Parameter Exploration for RLVR via Variational Learning

Vatsal Venkatkrishna INSAIT ✉, **Nico Carlsen Daheim**, **Iryna Gurevych** INSAIT
INSAIT INSAIT, Sofia University "St. Kliment Ohridski", Bulgaria
Ubiquitous Knowledge Processing Lab (UKP Lab),
Department of Computer Science, Technical University of Darmstadt and
National Research Center for Applied Cybersecurity ATHENE, Germany

c3po.insait.ai

Abstract

Exploration has been a focus of reinforcement learning research for a long time. Recently, there has been growing evidence that it is also an important ingredient in LLM reinforcement learning recipes that can significantly impact downstream performance. Many existing methods control exploration in the action-space, for example, using temperature scaling. However, these methods cannot reorder tokens but only influence the variance in the output distribution. This limits exploration and can lead to divergence or stalled training. Here, we investigate parameter-space exploration, where rollouts are generated by sampling different policies from a posterior that may each explore different rollouts. Sampling less or more diverse policies is then a complementary control lever over exploration. We introduce a family of methods called Perturbed Parameter Policy Optimization (3PO) which use different sampling strategies and different rollout grouping for reward estimation. Experiments on Olmo-3-1025-7B and Qwen2.5-Math-7B across mathematical reasoning and code generation tasks show that these approaches can improve the average downstream performance over standard GRPO at a near-identical FLOPs cost. Moreover, using multiple parameter samples generally produces fewer zero-advantage groups and malformed or incorrect rollouts during training than GRPO and action-space baselines. Overall, our work presents evidence that parameter-space exploration can improve reinforcement learning for LLMs.

1 Introduction

Reinforcement Learning with Verifiable Rewards (RLVR) is a powerful post-training paradigm that can be used to improve the reasoning capabilities of LLMs without relying on human demonstrations [31, 52]. Broadly, reinforcement learning could discover novel solutions by learning from autonomous experience [60, 1], but current RLVR algorithms are frequently limited by a model's capabilities prior to RL [8, 73, 68]. Since the model learns exclusively from self-generated rollouts, high-reward trajectories may simply have too low probability and are thus often not discovered [67]. This is particularly problematic in algorithms with grouped advantage calculations like Group-Relative Policy Optimization (GRPO) [52], where there is no learning signal whenever all rollouts for a prompt receive the same reward. Such zero-advantage groups stall learning and therefore increase the already high compute costs associated with RLVR [45, 63].

A common solution is to increase the softmax temperature during rollout generation to promote exploration [2, 76]. However, as illustrated in Figure 1, temperature can only control the variance

✉ Corresponding author: vatsal.venkatkrishna@insait.ai

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

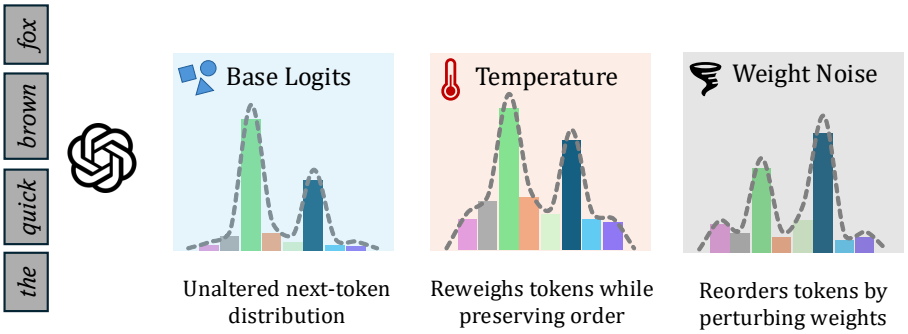


Figure 1: **Temperature versus Weight noise.** Temperature scaling controls the entropy of the token distribution but preserves relative token ordering. Adding noise to the model's weights can completely alter the token distribution, helping the model explore new regions of the solution space.

of the token distribution and not the relative token order, which limits the trajectories that can be explored. Moreover, it uniformly alters the probability of all tokens in the vocabulary, which can lead to degeneration at high temperatures when erroneous tokens become likely [21]. Other complementary strategies like advantage shaping [64, 75, 72] and entropy regularization [38, 25, 53] have been proposed, but rather aim to stabilize training than promote exploration.

Parameter-space exploration has a long track record in more classical RL [48, 50, 1]. One goal of parameter-space exploration has been to find high-quality trajectories from policies that are in the neighborhood of the current policy, thereby mitigating the issue that action-space noise might sample unlikely or low-quality outputs [12]. Additionally, weight perturbations can directly alter logit values and reorder token distributions as illustrated in Figure 1. Intuitively, if we could learn the amount of noise to add to each parameter, we could explore promising policies that each have promising alterations to token distributions and rollouts. However, to the best of our knowledge, such strategies have been underexplored for LLMs, possibly because noisy optimization itself is seldom used.

In this work, we present Perturbed Parameter Policy Optimization (3PO) which uses parameter-space exploration for RLVR by sampling models from an approximate posterior during training that can be updated during the reinforcement learning stage to adaptively learn noise injection. At each training step, one or more samples are obtained by adding learned noise to the learned parameters, which can be used to control the amount of exploration. We explore three strategies for noise injection. B3PO samples a single weight perturbation per gradient step, which makes it dependent on a single weight draw and can limit group diversity. M3PO calculates advantages separately for M model perturbations before a single weight update. C3PO divides each GRPO group across N independently sampled weights and calculates advantages on the full group to maximize diversity.

We validate 3PO on OLMo-3-1025-7B [46] and Qwen2.5-Math-7B [69] across mathematical reasoning and code generation tasks, which are widely used as RLVR testbeds [8, 10, 39, 64, 75]. Overall, we find that parameter exploration can help downstream performance and speed up convergence without using more rollouts than AdamW-based training. Our chunked noising approach, C3PO, achieves the best average performance across both model families. Broadly, using multiple model samples per batch performs better than using just a single sample, and rescues more zero-advantage groups than other action-space exploration baselines. We find consistent improvements in model performance that are concentrated on harder benchmarks like AIME [42] and LiveCodeBench [23]. Altogether, our work shows that parameter exploration via learned approximate posteriors can be helpful for RLVR by providing an additional control lever for exploration.

2 Background on RLVR and parameter exploration

2.1 RLVR for LLMs

Reinforcement Learning with Verifiable Rewards (RLVR) has recently become popular for aligning LLMs in problem settings where outputs of the model for a given problem can easily be verified, for example, in solving mathematical word problems, where a ground-truth solution exists. Overall,

the idea is to assign a reward r_i, for example, a binary $r_i \in \{0, 1\}$ to each sequence $\mathbf{y}^{(i)} \in \mathcal{V}^*$ that an LLM generates based on a problem $\mathbf{x}$, which is then subsequently used to train the model. Crucially, the outputs $\mathbf{y}^{(i)}$ are self-generated by the model being trained and not obtained via human demonstration. However, some human-provided ground-truth solution is still usually used to calculate r_i on a final answer at the end of the LLM's output $\mathbf{y}^{(i)}$.

Group Relative Policy Optimization (GRPO) [52] is a widely-used technique for RLVR training, where the learning signal is calculated based on advantages $A_{i,t} = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}$ using the shorthand $\mathbf{r} = (r_1, \ldots, r_G)$ for a group-size G. That is, a total of G rollouts are sampled for each step that are subsequently used for advantage calculation. Altogether, the GRPO loss takes the following form,

$$\mathcal{L}^{\text{GRPO}} = \mathop{\mathbb{E}}_{\mathbf{x}\sim\mathcal{D},\ \{\mathbf{y}^{(i)}\}_{i=1}^{G}\sim\pi_{\text{old}}(\cdot|\mathbf{x})} \left[\frac{1}{\sum_i |\mathbf{y}^{(i)}|} \sum_{i=1}^{G} \sum_{t=1}^{|\mathbf{y}^{(i)}|} \min\left\{ R_{i,t} A_{i,t}, \text{clip}(R_{i,t}, 1-\epsilon, 1+\epsilon) A_{i,t} \right\} \right]. \tag{1}$$

The ratio $R_{i,t} = \frac{\pi_{\boldsymbol{\theta}}(y_t^{(i)}|\mathbf{x}, \mathbf{y}_{<t}^{(i)})}{\pi_{\text{old}}(y_t^{(i)}|\mathbf{x}, \mathbf{y}_{<t}^{(i)})}$ is an importance sampling factor to account for cases where the language model that generates the rollouts (π_{old}) differs from the language model that is updated ($\pi_{\boldsymbol{\theta}}$). Here, we omit the KL-divergence term following prior work [72].

Despite the success of GRPO, several issues remain. For example, groups that lead to zero-advantages stall learning. This is the case when all rollouts achieve the same reward and either all lead to a correct or incorrect solution. Intuitively, this is connected to the problem of trading off exploration and exploitation. A good RL algorithm needs to mix failed attempts and successful attempts to sustain a learning signal, though achieving a good trade-off is often hard in practice [60]. This trade-off is similarly hard to establish in RLVR for LLM training. The entropy of the policy often collapses prematurely, and the groups sampled at each step become increasingly homogeneous, leading to a slowing or even plateauing of learning because the advantages become less informative [72].

As a consequence, a variety of approaches have been proposed to target this problem. Clip-higher [72] loosens the upper PPO clip to preserve probability mass on exploratory tokens, but it has no effect on on-policy methods as the importance sampling ratio vanishes. Another solution is to increase the softmax temperature during rollout generation in order to flatten the token distribution [2, 76], but this has limited benefits as it is rank-preserving and requires careful tuning to not overshoot. Explicit entropy regularization adds an entropy bonus to the loss [38, 25], and advantage-shaping methods reweight the learning signal toward tokens or trajectories deemed more informative [64, 75]. Other strategies are using larger rollout budgets [22, 19], prolonged training [37, 22, 71], or ordering training data by difficulty to gradually expand the reachable solution set [58, 51, 41].

Still, all of these methods operate in the action space. They reweigh or regularize token probabilities and do not alter the generators and policy's parameters. In this work, we explore an orthogonal axis: perturbing the policy's parameters directly as an additional axis of exploration in RLVR. By operating in the parameter space, our method is naturally compatible with previous action-space approaches.

2.2 Parameter space exploration

As previously established, the exploration-exploitation tradeoff is a well-studied problem in classical RL. While methods like ϵ-greedy [59, 65] and UCB [56, 24] work well for smaller problems, they tend to not scale well to problems with large action spaces, such as language modeling. Thus, there has been sustained interest in developing efficient algorithms for exploration [1].

One such method is parameter exploration, where noise is added to the policy's parameters rather than its actions [48, 50]. Early work used finite-difference gradients with fixed parameter perturbations [29]; subsequent extensions developed state-dependent and per-basis variants [61]. Salimans et al. [49] scaled the approach to deep policies via evolution strategies, and Plappert et al. [47] and Fortunato et al. [14] integrated parameter noise with modern deep RL algorithms, with the latter learning the noise magnitude jointly with the policy. Across continuous control benchmarks, episode-coherent parameter perturbation has consistently matched or outperformed action-space noise [57, 62]. Despite these successes, parameter-space exploration methods have largely been absent from RLVR for LLMs. Concurrent work [3] extends the method of Plappert et al. [47] to RLVR by adding Gaussian noise to the rollout generating policy and all rollouts are generated by

the same model. Their approach can be seen as a variant of the B3PO we present, where the noise added is controlled by an adaptive scheduler rather than a learned Hessian (cf. Eq. 4). Additionally, we study the effects of using multiple Monte-Carlo samples (M3PO) and a chunked noising approach (C3PO) to directly control group diversity in the GRPO objective, which we find is more effective.

2.3 Variational learning and IVON

In this work, we use variational learning for parameter exploration, because it naturally learns a distribution over neural network parameters. Given a family of distributions, usually Gaussians, $\mathcal{Q}$, generalized variational learning aims to find the following approximate posterior over parameters $\boldsymbol{\theta}$,

$$q^*(\boldsymbol{\theta}) = \arg\min_{q \in \mathcal{Q}} \sum_{i \in \mathcal{D}} \mathbb{E}_{\boldsymbol{\theta} \sim q}[\ell_i(\boldsymbol{\theta})/|\mathcal{D}|] + \lambda^{-1} \mathbb{D}_{\mathrm{KL}}(q \,\|\, p_0). \tag{2}$$

The objective is similar to standard training, as it also uses general losses [28, 26] but uses a smoothed loss via an expectation over q, as well as a KL regularizer towards a prior p_0. The prior can be chosen as $p_0 \propto \exp\{-\ell_0\}$, where ℓ_0 is a standard regularizer like weight decay to mirror conventional deep learning. The scaling factor $\lambda > 0$ is used to weigh data fit and regularization towards the prior which can be written as an entropy-regularized objective, where λ scales the entropy of q^* in the objective.

Many methods have been proposed to optimize Eq. 2. For example, Graves [17] and Bayes-by-Backprop [5] try to use conventional gradient descent methods to learn the mean and variance of a Gaussian $\mathcal{Q}$. More recently, natural-gradient-based methods have shown more success [27, 34]. In particular, the IVON optimizer [54] has been shown to provide similar performance as methods like AdamW [40] at little overhead in terms of runtime that is mainly due to the sampling implementation.

In IVON, at each step the loss is calculated at a perturbed point $\hat{\boldsymbol{\theta}}$, and not at the current mean $\mathbf{m}$, which corresponds to the point estimate learned by e.g. AdamW. Writing it out as the reparameterization,

$$\hat{\boldsymbol{\theta}} = \mathbf{m} + \sigma \odot \mathbf{z}; \quad \mathbf{z} \sim \mathcal{N}(0, I), \tag{3}$$

shows the similarity to weight-perturbation-based optimization strategies, which have been found to, among others, lead to improved generalization as they avoid sharp minima [66, 13, 44].

The diagonal variance in IVON is defined inversely proportional to the Hessian which is weighted by the effective sample size (ESS) $\lambda > 0$ from Eq. 2. The ESS can be tuned as an inverse temperature; for example, $\lambda = 10^9$ is common. Concretely, the variance is defined as follows:

$$\sigma = 1/\sqrt{\lambda(\mathbf{h} + \delta)}, \tag{4}$$

where $\mathbf{h}$ is the diagonal Hessian initialized to a constant h_0, and $\delta \geq 0$ is a weight-decay scaling. Choosing λ large thus reduces noise, while choosing λ small increases the amount of noise added to the parameters, which in turn increases diversity in predictions. We use this strategy to control exploration in reinforcement learning with verifiable rewards next.

3 Perturbed Parameter Policy Optimization

We now describe Perturbed Parameter Policy Optimization, which introduces exploration in the parameter space as an additional axis of exploration in RLVR by using noisy policies. Importantly, our method is compatible with existing approaches like temperature control [2, 76], clip-higher [72], and entropy regularization [38, 25]. We study three different strategies that are illustrated in Figure 2.

3.1 Batched noising (B3PO and M3PO)

Modern RL codebases such as `verl` [55] and SkyRL [6] decouple generation and gradient updates into separate inference and training engines, periodically synchronizing their weights at the start of every gradient step. Batched noising works well with such implementations, because a single weight perturbation $\widehat{\boldsymbol{\theta}}$ is sampled from the IVON posterior (Eq. 3) and synced with the rollout engine. B3PO is analogous to step-level temperature control [76], but acting in the weight space instead, and can be interpreted as a one Monte-Carlo sample approximation of Eq. 2 using Eq. 1 as the loss.

B3PO could have similar issues as standard GRPO if the single weight perturbation per step does not perform well (or produces homogeneous rollouts). One way to address this issue is to use multiple

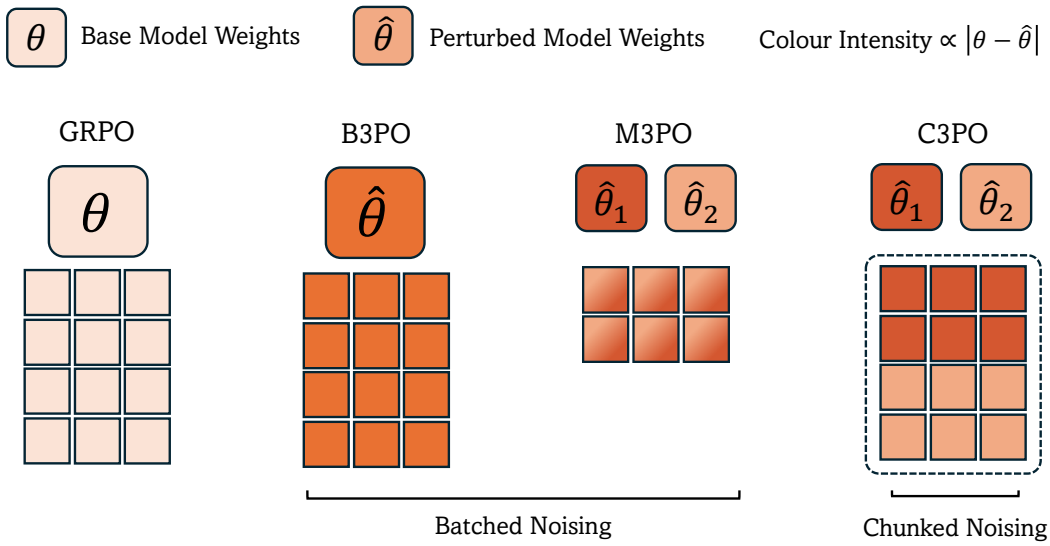


Figure 2: **Overview of noising strategies.** Squircles denote model weights (base $\boldsymbol{\theta}$ or perturbed $\hat{\boldsymbol{\theta}}$), while squares denote rollouts; each column corresponds to a rollout group for a different prompt. Color intensity denotes the amount of noise added. **GRPO** generates all G rollouts using a single shared model $\boldsymbol{\theta}$, with diversity arising only from sampling. **B3PO** samples a single perturbation $\hat{\boldsymbol{\theta}}$ per batch and reuses it across all prompts and groups. **M3PO** accumulates the loss over M perturbations to stabilize training. **C3PO** partitions rollouts across N perturbations and applies GRPO to the aggregated set, which now contains rollouts from multiple model samples.

MC samples for the expectation in Eq. 2, i.e., sampling M perturbations from the IVON posterior q. For each, we generate rollouts over the entire prompt batch and compute the GRPO loss in Eq. 1. The posterior is updated once at the end, with gradients averaged across all M perturbations. We refer to this strategy as M3PO. The pseudo-code for both methods is shown in Algorithm 1.

3.2 Chunked noising (C3PO)

GRPO's learning signal stems from within-group reward differences, but M3PO calculates advantages separately for each group instead of using one large group across samples that, intuitively, should be more diverse as it uses rollouts from different models to calculate advantages. We address this issue in C3PO, where we sample N weight samples using IVON, say with groups of size G/N, and then calculate advantages over the full group of G responses, whose diversity now reflects N distinct weight-space samples. The full procedure is given in Algorithm 2.

Empirically, we found that Seq-MIS [35], which uses sequence-level importance sampling and masking, was important to stabilize training and ensure convergence (see Appendix I for an ablation). A possible reason for this observation is that the policies that are used for gradient estimation differ from the N rollout generators, akin to training – inference mismatches in RLVR literature [70]. Caching and replaying the noise could be an alternative but comes at increased memory cost.

4 Experiments

4.1 Experimental setup

We use two base models, namely, `Olmo-3-1025-7B` [46] and `Qwen2.5-Math-7B` [69] for our experiments. We follow a two-stage post-training pipeline, beginning with a warm-start SFT followed by RLVR [31, 18]. For the SFT phase, we use a $\approx$ 2M-example subset of the Llama-Nemotron Post-Training Dataset [4] containing DeepSeek-R1 [18] responses covering math, code, general reasoning, and instruction following. We train for two epochs using IVON [54] for both Olmo3 and Qwen2.5-Math. This phase equips the model with initial reasoning capabilities. While it also produces a learned noise distribution that could be reused as an initialization for RLVR, throughout

	Method	AIME '24	AIME '25	AIME '26	MATH-500	AMC	Minerva	Average
Olmo	SFT	$16.67_{2.12}$	$21.25_{2.26}$	$12.50_{2.13}$	$78.08_{0.68}$	$51.25_{3.53}$	$30.51_{1.10}$	$35.04_{0.51}$
	GRPO	$25.42_{2.92}$	$22.08_{2.48}$	**19.17**$_{3.45}$	$84.93_{0.47}$	$69.38_{3.50}$	$36.99_{0.94}$	$42.99_{0.51}$
	GRPO (EntReg)	$23.75_{2.35}$	$23.33_{2.45}$	$17.92_{2.68}$	$84.60_{0.58}$	$68.44_{2.89}$	$36.58_{0.95}$	$42.44_{0.53}$
	GRPO (Polaris)	$26.25_{2.65}$	$20.83_{2.63}$	$15.83_{3.47}$	$84.88_{0.63}$	$64.38_{2.94}$	$36.63_{0.99}$	$41.47_{0.54}$
	GRPO (KL-Cov)	$24.17_{2.23}$	$24.58_{2.78}$	$16.25_{2.93}$	$83.88_{0.68}$	$66.56_{3.45}$	**38.28**$_{0.94}$	$42.29_{0.49}$
	B3PO	$25.42_{2.76}$	$25.00_{2.34}$	$15.00_{2.34}$	$83.30_{0.58}$	**70.31**$_{3.01}$	$37.36_{0.92}$	$42.73_{0.48}$
	M3PO	$25.83_{3.48}$	$24.58_{2.53}$	$17.92_{2.64}$	**85.05**$_{0.51}$	$69.38_{3.27}$	$36.58_{0.83}$	$43.22_{0.49}$
	C3PO	**27.50**$_{2.91}$	**26.25**$_{2.58}$	**19.17**$_{2.94}$	$84.53_{0.48}$	$70.00_{3.43}$	$36.81_{0.93}$	**44.04**$_{0.46}$
Qwen	SFT	$14.58_{2.50}$	$20.00_{3.12}$	$12.50_{2.39}$	$65.25_{1.93}$	$59.38_{2.98}$	$24.59_{1.15}$	$32.72_{0.36}$
	GRPO	$24.17_{3.25}$	$25.42_{3.51}$	$20.42_{3.05}$	$88.20_{0.96}$	$70.63_{3.01}$	$43.38_{1.20}$	$45.37_{0.35}$
	GRPO (EntReg)	$23.33_{3.34}$	$24.17_{3.67}$	$15.83_{3.19}$	$86.28_{0.99}$	$65.63_{3.06}$	$42.65_{1.18}$	$42.98_{0.34}$
	GRPO (Polaris)	$22.92_{3.46}$	$26.67_{3.35}$	$21.25_{3.14}$	$87.75_{1.02}$	$72.81_{2.97}$	$43.29_{1.26}$	$45.78_{0.32}$
	GRPO (KL-Cov)	$22.92_{3.13}$	$24.58_{3.58}$	$20.42_{3.20}$	$87.75_{0.94}$	$71.88_{2.63}$	$43.20_{1.24}$	$45.12_{0.37}$
	B3PO	$23.75_{3.14}$	$25.42_{3.65}$	$24.17_{3.19}$	$87.68_{0.96}$	$72.81_{2.62}$	$43.34_{1.21}$	$46.19_{0.37}$
	M3PO	$23.75_{3.11}$	$25.42_{3.44}$	$24.17_{3.29}$	**88.58**$_{0.95}$	**73.44**$_{3.00}$	$43.01_{1.20}$	$46.39_{0.36}$
	C3PO	**25.00**$_{3.21}$	**27.08**$_{3.12}$	**26.25**$_{2.93}$	$88.10_{0.93}$	$71.25_{2.67}$	**43.61**$_{1.22}$	**46.88**$_{0.38}$

Table 1: **A comparison of all methods across common mathematical reasoning benchmarks**. Methods with multiple noise samples outperform GRPO on average. Both models benefit most from C3PO's increased group diversity. The largest gains are on the harder AIME problems. We report the unbiased Pass@1 across 8 samples, and its standard error.

the main results, we initialize the IVON Hessian $\mathbf{h}$ from scratch with a constant h_0 rather than loading the SFT optimizer state to keep our results applicable to any off-the-shelf model and allow us to isolate the benefits of a learned initialization in Section 5.4.

Unless specified otherwise, we train our models using GRPO [52] on DAPO-MATH-17k [72], which contains approximately $17,000$ math problems with ground-truth answers. We evaluate on six widely used mathematical reasoning benchmarks: AIME 2024–26, AMC 2023, MATH-500 [20], and Minerva [32]. Across all experiments, we report the mean and standard error for Pass@1, calculated across $K=8$ rollouts using the unbiased estimator given by Chen et al. [7].

We compare 3PO to three action-space exploration baselines: adding an entropy regularization term (EntReg), raising sampling temperature when entropy collapses beyond a threshold (Polaris [2]) and applying a KL penalty to tokens with a high covariance between probability and advantage (KL-Cov [10]). Since these methods modify the token distribution while 3PO perturbs the weights, they are complementary rather than competing, and can be applied alongside 3PO.

For all methods, each prompt receives 16 rollouts per batch. For M3PO, we report an equal-compute comparison with AdamW, using $M=4$ MC samples per batch to generate rollouts and compute advantages with $G=4$ at a time. The resulting policy update is thus calculated on a total of $G\times M=16$ rollouts. For C3PO, we sample $N=4$ distinct policies per batch and accumulate 4 rollouts for each in a rollout buffer, computing advantages on 16 rollouts in total. We use $\lambda = 10^9$ for Olmo3 and $\lambda = 10^{10}$ for Qwen2.5-Math. We sweep these hyperparameters in Section 5 and Appendices C and D, and describe our training configuration in detail in Appendix A.

4.2 Main results

Parameter-space exploration outperforms action-space baselines. We report benchmark scores across all methods and models in Table 1. All 3PO variants outperform action-space exploration baselines on average. We hypothesize that these gains are driven by parameter-space exploration discovering fewer invalid rollouts than their action-space counterparts, since trajectories come from different plausible LLM policies. We probe this in Section 4.3. Notably, performance gains are concentrated on the harder AIME benchmarks, outperforming GRPO by up to 4.17 points on Olmo3 and 5.83 points on Qwen2.5-Math. We see similar trends on the relatively harder code generation task in Section 4.4. C3PO's gain over GRPO is robust to run-to-run variance, measured across 3 seeds (Appendix G).

On average, using multiple model samples outperforms single-model methods. B3PO outperforms GRPO on individual benchmarks such as AIME-2025, AMC, and Minerva, indicating that

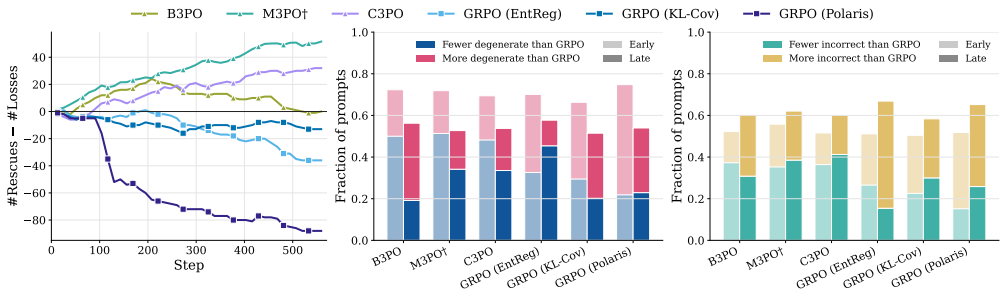


Figure 3: **(Left) Net zero-advantage groups rescued over training.** 3PO rescues more groups than GRPO and action-space baselines, with multiple perturbations yielding consistent improvements. **(Middle) Paired degeneracy rates vs GRPO.** Polaris degenerates due to flattening of the token distribution at high temperature. **(Right) Paired incorrect rollout rates vs GRPO.** While EntReg escapes degeneracy in late training, it buys diversity mostly through incorrect rollouts. Early and late refer to the early (first 33%) and late training steps (last 33%), respectively.

simply switching optimizers from AdamW to IVON during RLVR could buy modest improvements in some cases. However, its average performance does not improve meaningfully beyond that of the standard GRPO, possibly due to high noise sample variance and limited group diversity. While M3PO uses M MC samples to reduce noise sample variance and achieves moderate performance gains over B3PO for both model families, its effectiveness is limited due to a proportional decrease in its group size G. We show that lifting the equal-compute constraint further improves the performance of M3PO in Section 5.2. C3PO has the highest average performance in both model families, likely due to the increase in group diversity due to using multiple model samples. We show results with Olmo3 in the next sections and report results using Qwen2.5-Math in Appendices B to D.

4.3 Does 3PO really explore more?

In this section, we evaluate whether 3PO can increase exploration during RLVR. However, token-level entropy alone is not sufficient to measure this, since declining entropy could indicate two opposite behaviors. On the one hand, entropy can decrease due to a focus on high-probability – high-advantage tokens [10] which could be seen as a form of exploitation. On the other hand, low entropy can equally signal a policy learning incorrect trajectories [8]. We observe this phenomenon in Figure 7 when noise is scaled too far and where the steep decline in C3PO's entropy during training (Figure 6) is largely uninformative. Measures such as pass@k may also be misleading, as they measure the ability of a model to explore post-RLVR at inference, which is often traded for a higher pass@1 via the entropy – reward relationship mentioned above in a form of exploitation.[1]

To measure exploration during RLVR, we pair each method against GRPO and study how often each method "rescues" zero-advantage groups. A prompt is "rescued" if a method produces at least one correct rollout where GRPO produced none, and "lost" in the opposite case. A successful training-time exploration method should rescue more dead groups than it loses, thus sustaining a learning signal. We also measure the number of degenerate and incorrect rollouts produced by each method during training as compared to GRPO. For each prompt, a method "wins" over GRPO if it produces fewer malformed/incorrect rollouts during training. A degenerate (or malformed) rollout is one that does not produce an extractable answer at all, often repeating tokens until the generation limit, and yields no learning signal.

3PO rescues more zero-advantage groups than baselines. Figure 3 (Left) shows that 3PO tends to rescue more groups than it loses, while action-space methods fail to do so. Intuitively, 3PO's trajectories come from LLM policies that are likely under the posterior and which are less likely to generate invalid trajectories than the undirected exploration of action-space methods. Using IVON further improves the quality of these trajectories by learning the distribution over policies jointly during training (Eq. 4). Polaris loses significantly more groups than it rescues and its gap to GRPO widens during training. Although B3PO rescues dead groups in early training, its advantage

[†]GRPO's group is subsampled to 4 rollouts for its pairing with M3PO.

[1]Still, sampling from the learned posterior could yield inference-time improvements [9].

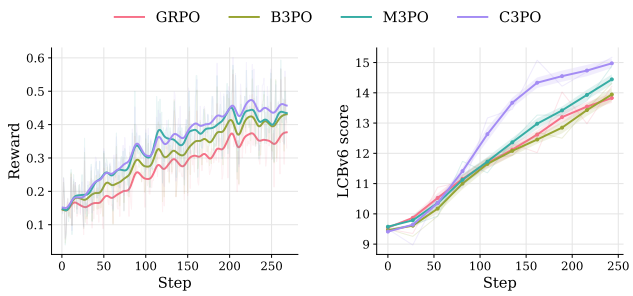


Figure 4: **Reward and accuracy curves for code generation.** All 3PO methods learn faster and achieve higher rewards than GRPO. C3PO reaches GRPO's final LCB-v6 score at step 162.

Method	ESS (λ)	Avg. Math Pass@1
B3PO	10^8	41.90 ± 0.51
	10^9	$\mathbf{42.73 \pm 0.48}$
	10^{10}	42.19 ± 0.53
M3PO	10^8	42.93 ± 0.51
	10^9	43.22 ± 0.49
	10^{10}	$\mathbf{43.59 \pm 0.51}$
C3PO	10^8	0.00 (collapsed)
	10^9	$\mathbf{44.04 \pm 0.46}$
	10^{10}	41.42 ± 0.54

Table 2: **Effects of scaling λ for Olmo3:** 10^9 is a good default value, while $\lambda = 10^8$ hurts performance. M3PO benefits from a higher λ.

diminishes in the later training steps. Conversely, M3PO and C3PO continue to rescue zero-advantage groups throughout training, highlighting the benefits of using multiple parameter perturbations.

Undirected action-space methods mostly yield malformed or incorrect rollouts. Action-space baselines fail to improve over GRPO on malformed and incorrect rollout generation rates. Polaris increases sampling temperature when entropy collapses beyond a threshold. Since temperature only flattens the token distribution (Figure 1), it consistently yields rollouts with no extractable answer (Figure 3 (Middle)). EntReg exhibits a different failure mode. Although its rollouts stay well-formed, even beating 3PO in late-training, its higher entropy is spent on mostly incorrect rollouts (Figure 3 (Right)). KL-Cov produces more degenerate rollouts than GRPO, while its incorrect rate is largely unchanged. B3PO also tends to produce degenerate rollouts later in training which might explain its diminishing returns in rescuing zero-advantage groups. M3PO and C3PO are the only methods to consistently produce fewer malformed and incorrect rollouts than GRPO which suggests that parameter-space exploration does produce more useful rollouts for RLVR.

4.4 Results for code generation

In this section, we study the effect of parameter-space exploration beyond mathematical reasoning by testing our methods on code generation tasks. Specifically, given a problem statement, the LLM must produce code that passes all hidden test cases. We train Olmo3 on the CodeR1-12K dataset [36], which contains verified coding problems from LeetCode and TACO [33]. We use the same hyperparameters as math detailed in Section 4.1 and utilize SandboxFusion [15] for executing untrusted LLM code. We evaluate all models on LiveCodeBench-v6 [23], which contains problems released between January 2025 and April 2025.

As seen in Figure 4, there is a consistent gap in the reward curves between our 3PO methods and GRPO. C3PO is the best-performing method in this regime, reaching an LCBv6 score of 15.17, compared to 13.77 for GRPO. C3PO's lead over GRPO emerges around step 100 and persists for the rest of training, which highlights the effectiveness of increased group diversity through our proposed chunked noising approach. M3PO narrowly outperforms B3PO and GRPO, but is probably limited by its lower group diversity due to a proportional decrease in its group size (Section 5.2). Similar to math, both M3PO and C3PO are more sample efficient and reach GRPO's final reward within the first 50% of the training steps.

Notably, 3PO methods reach a noticeably higher reward than GRPO on code generation as compared to mathematical reasoning. We hypothesize that this is because code generation is a harder task for the pre-RL model, as seen by the SFT model's performance on the LCBv6 benchmark (9.5%) compared to the math benchmarks (35.04%). Thus, increased exploration with 3PO could be more valuable here and aid in discovering high-reward trajectories. We observed a similar trend on the harder AIME benchmarks for math in Section 4.2, which suggests that parameter-space exploration is particularly effective for harder tasks where the pre-RL model is less capable and more prone to generating malformed or incorrect rollouts.

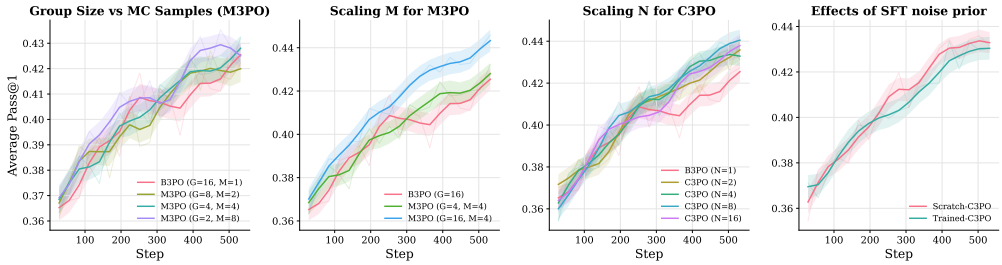


Figure 5: **(a) Equal-compute comparison:** Benefits of increasing M are offset by reducing G. **(b) Scaling M:** The advantages of more MC samples are fully realized without an equal-compute constraint. **(c) Scaling N:** $N > 1$ improves performance, further scaling is incremental. **(d) Learned noise prior:** Initializing $\mathbf{h}$ from the learned SFT prior has negligible effects.

5 Further analysis

We now ablate several components of our proposed approach. Specifically, we study how λ should be scaled (Section 5.1), how many MC samples (Section 5.2) and chunks (Section 5.3) are needed, and the impact of Hessian initialization (Section 5.4). We conclude with practitioner recommendations.

5.1 Impact of scaling ESS (λ)

The ESS is a crucial hyperparameter in our setup that controls the amount of noise added to the weights, and needs to be tuned to achieve good performance [11, 9]. We ablate $\lambda \in \{10^8, 10^9, 10^{10}\}$.

We find decreasing λ too far to 10^8 always hurts performance, and completely collapses training for C3PO. This is an expected result because C3PO samples multiple models in the group rollout, but the advantages are still calculated by a single model. Adding too much noise might stray the model too far. Conversely, increasing λ to 10^{10} has more nuanced results. For the batched noising methods B3PO and M3PO, it behaves similarly to $\lambda = 10^9$ with a steeper entropy decline than $\lambda = 10^8$. $\lambda = 10^{10}$ even slightly outperforms $\lambda = 10^9$ on average for M3PO (Table 2).

However, $\lambda = 10^{10}$ plateaus at a lower ceiling than the 10^9 C3PO run, as seen from its entropy and reward curves. A possible explanation is that M3PO has lower variance in the gradient update due to the averaging, and a higher λ reduces variance further by reducing the amount of noise added during optimization. On the other hand, C3PO benefits from increased noise due to greater group diversity. We present extended results for this ablation in Appendix C. It is also plausible that adaptive noise schedules could improve training more. We leave a detailed exploration of this for future work.

5.2 How does the number of MC samples M interplay with group size G?

M3PO comes at an M-fold compute cost because rollouts must be generated M times, once per noise sample (see Table 3). We offset this by shrinking the per-perturbation group to $G = 16/M$, holding the total rollout count per step fixed. This equal-compute correction introduces an important tradeoff: as M grows, each group shrinks, reducing the diversity in the GRPO group and the chances of sampling a high-reward response per prompt [22]. Figure 5(a) plots the result of sweeping $M \in \{1, 2, 4, 8\}$ under this constraint. The effects of variance reduction from increasing M are roughly cancelled by the reduced group diversity from a smaller G, and vice versa.

Without the equal-compute correction (Figure 5(b)), M3PO steadily and substantially outperforms the equal-compute run, as well as B3PO with $G = 16$. This indicates that reduced gradient-update variance does benefit M3PO when compute is unconstrained, and might explain its advantage over B3PO in Table 1. We provide further results in Appendix D.

5.3 What chunk size N is required?

We ablate the chunk size $N \in \{1, 2, 4, 8, 16\}$ to analyze how many models are required to see the benefit of increased diversity. Note that $N = 1$ reduces to B3PO.

Even modest group diversification yields a consistent benefit, as all runs with $N \geq 2$ outperform the single-perturbation baseline (Figure 5(c)). Scaling beyond $N = 2$ improves performance in earlier training steps but yields smaller gains in final scores. We plot the reward and entropy curves in Appendix E.

5.4 Does learning the noise improve performance?

The experiments so far initialize the IVON Hessian $\mathbf{h}$ to a constant h_0 at the start of RLVR, discarding the optimizer state obtained during the warm-start SFT phase. A natural hypothesis is that loading this learned Hessian, which encodes per-parameter variance information from the SFT phase, would yield a better posterior and improve downstream performance. We test this hypothesis by comparing two C3PO runs on Olmo3 that share the same SFT checkpoint. Note that $\mathbf{h}$ continues to be updated during RLVR in both cases, and replacing it with fixed isotropic noise of matched magnitude drops C3PO below GRPO (Appendix H).

Figure 5(d) shows that the two configurations track each other closely throughout training, suggesting that a learned distribution from SFT does not provide a clear advantage. This may be because the SFT phase is relatively short, and the loaded Hessian is nearly isotropic at the end. Concretely, $< 8\%$ of the values lie outside the $[0.9h_0, 1.1h_0]$ range. Learning the distribution over a longer SFT phase or from pre-training itself could be more beneficial, but we need to leave this experiment for future work due to computational constraints. Nonetheless, our findings transfer to any base model without additional training. We provide detailed results in Appendix F.

Practical recommendations. Based on our experiments so far, we conclude that using multiple parameter samples per step outperforms a single perturbation. Due to the G–M tradeoff, we recommend using C3PO with $N = 2-4$ perturbations and $\mathbf{h}$ initialized to a constant h_0. $\lambda = 10^9$ is a good default value, but must be tuned per model. Further hyperparameter details are in Appendix A.

6 Conclusion

We introduce Perturbed Parameter Policy Optimization (3PO), a family of parameter-space exploration strategies for RLVR. By sampling weights from a learned posterior at rollout time, 3PO provides an additional lever for exploration. Across our experiments, drawing multiple model samples per gradient step consistently outperforms action-space baselines. We show that our proposed methods sustain a learning signal by rescuing more zero-advantage groups than action-space baselines. While M3PO tends to reduce noise sample variance, it either requires more compute or suffers from reduced group diversity. Our chunked noising approach, C3PO, has the best average performance across both model families. We find it useful to tune the amount of added noise through λ to control the extent of exploration during RLVR. While a noise prior can be learned before RL, a constant initialization offers similar improvements, making our methods compatible as a drop-in replacement for checkpoints lacking variational training. Overall, our work shows that parameter-space exploration can be used to improve reinforcement learning for large language models.

7 Limitations

We expect the benefits of 3PO to grow with model scale but our work is limited by compute. Gan and Isola [16] show that large LLMs' pretraining neighborhoods are densely populated with task-specific experts. Thus, larger models should tolerate a lower λ before training collapses, allowing for more exploration: an exciting direction for future work that we were unable to explore. In addition, M3PO and C3PO currently have $\approx 1.5\times$ higher wall-clock times than GRPO (Appendix A), despite equivalent computational requirements in pseudocode. This stems from processing small prompt batches between vLLM model updates, which reduces parallelism benefits [30]. Overall, modern RL code stacks all operate using a single model, and it is often difficult to implement changes for variational learning. We also did not explore scaling the number of MC samples and rollouts further for test-time scaling due to limited resources. As with other IVON-based methods, performance is sensitive to the ESS λ, which must be tuned. Finally, using more directed exploration through structured, non-diagonal posterior estimators [43] instead of IVON, or alternative sampling strategies to Monte-Carlo sampling are a promising direction for future work.

Acknowledgements and Disclosures

Funding statement. We gratefully acknowledge the support from (1) the Ministry of Education and Science of Bulgaria (support for INSAIT, part of the Bulgarian National Roadmap for Research Infrastructure), (2) the hessian.AI Service Center (funded by the Federal Ministry of Research, Technology and Space, BMFTR, grant no. 16IS22091), (3) the hessian.AI Innovation Lab (funded by the Hessian Ministry for Digital Strategy and Innovation, grant no. S-DIW04/0013/003), and (4) the German Federal Ministry of Research, Technology, and Space and the Hessian Ministry of Higher Education, Research, Science, and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

Competing interests. The authors declare no competing interests.

References

[1] Susan Amin, Maziar Gomrokchi, Harsh Satija, Herke van Hoof, and Doina Precup. A survey of exploration methods in reinforcement learning. *CoRR*, abs/2109.00157, 2021. URL https://arxiv.org/abs/2109.00157.

[2] Chenxin An, Zhihui Xie, Xiaonan Li, Lei Li, Jun Zhang, Shansan Gong, Ming Zhong, Jingjing Xu, Xipeng Qiu, Mingxuan Wang et al. Polaris: A post-training recipe for scaling reinforcement learning on advanced reasoning models, 2025. URL https://hkunlp.github.io/blog/2025/Polaris.

[3] Bizhe Bai, Xinyue Wang, Peng Ye, and Tao Chen. Learning to explore with parameter-space noise: A deep dive into parameter-space noise for reinforcement learning with verifiable rewards. *CoRR*, abs/2602.02555, 2026. doi: 10.48550/ARXIV.2602.02555. URL https://doi.org/10.48550/arXiv.2602.02555.

[4] Akhiad Bercovich, Itay Levy, Izik Golan, Mohammad Dabbah, Ran El-Yaniv, Omri Puny, Ido Galil, Zach Moshe, Tomer Ronen, Najeeb Nabwani et al. Llama-nemotron: Efficient reasoning models, 2025. URL https://arxiv.org/abs/2505.00949.

[5] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1613–1622, Lille, France, 07–09 Jul 2015. PMLR. URL https://proceedings.mlr.press/v37/blundell15.html.

[6] Shiyi Cao, Sumanth Hegde, Dacheng Li, Tyler Griggs, Shu Liu, Eric Tang, Jiayi Pan, Xingyao Wang, Akshay Malik, Graham Neubig et al. Skyrl-v0: Train real-world long-horizon agents via reinforcement learning, 2025. URL https://github.com/NovaSky-AI/SkyRL.

[7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman et al. Evaluating large language models trained on code. *CoRR*, abs/2107.03374, 2021. URL https://arxiv.org/abs/2107.03374.

[8] Peter Chen, Xiaopeng Li, Ziniu Li, Wotao Yin, Xi Chen, and Tianyi Lin. Exploration vs exploitation: Rethinking RLVR through clipping, entropy, and spurious reward. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=sE8DCSJTzd.

[9] Bai Cong, Nico Daheim, Yuesong Shen, Rio Yokota, Mohammad Emtiyaz Khan, and Thomas Möllenhoff. Improving LoRA with variational learning, 2025. URL https://arxiv.org/abs/2506.14280.

[10] Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen Fan, Huayu Chen, Weize Chen et al. The entropy mechanism of reinforcement learning for reasoning language models. *CoRR*, abs/2505.22617, 2025. doi: 10.48550/ARXIV.2505.22617. URL https://doi.org/10.48550/arXiv.2505.22617.

[11] Nico Daheim, Clara Meister, Thomas Möllenhoff, and Iryna Gurevych. Uncertainty-aware decoding with minimum bayes risk. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=hPpyUv1XyQ.

[12] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A survey on policy search for robotics. *Found. Trends Robotics*, 2(1-2):1–142, 2013. doi: 10.1561/2300000021. URL https://doi.org/10.1561/2300000021.

[13] Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=6Tm1mposlrM.

[14] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Matteo Hessel, Ian Osband, Alex Graves, Volodymyr Mnih, Rémi Munos, Demis Hassabis et al. Noisy networks for exploration. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL https://openreview.net/forum?id=rywHCPkAW.

[15] Bytedance-Seed Foundation-Code-Team, Yao Cheng, Jianfeng Chen, Jie Chen, Li Chen, Liyu Chen, Wentao Chen, Zhengyu Chen, Shijie Geng, Aoyan Li et al. FullStack Bench: Evaluating LLMs as full stack coders, 2025. URL https://arxiv.org/abs/2412.00535.

[16] Yulu Gan and Phillip Isola. Neural thickets: Diverse task experts are dense around pretrained weights. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=92oF5bU4cU.

[17] Alex Graves. Practical variational inference for neural networks. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper_files/paper/2011/file/7eb3c8be3d411e8ebfab08eba5f49632-Paper.pdf.

[18] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi et al. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081):633–638, 2025. doi: 10.1038/S41586-025-09422-Z. URL https://doi.org/10.1038/s41586-025-09422-z.

[19] Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen et al. Skywork open reasoner 1 technical report. *CoRR*, abs/2505.22312, 2025. doi: 10.48550/ARXIV.2505.22312. URL https://doi.org/10.48550/arXiv.2505.22312.

[20] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In Joaquin Vanschoren and Sai-Kit Yeung, editors, *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract-round2.html.

[21] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=rygGQyrFvH.

[22] Jian Hu, Mingjie Liu, Ximing Lu, Fang Wu, Zaid Harchaoui, Shizhe Diao, Yejin Choi, Pavlo Molchanov, Jun Yang, Jan Kautz et al. BroRL: Scaling reinforcement learning via broadened exploration. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=KmS7pdFBEh.

[23] Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=chfJJYC3iL.

[24] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *J. Mach. Learn. Res.*, 11:1563–1600, 2010. doi: 10.5555/1756006.1859902. URL https://dl.acm.org/doi/10.5555/1756006.1859902.

[25] Yuxian Jiang, Yafu Li, Guanxu Chen, Dongrui Liu, Yu Cheng, and Jing Shao. Rethinking entropy regularization in large reasoning models. *CoRR*, abs/2509.25133, 2025. doi: 10.48550/ARXIV.2509.25133. URL https://doi.org/10.48550/arXiv.2509.25133.

[26] Mohammad Emtiyaz Khan and Håvard Rue. The bayesian learning rule. *Journal of Machine Learning Research*, 24(281):1–46, 2023. URL http://jmlr.org/papers/v24/22-0291.html.

[27] Mohammad Emtiyaz Khan, Didrik Nielsen, Voot Tangkaratt, Wu Lin, Yarin Gal, and Akash Srivastava. Fast and scalable Bayesian deep learning by weight-perturbation in Adam. In Jennifer G. Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Proceedings of Machine Learning Research, pages 2616–2625. PMLR, 2018. URL http://proceedings.mlr.press/v80/khan18a.html.

[28] Jeremias Knoblauch, Jack Jewson, and Theodoros Damoulas. Generalized variational inference: Three arguments for deriving new posteriors. *arXiv*, 1904.02063, 2019. URL https://arxiv.org/abs/1904.02063.

[29] Nate Kohl and Peter Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *Proceedings of the 2004 IEEE International Conference on Robotics and Automation, ICRA 2004, April 26 - May 1, 2004, New Orleans, LA, USA*, pages 2619–2624. IEEE, 2004. doi: 10.1109/ROBOT.2004.1307456. URL https://doi.org/10.1109/ROBOT.2004.1307456.

[30] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In Jason Flinn, Margo I. Seltzer, Peter Druschel, Antoine Kaufmann, and Jonathan Mace, editors, *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*, pages 611–626. ACM, 2023. doi: 10.1145/3600006.3613165. URL https://doi.org/10.1145/3600006.3613165.

[31] Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James Validad Miranda, Alisa Liu, Nouha Dziri, Xinxi Lyu et al. Tulu 3: Pushing frontiers in open language model post-training. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=i1uGbfHHpH.

[32] Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay V. Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo et al. Solving quantitative reasoning problems with language models. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/18abbeef8cfe9203fdf9053c9c4fe191-Abstract-Conference.html.

[33] Kaixin Li. Verified TACO problems, 2024. URL https://huggingface.co/datasets/likaixin/TACO-verified.

[34] Wu Lin, Mark Schmidt, and Mohammad Emtiyaz Khan. Handling the positive-definite constraint in the Bayesian learning rule. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, pages 6116–6126. PMLR, 2020. URL http://proceedings.mlr.press/v119/lin20d.html.

[35] Jiacai Liu, Yingru Li, Yuqian Fu, Jiawei Wang, Qian Liu, and Zhuo Jiang. When speed kills stability: Demystifying RL collapse from the training-inference mismatch. https://richardli.xyz/rl-collapse, September 2025.

[36] Jiawei Liu and Lingming Zhang. Code-R1: Reproducing R1 for code with reliable rewards. https://github.com/ganler/code-r1, 2025.

[37] Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong. ProRL: Prolonged reinforcement learning expands reasoning boundaries in large language models. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=YPsJha5HXQ.

[38] Zhuang Liu, Xuanlin Li, Bingyi Kang, and Trevor Darrell. Regularization matters in policy optimization - an empirical study on continuous control. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=yr1mzrH3IC.

[39] Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding r1-zero-like training: A critical perspective. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=5PAF7PAY2Y.

[40] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

[41] Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa et al. Deepscaler: Surpassing o1-preview with a 1.5b model by scaling rl. https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2, 2025. Notion Blog.

[42] Mathematical Association of America. American Invitational Mathematics Examination, 2026. URL https://maa.org/maa-invitational-competitions/.

[43] Adrian Robert Minut, Nico Daheim, Marco Miani, Mohammad Emtiyaz Khan, Wu Lin, and Thomas Möllenhoff. Soap-bubbles: Structured weight uncertainty for neural networks, 2026. URL https://arxiv.org/abs/2606.23357.

[44] Thomas Möllenhoff and Mohammad Emtiyaz Khan. SAM as an optimal relaxation of Bayes. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=k4fevFqSQcX.

[45] Michael Noukhovitch, Shengyi Huang, Sophie Xhonneux, Arian Hosseini, Rishabh Agarwal, and Aaron Courville. Asynchronous RLHF: Faster and more efficient off-policy RL for language models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=FhTAG591Ve.

[46] Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison et al. Olmo 3, 2025. URL https://arxiv.org/abs/2512.13961.

[47] Matthias Plappert, Rein Houthooft, Prafulla Dhariwal, Szymon Sidor, Richard Y. Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL https://openreview.net/forum?id=ByBAl2eAZ.

[48] Thomas Rückstieß, Frank Sehnke, Tom Schaul, Daan Wierstra, Yi Sun, and Jürgen Schmidhuber. Exploring parameter space in reinforcement learning. *Paladyn J. Behav. Robotics*, 1(1):14–24, 2010. doi: 10.2478/S13230-010-0002-4. URL https://doi.org/10.2478/s13230-010-0002-4.

[49] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *CoRR*, abs/1703.03864, 2017. URL http://arxiv.org/abs/1703.03864.

[50] Frank Sehnke, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. Parameter-exploring policy gradients. *Neural Networks*, 23(4):551–559, 2010. doi: 10.1016/J.NEUNET.2009.12.004. URL https://doi.org/10.1016/j.neunet.2009.12.004.

[51] Amrith Setlur, Matthew Y. R. Yang, Charlie Victor Snell, Jeremiah Greer, Ian Wu, Virginia Smith, Max Simchowitz, and Aviral Kumar. e3: Learning to explore enables extrapolation of test-time compute for LLMs. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=aID0dZmMmM.

[52] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *CoRR*, abs/2402.03300, 2024. doi: 10.48550/ARXIV.2402.03300. URL https://doi.org/10.48550/arXiv.2402.03300.

[53] Han Shen. On entropy control in LLM-RL algorithms. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=LqazVN5epT.

[54] Yuesong Shen, Nico Daheim, Bai Cong, Peter Nickl, Gian Maria Marconi, Clement Bazan, Rio Yokota, Iryna Gurevych, Daniel Cremers, Mohammad Emtiyaz Khan et al. Variational learning is effective for large deep networks. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp, editors, *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, Proceedings of Machine Learning Research, pages 44665–44686. PMLR / OpenReview.net, 2024. URL https://proceedings.mlr.press/v235/shen24b.html.

[55] Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. HybridFlow: A flexible and efficient RLHF framework. In *Proceedings of the Twentieth European Conference on Computer Systems, EuroSys 2025, Rotterdam, The Netherlands, 30 March 2025 - 3 April 2025*, pages 1279–1297. ACM, 2025. doi: 10.1145/3689031.3696075. URL https://github.com/verl-project/verl.

[56] Alexander L. Strehl and Michael L. Littman. A theoretical analysis of model-based interval estimation. In *Proceedings of the 22nd International Conference on Machine Learning*, ICML '05, page 856–863, New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1595931805. doi: 10.1145/1102351.1102459. URL https://doi.org/10.1145/1102351.1102459.

[57] Freek Stulp and Olivier Sigaud. Path integral policy improvement with covariance matrix adaptation. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012, Edinburgh, Scotland, UK, June 26 - July 1, 2012*. icml.cc / Omnipress, 2012. URL http://icml.cc/2012/papers/171.pdf.

[58] Yiyou Sun, Yuhan Cao, Pohao Huang, Haoyue Bai, Hannaneh Hajishirzi, Nouha Dziri, and Dawn Song. RL grokking recipe: How does RL unlock and transfer new algorithms in LLMs? In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=CJJ8VxOWbG.

[59] Richard S. Sutton. Generalization in reinforcement learning: Successful examples using sparse coarse coding. In David S. Touretzky, Michael Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8, NIPS, Denver, CO, USA, November 27-30, 1995*, pages 1038–1044. MIT Press, 1995. URL http://papers.nips.cc/paper/1109-generalization-in-reinforcement-learning-successful-examples-using-sparse-coarse-coding.

[60] Richard S. Sutton and Andrew G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press, 1998. ISBN 978-0-262-19398-6. URL http://www.incompleteideas.net/book/first/the-book.html.

[61] Evangelos A. Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *J. Mach. Learn. Res.*, 11:3137–3181, 2010. doi: 10.5555/1756006.1953033. URL https://dl.acm.org/doi/10.5555/1756006.1953033.

[62] Herke van Hoof, Daniel Tanneberg, and Jan Peters. Generalized exploration in policy search. *Mach. Learn.*, 106(9-10):1705–1724, 2017. doi: 10.1007/S10994-017-5657-1. URL https://doi.org/10.1007/s10994-017-5657-1.

[63] Vatsal Venkatkrishna, Indraneil Paul, and Iryna Gurevych. Aletheia: What makes RLVR for code verifiers tick? *Transactions on Machine Learning Research*, 2026. ISSN 2835-8856. URL https://openreview.net/forum?id=3rVrBGp0mr.

[64] Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xiong-Hui Chen, Jianxin Yang, Zhenru Zhang et al. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for LLM reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=yfcpdY4gMP.

[65] Christopher John Cornish Hellaby Watkins. *Learning from Delayed Rewards*. PhD thesis, King's College, Cambridge, 1989.

[66] Dongxian Wu, Shu-Tao Xia, and Yisen Wang. Adversarial weight perturbation helps robust generalization. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/1ef91c212e30e14bf125e9374262401f-Abstract.html.

[67] Fang Wu, Weihao Xuan, Ximing Lu, Zaïd Harchaoui, and Yejin Choi. The invisible leash: Why RLVR may not escape its origin. *CoRR*, abs/2507.14843, 2025. doi: 10.48550/ARXIV.2507.14843. URL https://doi.org/10.48550/arXiv.2507.14843.

[68] Mingqi Wu, Zhihao Zhang, Qiaole Dong, Zhiheng Xi, Jun Zhao, Senjie Jin, Xiaoran Fan, Yuhao Zhou, Huijie Lv, Ming Zhang et al. Reasoning or memorization? unreliable results of reinforcement learning due to data contamination. In Sven Koenig, Chad Jenkins, and Matthew E. Taylor, editors, *Fortieth AAAI Conference on Artificial Intelligence, Thirty-Eighth Conference on Innovative Applications of Artificial Intelligence, Sixteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2026, Singapore, January 20-27, 2026*, pages 33944–33952. AAAI Press, 2026. doi: 10.1609/AAAI.V40I40.40687. URL https://doi.org/10.1609/aaai.v40i40.40687.

[69] An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin et al. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement, 2024. URL https://arxiv.org/abs/2409.12122.

[70] Feng Yao, Liyuan Liu, Dinghuai Zhang, Chengyu Dong, Jingbo Shang, and Jianfeng Gao. Your efficient rl framework secretly brings you off-policy rl training, August 2025. URL https://fengyao.notion.site/off-policy-rl.

[71] Xinhao Yao, Lu Yu, Xiaolin Hu, Fengwei Teng, Qing Cui, Jun Zhou, and Yong Liu. The debate on RLVR reasoning capability boundary: Shrinkage, expansion, or both? A two-stage dynamic view. *CoRR*, abs/2510.04028, 2025. doi: 10.48550/ARXIV.2510.04028. URL https://doi.org/10.48550/arXiv.2510.04028.

[72] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, YuYue, Weinan Dai, Tiantian Fan, Gaohong Liu, Juncai Liu et al. DAPO: An open-source LLM reinforcement learning system at scale. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=2a36EMSSTp.

[73] Rosie Zhao, Alexandru Meterez, Sham M. Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach. Echo chamber: RL post-training amplifies behaviors learned in pretraining. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=dp4KWuSDzj.

[74] Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang et al. Group sequence policy optimization, 2025. URL https://arxiv.org/abs/2507.18071.

[75] Xinyu Zhu, Mengzhou Xia, Zhepei Wei, Wei-Lin Chen, Danqi Chen, and Yu Meng. The surprising effectiveness of negative reinforcement in LLM reasoning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=ftVlLG9cks.

[76] Haomin Zhuang, Yujun Zhou, Taicheng Guo, Yue Huang, Fangxu Liu, Kai Song, and Xiangliang Zhang. Exploring multi-temperature strategies for token- and rollout-level control in RLVR. *CoRR*, abs/2510.08892, 2025. doi: 10.48550/ARXIV.2510.08892. URL https://doi.org/10.48550/arXiv.2510.08892.

A Hyperparameter Settings

Across our experiments, we use the `verl` Python library [55] for training and vLLM [30] for efficient inference during rollouts and evaluation. We use an on-policy training setup by performing a single gradient update per prompt batch to avoid performance degradations often associated with batch-online methods [63, 74]. Unless otherwise stated, we evaluate with temperature $\tau = 0.6$, top-$p = 0.95$, top-$k = 50$, and $K = 8$ responses, computing pass@k with the unbiased estimator of Chen et al. [7]. We train on an internal cluster of 8 NVIDIA H200 GPUs with 144 GB of vRAM, and each algorithm takes $24-40$ hours to run. We report the per-step FLOPs and times in Table 3. Despite having near-identical FLOPs cost, our approaches that sample multiple models tend to be $\approx 1.5\times$ slower. This is due to inefficient solutions for sampling from multiple models. Theoretically, the only additional compute over GRPO is the $O(\text{model_size})$ elementwise noise operations from Eq. 3 and the recomputation of prefills for each of the N (or M) perturbations, both comparatively cheap. As expected, removing the equal-compute constraint on M3PO ($G = 16, M = 4$) raises the FLOPs cost by a factor of M.

For Polaris, we increased the sampling temperature from 1.0 to 1.4 when entropy reached 85% of its original value. We experimented with other values like 75% and 50%, but these thresholds either triggered the switch too late and did not allow training to recover, or did not trigger at all. We used an entropy coefficient of $\gamma = 10^{-3}$ for EntReg, and additionally experimented with $\gamma = 5 \times 10^{-3}$ that was too high and completely collapsed training. We use the default hyperparameters swept by the original authors for KL-Cov [10].

The $\approx 1.5\times$ wall-clock overhead is entirely a systems cost stemming from suboptimal RL stacks comprising unfused IVON kernels. The larger cost comes from inefficient multi-model sampling in vLLM. Currently, one must perform N or M inference operations, which reduces the speed gains from vLLM's continuous batching. Implementing more efficient multi-model sampling strategies could significantly reduce this overhead. Moreover, all 3PO variants also converge faster than GRPO (Figures 4 and 6), so their effective cost-to-target could even fall below GRPO once these implementations mature. 3PO also introduces no additional memory complexity over AdamW since IVON's optimizer state consists of a momentum and a diagonal Hessian, occupying the same footprint as AdamW's two moment estimates.

Method	**TFLOPs /step**	**Time (mins) /step**
GRPO	51,719	2.72
B3PO	51,719	2.76
M3PO ($G = 4, M = 4$)	51,860	4.02
C3PO ($N = 4$)	52,140	4.18
M3PO ($G = 16, M = 4$)	206,877	9.63

Table 3: **FLOPs and wall clock time per step for each algorithm.** All methods in our main tables have near-identical FLOPs cost. The $1.5\times$ wall-clock gap is a systems artifact of suboptimal multi-model RL infrastructure.

Warm-start SFT phase. We use the IVON optimizer with a learning rate of 50.0, weight decay 10^{-8}, $\lambda = 10^{10}$, $\beta_1 = 0.9$, $\beta_2 = 0.9999$, $h_0 = 0.001$, and clipping radius 0.001. We filter the Llama-Nemotron Post-Training Dataset [4] to retain DeepSeek-R1 responses with a context length of up to 4096 tokens. We use a cosine-decay learning rate schedule with a 10% linear warmup, decaying to 10% of the initial value by the end of training.

RLVR phase. We use a token-level GRPO formulation [72], with the lower and upper PPO importance-sampling clip bounds both set to 0.2. We use a binary correctness reward $(+1/0)$, a batch size of 32 prompts with 16 rollouts each, and one gradient step per rollout batch. For the AdamW

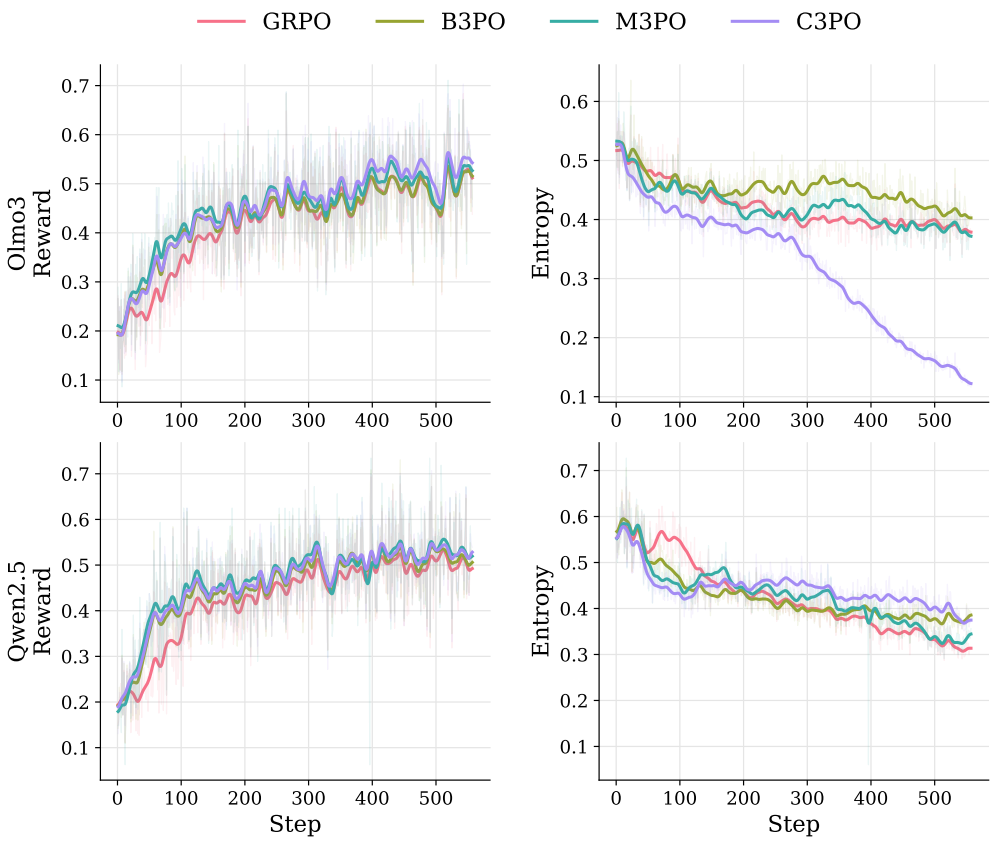


Figure 6: **Reward and entropy curves for math**. All 3PO methods converge faster than GRPO. Olmo3-C3PO's entropy declines rapidly, but all methods have similar profiles for Qwen2.5-Math.

runs, we use a learning rate of 10^{-6}, weight decay 0.1, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. For IVON, we retain the same hyperparameter configuration as in the warm-start SFT phase, except with a learning rate of 1.0 and $\lambda = 10^9$ for Olmo3 and $\lambda = 10^{10}$ for Qwen2.5-Math. We mask out the sequence-level importance weights that lie outside $[0.5, 2.0]$ for C3PO. All runs use a constant learning rate schedule with a 6% linear warmup. Rollouts are sampled at a temperature of 1.0 during training, and allowed to generate a maximum of 4096 tokens for 559 steps for mathematical reasoning and 270 steps for code generation.

B Reward and entropy curves

We plot the reward and entropy curves for our proposed methods and GRPO across both model families in Figure 6. Similar to the observations for code generation, our proposed methods consistently converge faster than GRPO across both model families. For Olmo3, C3PO's entropy declines rapidly, suggesting that it is able to concentrate probability mass on high-reward regions of the solution space. For Qwen2.5-Math, all methods maintain largely similar entropy profiles, indicating that the improved exploration is not simply due to increased randomness in the policy. As mentioned in the main text, we advise against using entropy as a metric for exploration, as it is not always indicative of downstream performance.

C Impact of scaling ESS (λ)

In the main text, we discussed the importance of λ as a crucial control knob for the amount of noise added to the weights, and briefly ablated its effect on our methods for Olmo3. Here, we present detailed results from this ablation. As noted in Section 5.1, $\lambda = 10^9$ is a good default value for

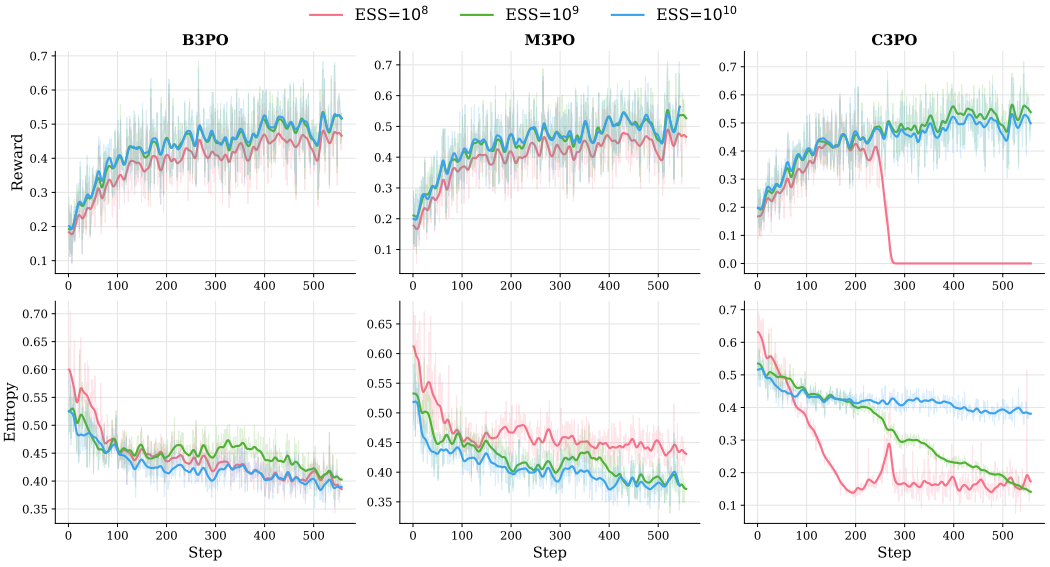


Figure 7: λ **scaling curves across all 3PO methods.** Adding too much noise with small λ hurts performance, while large λ values can make the sampled models too similar.

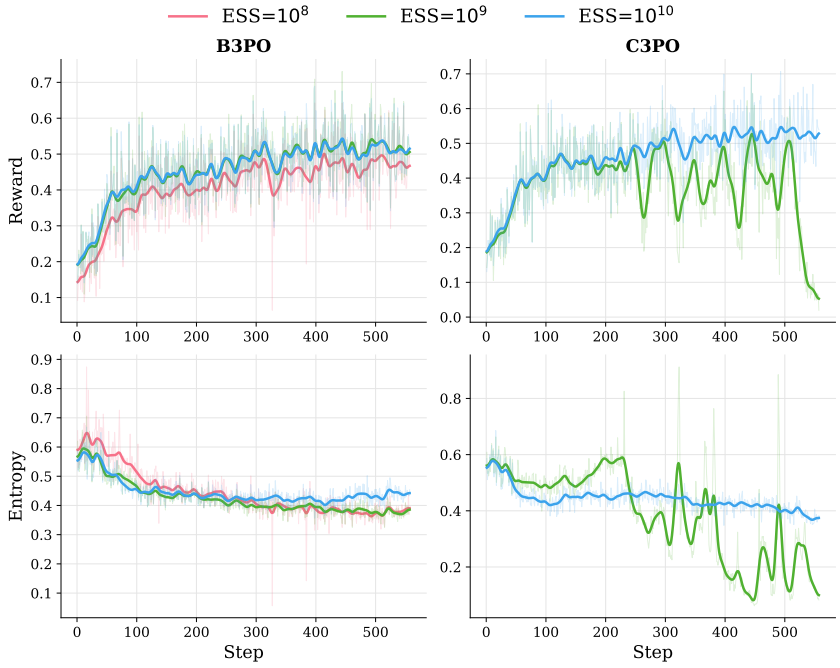


Figure 8: **Effects of scaling λ for Qwen2.5-Math.** Qwen2.5-Math is more sensitive to small λ, where overexploration causes the curves to oscillate. Tuning λ is therefore important.

Olmo3, yielding consistent gains for all three methods. $\lambda = 10^8$ tends to be very unstable, generally underperforming higher values and completely collapsing for C3PO. The entropy curve for C3PO at $\lambda = 10^8$ is evidence for why entropy curves alone are not reliable indicators of exploration.

We additionally reproduce the same comparison on B3PO and C3PO for Qwen2.5-Math in Figure 8. Qwen2.5-Math is less robust to small λ than Olmo3: while B3PO behaves similarly across the two models, C3PO is more sensitive to λ because it samples a fresh model many more times during

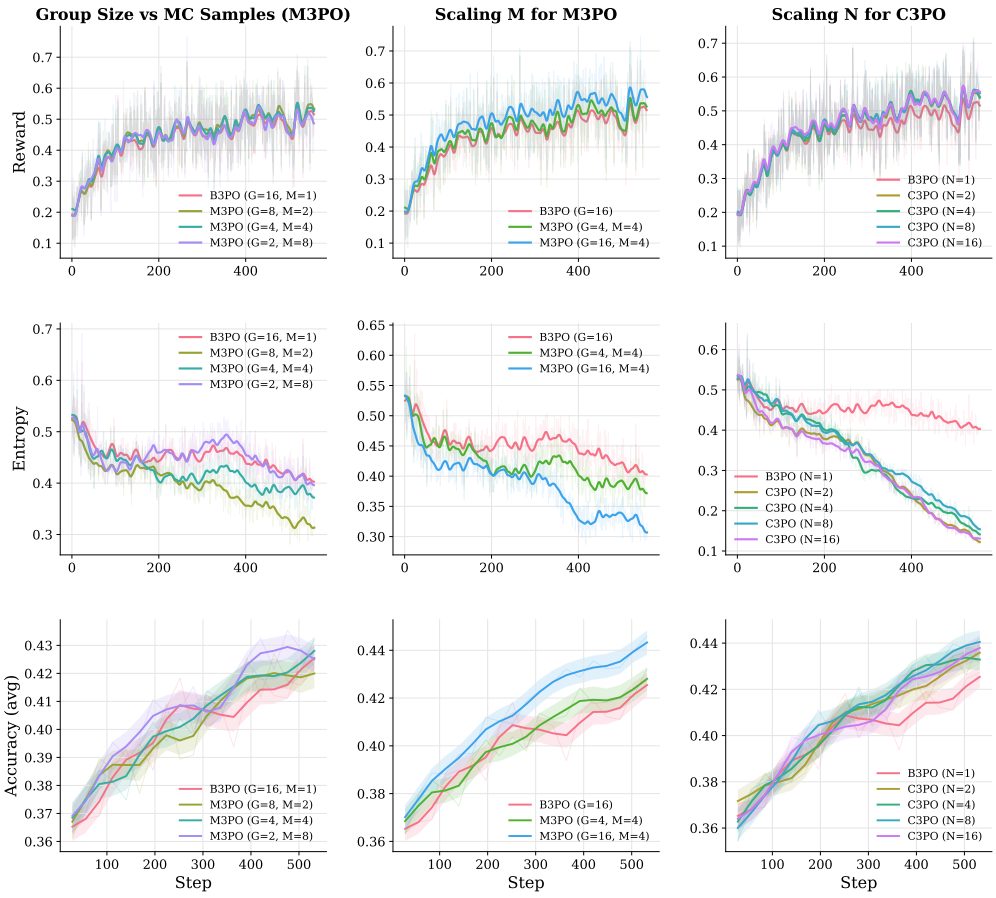


Figure 9: **Detailed results for ablating MC samples and chunk size.** Using $M > 1$ at a constant rollout budget lowers entropy, but average reward and downstream pass@1 are largely unchanged.

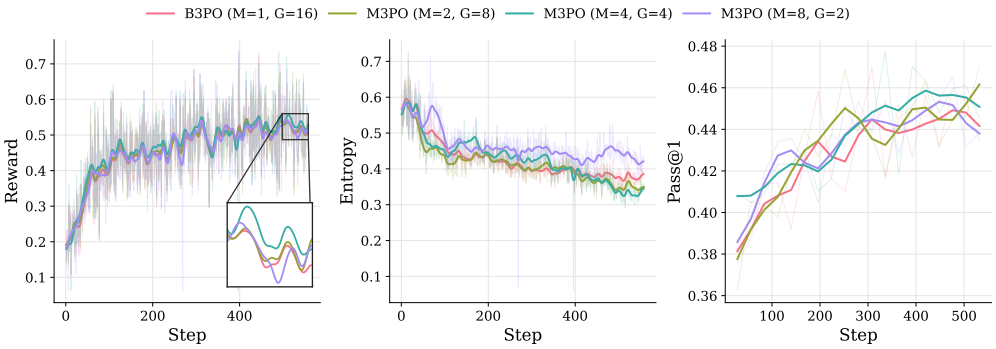


Figure 10: G–M **tradeoff on Qwen2.5-Math.** The variance reduction from larger M does not compensate for the loss in group diversity from smaller G, and vice versa.

rollout generation, producing very erratic reward and entropy curves at low λ. The default $\lambda = 10^9$ is too noisy for C3PO with Qwen2.5-Math, and 10^{10} works better, which we use in the main text.

D Ablating the number of Monte Carlo samples in M3PO

In Section 5.2, we analyzed the tradeoff between the group size G and the number of MC samples M for M3PO. Increasing M reduces gradient variance but proportionally increases compute, so we

shrink G to keep the rollout budget fixed. However, smaller G also reduces group diversity, which is crucial for grouped-advantage algorithms like GRPO [19, 22]. We presented the pass@1 scores in Figure 5(a-b); here we additionally plot the corresponding reward and entropy curves in Figure 9(a-b). Moving from $M = 1$ to $M = 2$ markedly lowers the entropy curve, but the reward curve plateaus only slightly higher and average downstream performance is unaffected. Scaling M further (at the cost of G) progressively closes the entropy gap. However, removing the equal-compute constraint substantially boosts M3PO, lifting reward and trading more entropy for downstream performance. This suggests that scaling M without shrinking the rollout group could be a viable strategy for improving performance, but it comes with a proportional computational cost.

We also study the G–M tradeoff for M3PO on Qwen2.5-Math. Following Section 5.2, we hold the total rollout count fixed, varying $M \in \{1, 2, 4, 8\}$ and reducing the group size proportionally. Figure 10 shows the same trend as on Olmo3: the gain from variance reduction is roughly cancelled by the loss in group diversity. As discussed in Section 5.2, realizing the variance-reduction benefit requires investing additional compute.

E Scaling the chunk size in C3PO

In this section, we examine the effect of varying the chunk size N in C3PO in more detail. Figure 9(c) shows the reward and entropy curves corresponding to the analysis in Section 5.3, and the trends mirror those of the pass@1 curves. While larger N yields marginally higher reward, even $N = 2$ captures almost all of the performance gains. The entropy trend matches Figure 6: $N > 1$ sharply increases group diversity, enabling the policy to keep discovering high-reward regions of the solution space and concentrate probability mass there, trading entropy for downstream performance [10, 71]. Together, these results suggest that setting N between 2 and 4 and relying on temperature sampling for the remaining rollouts is an effective strategy for improving downstream performance.

F Effects of a learned noise prior on other algorithms

In Section 5.4, we analyzed the impact of initializing the Hessian $\mathbf{h}$ with a learned prior obtained during the warm-start SFT phase for C3PO. We present detailed results for this comparison across the other 3PO methods in Figure 11. Intuitively, a learned prior should stabilize learning more than initializing from scratch and thereby improve performance. We do not, however, observe this behavior for any of our methods; for B3PO, the learned prior even slows convergence. We observed a similar pattern when using Qwen2.5-Math as well. A likely reason is that our SFT phase is relatively short, leaving the Hessian largely isotropic at the start of RL; learning the Hessian from pretraining itself is a promising direction for future work. In the meantime, the absence of a measurable advantage from the learned prior implies that 3PO can be applied to any off-the-shelf checkpoint without first running an IVON-based SFT phase to calibrate the noise distribution.

G Multi-seed robustness

Here, we validate the robustness of our methods to different seeds. However, because rerunning each method multiple times was not possible due to computational load we only compare GRPO and C3PO on Olmo3 across three seeds on mathematical reasoning tasks. C3PO improves over GRPO by an average of $+0.89$ points across the six benchmarks, positive on all three seeds, with a paired-t $p = 0.027$. The gains concentrate on harder benchmarks, with AIME'24 improving by $+1.8$ points ($p = 0.023$). These improvements therefore survive run-to-run variance, and combined with the code-generation results (Section 4.4) reinforce that our parameter-space exploration methods have the greatest benefits for harder tasks.

H Fixed isotropic noise vs. learned curvature

To isolate the contribution of the Hessian-scaled noise (Eq. 4), we run C3PO with isotropic noise of matched magnitude, replacing per-parameter variance with a single global scale. This variant attains an average accuracy of 42.12, below GRPO's 42.99 and well below full C3PO. As seen in Figure 12

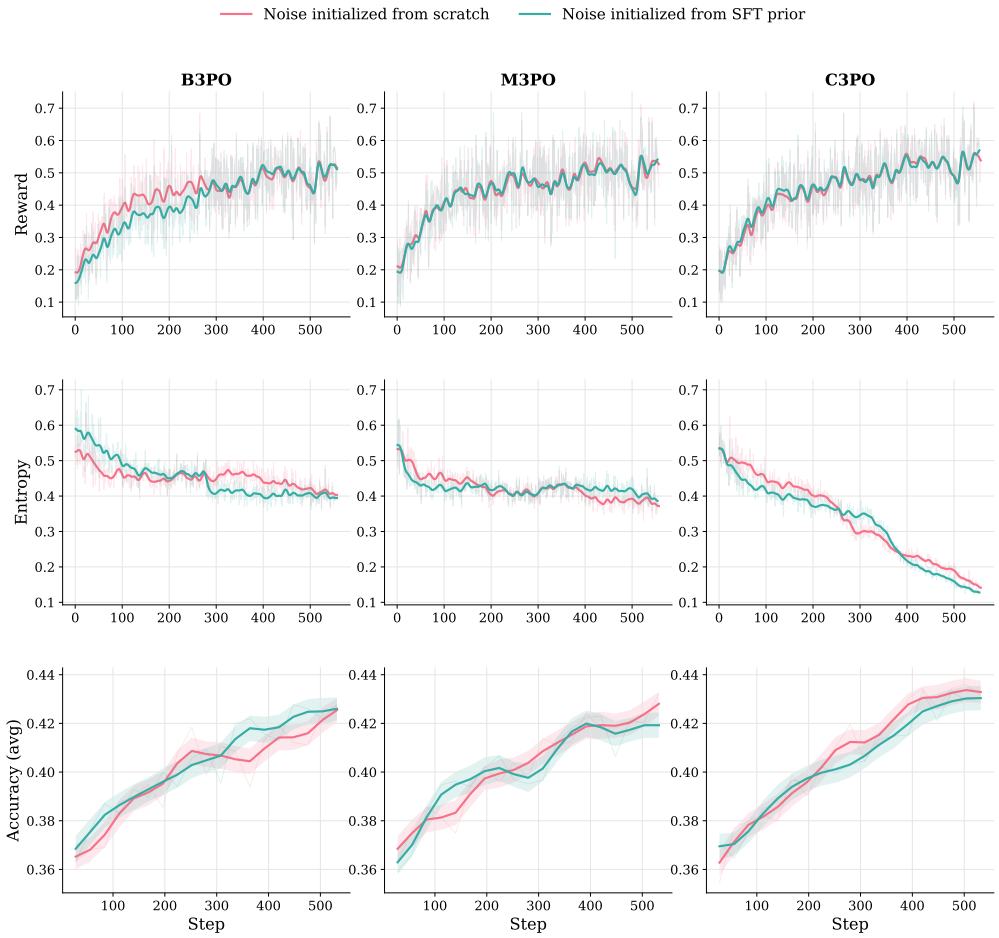


Figure 11: **Effects of a learned noise prior on other algorithms.** A learned noise prior gives no clear benefit for any 3PO variant and slows B3PO's convergence, possibly because the Hessian remains nearly isotropic after SFT.

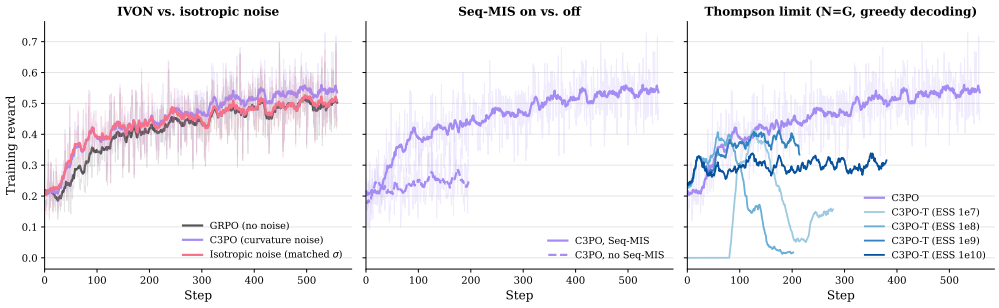


Figure 12: **(Left) IVON vs. fixed isotropic noise.** Isotropic noise climbs faster early but plateaus at a level similar to GRPO. **(Middle) Effect of Seq-MIS correction.** Training completely stalls without the correction due to training-inference mismatch. **(Right) Effect of Thompson sampling.** Sampling a fresh model for each rollout and acting greedily under it is unstable.

(Left), its reward does climb faster than GRPO early in training but converges to a similar plateau, whereas full C3PO plateaus above GRPO in late training.

I Ablating the Seq-MIS correction for C3PO

C3PO generates the rollouts of a group from N distinct perturbed models, which introduces a training – inference mismatch relative to the single policy assumed by the GRPO ratio in Eq. 1. To show that correcting this mismatch is necessary, we ran an early C3PO configuration that did not apply any such correction (Figure 12 (Middle)). This run's training reward curve stayed essentially flat throughout, in contrast to the steadily rising reward of our corrected runs. This mirrors the instability that the broader literature attributes to the training – inference mismatch in RLVR [70, 35], and motivates the correction used in all of our main results.

J A limiting case of C3PO

One limiting case of C3PO is to use $N = G$ with temperature $\tau = 0.0$, which would draw a fresh model for each rollout and decode greedily under it. However, we found this variant was very unstable (Figure 12 (Right)), with the reward curve oscillating wildly. Our most stable run at $\lambda = 10^{10}$ performed markedly worse than C3PO with temperature sampling. Thus, our methods in the main text perform exploration in both the action and parameter spaces: by first sampling a policy from the IVON posterior and further sampling actions under this policy.

K Algorithms

Require: Dataset $\mathcal{D}$, train policy π_θ, rollout policy π_μ
Constants: group size G, temperature τ, Monte-Carlo samples M

$(\mathbf{m},\, \boldsymbol{\sigma}) \leftarrow \mathrm{IVON}(\theta_0)$
for each batch $\mathcal{B} \sim \mathcal{D}$ **do**
  $\mathcal{L} \leftarrow 0$
  for $j = 1, \ldots, M$ **do** ▷ Optional. $M = 1$ for B3PO
    Sample $\hat{\theta}_j \leftarrow \mathbf{m} + \boldsymbol{\sigma} \odot \mathbf{z},\ \mathbf{z} \sim \mathcal{N}(0,\, I)$
    $\pi_\mu,\, \pi_\theta \leftarrow \hat{\theta}_j$ ▷ sync rollout and train policy
    for each $x \in \mathcal{B}$ **do**:
      sample $\{o_i\}_{i=1}^{G} \sim \pi_\mu(\cdot \mid x, \tau)$ ▷ generate G completions per prompt
    end for
    Compute rewards r_i and advantages $A_{i,t}$
    $\mathcal{L} \leftarrow \mathcal{L} + \mathcal{L}_{\mathrm{GRPO}}(\pi_\theta, A_{i,t})$ (cf. Eq. 1)
  end for
  $\mathbf{m}, \boldsymbol{\sigma} \leftarrow \mathrm{IVON.step}(\nabla_\theta \mathcal{L})$ ▷ Update posterior
end for

Algorithm 1: **Batched noising methods (B3PO and M3PO).** Weight perturbations $\hat{\theta}$ are sampled once per gradient step; all G rollouts for a batch of prompts are generated from the same model π_μ. Model weights are updated by accumulating losses over M Monte Carlo samples; $M = 1$ for B3PO.

Require: Dataset $\mathcal{D}$, train policy π_θ, rollout policy π_μ
Constants: group size G, temperature τ, chunk size N

$(\mathbf{m},\, \boldsymbol{\sigma}) \leftarrow \mathrm{IVON}(\theta_0)$
for each batch $\mathcal{B} \sim \mathcal{D}$ **do**
  $\mathcal{R} \leftarrow \emptyset$ ▷ initialize response buffer
  for $n = 1, \ldots, N$ **do**
    Sample $\hat{\theta}_n \leftarrow \mathbf{m} + \boldsymbol{\sigma} \odot \mathbf{z},\ \mathbf{z} \sim \mathcal{N}(0,\, I)$
    $\pi_\mu \leftarrow \hat{\theta}_n$ ▷ sync rollout policy
    for each $x \in \mathcal{B}$ **do**:
      Sample $\left\{o_i^{(n)}\right\}_{i=1}^{G/N} \sim \pi_\mu(\cdot \mid x, \tau)$;
      $\mathcal{R} \leftarrow \mathcal{R} \cup \left\{o_i^{(n)}\right\}$
    end for
  end for ▷ $\mathcal{R}$ accumulates G rollouts per prompt from N perturbations
  Compute rewards r_i and advantages $A_{i,t}$
  Sample $\hat{\theta}' \leftarrow \mathbf{m} + \boldsymbol{\sigma} \odot \mathbf{z}',\ \mathbf{z}' \sim \mathcal{N}(0,\, I);\ \pi_\theta \leftarrow \hat{\theta}'$ ▷ fresh draw for the update
  $\rho_i^{(n)} \leftarrow \pi_\theta\big(o_i^{(n)} \mid x\big) / \pi_{\hat{\theta}_n}\big(o_i^{(n)} \mid x\big)$; mask $o_i^{(n)}$ if $\rho_i^{(n)} \notin [0.5, 2]$ ▷ Seq-MIS
  $\mathcal{L} \leftarrow \mathcal{L}_{\mathrm{GRPO}}(\pi_\theta,\, A_{i,t})$ on unmasked rollouts, with $\pi_{\text{old}} = \pi_{\hat{\theta}_n}$ (cf. Eq. 1)
  $\mathbf{m}, \boldsymbol{\sigma} \leftarrow \mathrm{IVON.step}(\nabla_\theta \mathcal{L})$ ▷ Update posterior
end for

Algorithm 2: **Chunked noising approach (C3PO).** Rollouts are gathered in a buffer $\mathcal{R}$ across N independent weight draws $\hat{\theta}_n$, each generating G/N rollouts. The GRPO advantage is calculated on the accumulated buffer.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction clearly state the paper's contributions: identifying parameter-space exploration via variational learning as an additional axis of exploration in RLVR (Sections 2.1 and 2.2), proposing 3PO with three noising strategies in Section 3, and showing empirically that using multiple parameter samples per gradient step improves downstream performance on mathematical reasoning across two model families. These claims are substantiated by the main results in Section 4.2 and the analysis in Section 5.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 7 discusses several limitations of our work: the inability to scale to larger model families due to compute constraints (despite our hypothesis that bigger models would tolerate lower λ), the engineering bottleneck imposed by the single-model assumption in current RL frameworks and the higher wall-clock times of our methods due to inefficient implementations relative to AdamW.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper is primarily empirical and does not include formal theoretical results such as theorems, lemmas, or proofs.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: All hyperparameters for the warm-start SFT and RLVR phases are detailed in Appendix A, including optimizer choice, learning-rate schedule, λ, h_0, batch sizes, and PPO clip bounds. The training dataset (DapoMath-17k), evaluation benchmarks, base models, and inference settings (temperature, top-p, top-k, number of samples) are specified in Section 4.1 and Appendix A. Pseudocode for the proposed algorithms (B3PO/M3PO and C3PO) is provided in Appendix K.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We release our code publicly at `https://github.com/insait-institute/C3PO`. All training datasets (DapoMath-17k, Llama-Nemotron Post-Training Dataset) and base models (Olmo-3-1025-7B, Qwen2.5-Math-7B) are publicly available and appropriately cited. Hyperparameters and inference settings required to reproduce our results are provided in Appendix A, and pseudocode for the proposed algorithms is given in Appendix K.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 4.1 describes the experimental setup, including base models, datasets, training pipeline, and evaluation benchmarks. Appendix A provides complete optimizer hyperparameters for both the warm-start SFT and RLVR phases, along with inference settings used for evaluation. Algorithmic details for our proposed methods are given in Section 3 with pseudocode in Appendix K.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Across our figures and tables we report the mean Pass@1 with the standard error of the mean (SEM). Per-prompt Pass@1 is computed using the unbiased pass@k estimator of Chen et al. [7] with $K = 8$ rollouts (Appendix A); the SEM is computed across the prompts in each benchmark as $\sigma/\sqrt{N}$, where σ is the sample standard deviation of per-prompt Pass@1 and N is the number of prompts in the benchmark. We verify the robustness of our method to training noise by comparing C3PO and GRPO across three training seeds in Appendix G.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendix A specifies that all experiments are trained on an internal cluster of $8\times$ NVIDIA H200 GPUs, and the average amount of time taken per run. Table 3 additionally reports the wall-clock time per step for our proposed methods.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics. The research focuses on a training method for mathematical reasoning in LLMs and raises no ethical concerns specific to the Code of Ethics.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [N/A]

Justification: This work presents a foundational RLVR training method applied to mathematical reasoning. There is no direct path to specific negative societal impacts beyond those inherent to general improvements in LLM capability, and the paper does not introduce new data-collection procedures or deploy models in any public-facing system.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper releases a training method and accompanying mathematical reasoning models. Mathematical reasoning models trained on public benchmarks do not pose a high risk for misuse, and therefore no special safeguards beyond standard open-source licensing (CC-BY-4.0) are required.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All datasets (DapoMath-17k, Llama-Nemotron Post-Training Dataset), base models (Olmo-3-1025-7B, Qwen2.5-Math-7B), evaluation benchmarks (AIME 2024–26, AMC 2023, MATH-500, Minerva), and software libraries (`verl`, vLLM, IVON) are cited with their original papers. All assets used are publicly available and permissively licensed for research use.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: We release our code publicly at `https://github.com/insait-institute/C3PO`, sufficient to reproduce our experiments. The paper documents the training procedure (Sections 3 and 4.1), hyperparameters (Appendix A), pseudocode (Appendix K), and datasets used to produce the reported results.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.

- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects. All experiments are conducted on publicly available datasets and benchmarks.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or research with human subjects, so IRB approval is not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: We do not use LLMs as a non-standard component of our core methodology.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.