^2 &= 0, \\ j &\neq i. \end{aligned}$$

Let $c = V_{i,t} - \sum_{k \neq i} w_{ik} V_{k,t}$. Then

$$w_{ij} = \frac{c V_{j,t}}{\sigma_j^2}.$$

Substituting this expression into the definition of c gives

$$c = V_{i,t} - c \sum_{k \neq i} \frac{V_{k,t}^2}{\sigma_k^2}, \quad c = \frac{V_{i,t}}{1 + \sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}.$$

Therefore

$$w_{ij} = \frac{V_{i,t} V_{j,t} / \sigma_j^2}{1 + \sum_{k \neq i} V_{k,t}^2 / \sigma_k^2}, \quad j \neq i,$$

which proves the claim. The objective is strictly convex whenever $\sigma_j^2 > 0$, so the minimizer is unique. $\square$