

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 VCRL: VARIANCE-BASED CURRICULUM REINFORCE- MENT LEARNING FOR LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Policy-based reinforcement learning currently plays an important role in improving LLMs on mathematical reasoning tasks. However, existing rollout-based reinforcement learning methods (GRPO, DAPO, GSPO, etc.) fail to explicitly consider LLMs' learning ability for samples of different difficulty levels, which is contrary to the human cognitive process of mathematical reasoning tasks from easy to difficult. Intuitively, we find that the variance of the rollout group's reward in Reinforcement Learning with Verifiable Rewards (RLVR) partly reflects the difficulty of the current sample for LLMs. Samples that are too easy or too difficult have a lower variance, while samples with moderate difficulty have a higher variance. Based on this, we propose VCRL, a curriculum reinforcement learning framework that dynamically controls the difficulty of training samples based on the variance of group rewards. Experiments on five mathematical benchmarks and two models reveal the advantages of VCRL over the current LLM RL baselines. Code is available at <https://anonymous.4open.science/r/VCRL-BD7E>.

1 INTRODUCTION

The new generation of large language models (LLMs) that use long Chain-of-Thoughts (CoTs) for reasoning (Xu et al., 2025a) have achieved remarkable results in information extraction (Zhang et al., 2025d; Jiang et al., 2024; 2025a), mathematics (Wang et al., 2025a), code (Yang et al., 2025b), and agent (Gao et al., 2025; Zhang et al., 2025b) fields, including GPT-5¹, GPT-OSS (Agarwal et al., 2025), DeepSeek-R1 (Guo et al., 2025), and Kimi k1.5 (Team et al., 2025). A notable feature of this type of LLMs is the phenomenon called Test-Time Scaling (TTS) (Zhang et al., 2025c), which generates long CoTs to scale performance. Reinforcement Learning with Verifiable Rewards (RLVR) (Mroueh, 2025) has been proven to be an effective technique for achieving TTS in the post-training process.

Recently, Reinforcement Learning (RL) methods have shown significantly better generalization performance in improving LLM reasoning capabilities compared to traditional Supervised Fine-Tuning (SFT) (Chu et al., 2025). SFT relies on high-quality, labeled data from human annotations or stronger model distillation, while RL relies primarily on the model's own exploration. Rollout-based reinforcement learning methods represented by Group Relative Policy Optimization (GRPO) (Shao et al., 2024) require the model to generate multiple trajectories for each training sample and learn based on the rewards of the generated trajectories, which can continuously expand the boundaries of LLMs' capabilities through the continuous RL process with diverse training samples.

However, existing rollout-based RL methods do not consider how well the model's current abilities match the difficulty of training samples. In human learning, people usually start with easy tasks and move to harder ones, an approach called Curriculum Learning (CL) (Wang et al., 2022; Soviany et al., 2022). Rollout-based RL methods have the model explore rollouts generated by the training samples, without considering if those samples are easy or hard. This does not help LLMs learn efficiently from samples with different levels of difficulty. Also, the model's skills change during RL training, so the difficulty of training samples can vary for the model at different stages. Because of this, pre-sorting training samples by fixed difficulty is not effective.

¹<https://openai.com/index/gpt-5-system-card>

To address these limitations, we introduce a curriculum reinforcement learning framework called VCRL. It dynamically adjusts the difficulty of training samples based on the variance of group rewards. We find that the variance in rollout group rewards in RLVR partly reflects how hard a sample is for LLMs. With RLVR’s current sparse reward system, samples that are too hard often get only 0 rewards, leading to low variance; this also happens with samples that are too easy. When samples are more uncertain, such as when half of the rollouts receive a reward of 1 and the other half receive 0, the model is at a key learning point for that sample. VCRL uses **Variance-based Dynamic Sampling** to select these samples for training, helping control the quality of the training batch. Group variance also gives a way to measure sample difficulty for the current state of the model. Therefore, VCRL uses **Replay Learning** with a memory bank to further boost training efficiency.

Our contributions are as follows:

- We introduce VCRL, a curriculum reinforcement learning framework that adjusts the difficulty of training samples based on the variance of group rewards. By focusing on samples with high reward variance, VCRL selects those most valuable for current model training.
- Building on group variance, we further introduce Replay Learning with a memory bank to control training stability and improve training efficiency. By updating and utilizing the memory bank, VCRL ensures high variance of samples in the training batch, thus achieving higher training value.
- We conduct extensive experiments on five benchmark datasets to justify VCRL’s advantage on LLM’s efficient Test-Time Scaling over some SOTA RL methods. Our results show consistent performance gains across different models, validating the effectiveness and robustness of our VCRL.

2 PRELIMINARIES

In this section, we review the current policy-based reinforcement learning methods in LLM, especially the rollout-based like GRPO and some variants.

2.1 PROXIMAL POLICY OPTIMIZATION (PPO)

PPO (Schulman et al., 2017) limits the update of the current policy to the proximal region of the old policy through the clipping mechanism. Specifically, given a dataset \mathcal{D} , x is the query and y is the response. For the policy model π_θ parameterized by θ , the likelihood by the policy π_θ is given by $\pi_\theta(y|x) = \prod_{t=1}^{|y|} \pi_\theta(y_t|x, y_{<t})$, where $|y|$ is the number of tokens in y . In RLVR, there is a verifier r that can score a given query-response pair (x, y) and obtain a reward $r(x, y) \in [0, 1]$. PPO optimizes the following objective for policy optimization to update the actor in the proximal region of the old policy $\pi_{\theta_{\text{old}}}$:

$$\mathcal{J}_{\text{PPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi_{\theta_{\text{old}}}(\cdot|x)} \left[\frac{1}{|y|} \sum_{t=1}^{|y|} \min \left(r_t(\theta) \hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right], \quad (1)$$

where the importance ratio of the token y_t is given by $r_t(\theta) = \frac{\pi_\theta(y_t|x, y_{<t})}{\pi_{\theta_{\text{old}}}(y_t|x, y_{<t})}$, ϵ is the clipping range of the importance ratio, and the advantage \hat{A}_t of y_t is estimated using a value model estimated by Generalized Advantage Estimator (GAE) (Schulman et al., 2016).

PPO relies on the value model to evaluate the current state. Typically, the value model and the trained model have similar structures and parameters, resulting in significant computational and memory costs. Furthermore, the accuracy of the value model itself limits the effectiveness of the PPO algorithm, especially for long response and sparse reward in complex tasks for LLM.

2.2 GROUP RELATIVE POLICY OPTIMIZATION (GRPO) AND VARIANTS

GRPO (Shao et al., 2024) calculates the relative advantages of each response within a group of responses generated by LLM to the same query, eliminating the need to the value model. Specifically,

108 GRPO optimizes the following objective for policy optimization to update the actor within the group
 109 of responses (we omit the KL regularization term for brevity):
 110

$$111 \quad \mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|x)} \\ 112 \quad \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t} \right) \right], \quad (2)$$

116 where G is the number of generated responses to the same query x , the importance ratio $r_{i,t}(\theta)$ and
 117 advantage $\hat{A}_{i,t}$ of token $y_{i,t}$ are given by
 118

$$119 \quad r_{i,t}(\theta) = \frac{\pi_{\theta}(y_{i,t}|x, y_{i,<t})}{\pi_{\theta_{\text{old}}}(y_{i,t}|x, y_{i,<t})}, \quad \hat{A}_{i,t} = \frac{r(x, y_i) - \text{mean}(\{r(x, y_i)\}_{i=1}^G)}{\text{std}(\{r(x, y_i)\}_{i=1}^G)}. \quad (3)$$

122 Based on GRPO, Decoupled Clip and Dynamic sampling Policy Optimization (DAPO) (Yu et al.,
 123 2025) removes the KL divergence regularization and introduces the clip-higher and dynamic sam-
 124 pling with token-level loss, further improving the training stability and performance for LLMs.
 125 Specifically, DAPO optimizes the following objective for policy optimization:

$$126 \quad \mathcal{J}_{\text{DAPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|x)} \\ 127 \quad \left[\frac{1}{\sum_{i=1}^G |y_i|} \sum_{i=1}^G \sum_{t=1}^{|y_i|} \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}}) \hat{A}_{i,t} \right) \right], \quad (4)$$

$$131 \quad \text{s.t. } 0 < |\{y_i \mid \text{is_equivalent}(y^*, y_i)\}| < G,$$

132 where ϵ_{low} and ϵ_{high} are the low and high clipping bound for the importance ratio respectively, and
 133 y^* is the correct answer.
 134

135 Based GRPO, Group Sequence Policy Optimization (GSPO) (Zheng et al., 2025) uses sequence-
 136 level importance ratio to replace the original token-level importance ratio to match the sentence-level
 137 reward in the generation task and optimization objective, thus achieving remarkable improvements.
 138 Specifically, GSPO optimizes the following objective for policy optimization:

$$139 \quad \mathcal{J}_{\text{GSPO}}(\theta) = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|x)} \left[\frac{1}{G} \sum_{i=1}^G \min \left(s_i(\theta) \hat{A}_i, \text{clip}(s_i(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) \right], \quad (5)$$

142 where the group-based advantage estimation and importance ratio are given by
 143

$$144 \quad \hat{A}_i = \frac{r(x, y_i) - \text{mean}(\{r(x, y_i)\}_{i=1}^G)}{\text{std}(\{r(x, y_i)\}_{i=1}^G)}, \quad s_i(\theta) = \left(\frac{\pi_{\theta}(y_i|x)}{\pi_{\theta_{\text{old}}}(y_i|x)} \right)^{\frac{1}{|y_i|}}. \quad (6)$$

147 3 VARIANCE-BASED CURRICULUM REINFORCEMENT LEARNING

149 In this section, we introduce Variance-based Curriculum Reinforcement Learning (VCRL), shown
 150 in Figure 1. First, we explain Variance-based Dynamic Sampling and how it helps identify the
 151 difficulty and value of training samples. Next, we combine Replay Learning with a memory bank to
 152 focus training on high-value samples, which improves RL training efficiency and stability.
 153

154 3.1 VARIANCE-BASED DYNAMIC SAMPLING

156 As discussed above, existing rollout-based RL methods do not properly match model capabilities
 157 with sample difficulty during training. This problem mainly shows up in two ways:
 158

- 159 **1. Dynamic Model Parameters:** During training, gradient backpropagation is performed us-
 160 ing the objective function calculated from the training samples. This updates the model pa-
 161 rameters to improve its performance on current samples. Model parameters keep changing,
 162 so the model may perform differently on the same samples at different stages of training.

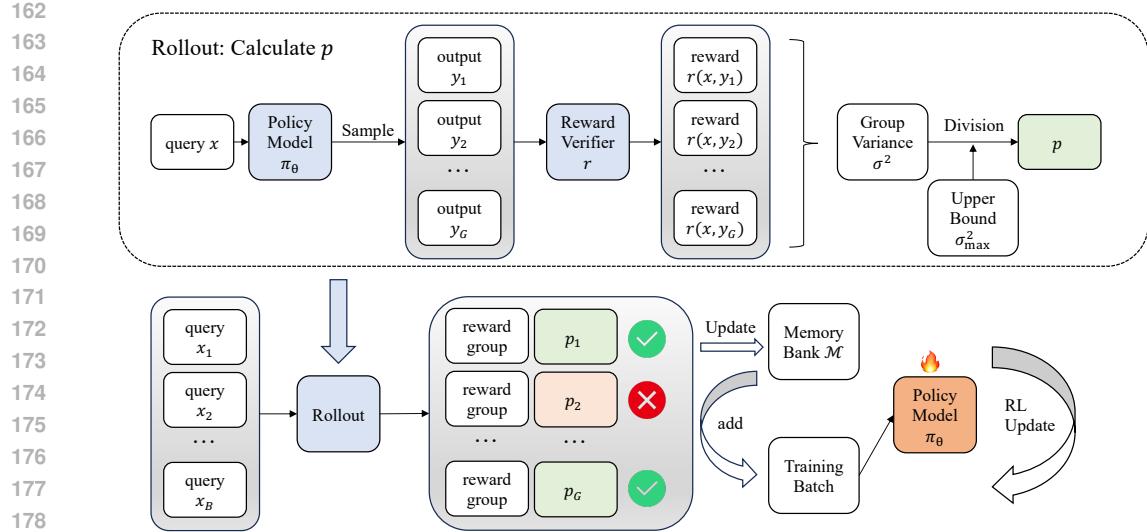


Figure 1: An illustration of the proposed VCRL method. For rollout-based RL training, VCRL first calculates our proposed p for each query’s rollout results and filters queries based on their p . VCRL then uses the existing memory bank \mathcal{M} to update and add training samples. Finally, VCRL performs the standard RL update for this training batch.

2. **Unordered Sample Difficulty:** For most training datasets and algorithms, the difficulty of training samples is not considered. Some tasks, like search (Feng et al., 2025b; Hao et al., 2025) and tool use (Shen, 2024; Lu et al., 2025), are hard to define by difficulty. Also, sorting samples by difficulty requires a lot of data preprocessing. As a result, most datasets include samples that are not ordered by difficulty.

Dynamic model parameters and unordered sample difficulty make it too expensive and hard to use ordered training samples based on predefined difficulty. Samples that are hard for the model early in training often become easier later. So, the indicator of training sample difficulty must be adjusted dynamically as the model changes.

Multiple rollouts for the same query can help measure how hard a training sample is for the current model. Formally, for a query x , if it is too easy for model π_θ , then $\mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [r(x, y)] \approx 1$. If x is too hard, then $\mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [r(x, y)] \approx 0$. Both easy and hard samples have low group reward variance. So, we can use the variance of group rewards to pick samples that are better suited for the current model. Samples with higher variance are neither too easy nor too hard, meaning the difference between the probabilities of positive and negative outcomes is small.

In RLVR, for the binary reward distribution, the group variance for the query x is given by

$$\text{Var}_{y \sim \pi_\theta(\cdot|x)} (r(x, y)) = \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [(r(x, y) - \mathbb{E}_{y \sim \pi_\theta(\cdot|x)} [r(x, y)])^2]. \quad (7)$$

If there are k rollouts with a reward of 1, the unbiased estimator of the group variance can be written as

$$\begin{aligned} \sigma^2 &= \frac{1}{G-1} \sum_{i=1}^G \left[r(x, y_i) - \frac{1}{G} \sum_{i=1}^G r(x, y_i) \right]^2 \\ &= \frac{1}{G-1} \sum_{i=1}^G \left[r(x, y_i) - \frac{k}{G} \right]^2 \\ &= \frac{k(G-k)}{G(G-1)}. \end{aligned} \quad (8)$$

216 When $k = \lfloor \frac{G}{2} \rfloor$, the estimator achieves the maximum value:
 217

$$218 \quad \sigma_{\max}^2 = \begin{cases} \frac{G}{4(G-1)}, & G \text{ is even,} \\ 219 \quad \frac{G+1}{4G}, & G \text{ is odd.} \end{cases} \quad (9)$$

221 Obviously, the group variance cannot exceed σ_{\max}^2 in any case, so we can use the normalized group
 222 variance $p = \frac{\sigma^2}{\sigma_{\max}^2}$ to measure the value of the current query x for the model π_θ . Training with
 223 samples that have high p helps the model learn areas where it is less skilled, which improves the
 224 model more effectively than using unordered samples. See Appendix Section B for more discussion.
 225

226 3.2 REPLAY LEARNING

228 Based on the normalization value p discussed above, we can dynamically sample queries during
 229 training using threshold rules. This helps ensure that each training sample has high value for the
 230 model. For unordered training datasets, each sampled query can only obtain its p value after a long
 231 rollout, so we use variance-based dynamic sampling. Calculating p for each training sample requires
 232 significant computational resources and time, which can be expensive if used only for sampling.

233 To address this, we propose building a high-value memory bank using p and maintaining it with a
 234 momentum update method. This lets us apply curriculum learning with data replay based on group
 235 variance, as shown in Algorithm 1. Specifically, each time we sample from the training set \mathcal{D} , we
 236 get a query batch $\{x_j\}_{j=1}^B$, where B is the batch size. First, we get the corresponding response set
 237 $\{y_{j,i}\}_{i=1}^G$ and reward set $\{r(x_j, y_{j,i})\}_{i=1}^G$, then calculate p_j for each query x_j . If $p_j \geq \kappa$, where
 238 $\kappa \in [0, 1]$ is a predefined threshold, we keep the query x_j . Otherwise, we remove it from the batch
 239 and perform variance-based dynamic sampling.

240 Suppose M queries are removed from a batch of B . To keep the batch size unchanged, we replace
 241 the missing M queries by sampling queries from the memory bank \mathcal{M} . The memory bank \mathcal{M} is
 242 implemented as a priority queue, where each entry is a query x_j , and the priority $P(x_j)$ is updated
 243 based on momentum and the number of steps since it was last accessed, $\beta(x_j)$:

$$244 \quad P(x_j) \leftarrow \alpha P(x_j) + (1 - \alpha) \beta(x_j), \quad (10)$$

245 where α is the momentum constant and the $P(x_j)$ is initialized using p_j .

247 The proposed VCRL based on GRPO optimizes the following objective for policy optimization:

$$249 \quad \mathcal{J}_{\text{VCRL}}(\theta) = \mathbb{E}_{x \sim \mathcal{D} \cup \mathcal{M}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}(\cdot|x)}} \\ 250 \quad \left[\frac{1}{G} \sum_{i=1}^G \frac{\mathbb{I}\left(p_i = \frac{\sigma_i^2}{\sigma_{\max}^2} \geq \kappa\right)}{|y_i|} \sum_{t=1}^{|y_i|} \min\left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t}\right) \right], \\ 253 \quad (11)$$

254 where the calculation of p_i and the memory bank \mathcal{M} mechanism are as described above, and $\mathbb{I}(\cdot)$
 255 is the indicator function. See Appendix Section C for a comparison of theoretical perspectives on
 256 GRPO and VCRL.

258 4 EXPERIMENTS

261 4.1 EXPERIMENTAL SETUP

263 **Benchmarks.** In this work, we focus specifically on mathematical reasoning tasks to evaluate our
 264 VCRL algorithm. For mathematical reasoning tasks, we use AIME-2024², AIME-2025³, MATH500
 265 (Lightman et al., 2024), OlympiadBench (He et al., 2024), and AMC23⁴. Among them, AIME-2024
 266 and AIME-2025 are used as high-difficulty benchmarks to effectively evaluate the performance of
 267 VCRL and other baseline RL methods in multiple difficulty levels.

268 ²https://huggingface.co/datasets/Maxwell-Jia/AIME_2024

269 ³https://huggingface.co/datasets/yentinglin/aime_2025

⁴<https://huggingface.co/datasets/AI-MO/aimo-validation-amc>

270

Algorithm 1: VCRL: Variance-based Curriculum Reinforcement Learning

271

Require: Training Set \mathcal{D} , Reward Verifier r , p -threshold κ , Policy Model π_θ , Momentum Constant α , Training Batch Size B , Rollout Group Size G

- 1: Initialize $\mathcal{M} \leftarrow \text{PriorityQueue}()$
- 2: **while** Training **do**
- 3: Sample $\{x_j\}_{j=1}^B \sim \mathcal{D}$, $M \leftarrow 0$
- 4: **for** $j = 1$ to B **do**
- 5: Sample $\{y_{j,i}\}_{i=1}^G \sim \pi_\theta(\cdot|x_j)$
- 6: Calculate Reward $\{r(x_j, y_{i,j})\}_{i=1}^G$
- 7: Calculate p_j for x_j
- 8: **if** $p_j < \kappa$ **then**
- 9: Remove x_j from Training Batch
- 10: $M \leftarrow M + 1$
- 11: **end if**
- 12: **end for**
- 13: Pop M queries from \mathcal{M} and add them to the Training Batch
- 14: **for** $x \in \mathcal{M}$ **do**
- 15: $\beta(x) \leftarrow \beta(x) + 1$
- 16: $P(x) \leftarrow \alpha P(x) + (1 - \alpha)\beta(x)$
- 17: **end for**
- 18: Apply RL update using the Augmented Training Batch \mathcal{B}
- 19: **for** $x \in \mathcal{B}$ **do**
- 20: Calculate p for x
- 21: **if** $p \geq \kappa$ **then**
- 22: Push x into \mathcal{M} with priority $P(x) = p$ and $\beta(x) = 0$
- 23: **end if**
- 24: **end for**
- 25: **end while**

297

298

299

Implementation Details. For training dataset, we use DAPO-Math-17K⁵ to improve training stability, which consists of 17K prompts, each paired with an integer as the answer. We implement VCRL and conduct all experiments based on the verl (Sheng et al., 2025) framework. For hyperparameters, we utilize the AdamW (Loshchilov & Hutter, 2019) optimizer with a constant learning rate of 1×10^{-6} . For rollout, the prompt batch size is $B = 128$ and we sample $G = 16$ responses for each prompt. For training, we train 500 steps to ensure convergence. The maximum number of tokens for generation is set to 4,096 tokens. For evaluation on benchmarks, we repeat the evaluation set for 16 times and report avg@16 for the stability of the results. The inference hyperparameters of evaluation are set to temperature 0.6 and top-p 0.95. For VCRL, we set the variance threshold κ to 0.3 in first 20 steps and 0.8 in remaining steps, and the momentum constant α is set to 0.9. We implement VCRL based on GRPO’s RL update. For memory bank, we allow up to 2 replays for the same sample to ensure the diversity of training sample. We conduct all experiments on a server with 8×NVIDIA H20-3e GPUs and an Intel® Xeon® Platinum 8575C CPU.

311

312

313

314

315

316

Baselines and Models. We mainly use GRPO (Shao et al., 2024), DAPO (Yu et al., 2025) and GSPO (Zheng et al., 2025) as the baselines for our VCRL comparison. For Clip-Higher mechanism in DAPO, we set the clipping parameter ϵ_{low} to 0.2 and ϵ_{high} to 0.28, which is aligned with the DAPO setting in the original paper. For GSPO, we set the clipping parameter ϵ to 0.0003. For models, we use the Qwen3 (Yang et al., 2025a) series models for training, including *Qwen3-4B-Base* and *Qwen3-8B-Base*.

317

318

4.2 MAIN RESULTS

319

320

321

322

323

We conduct a comprehensive evaluation of our proposed method, VCRL, against several strong LLM RL baselines on a diverse suite of mathematical reasoning benchmarks. As detailed in Table 1, the experiments are performed on two models, *Qwen3-4B-Base* and *Qwen3-8B-Base*, to assess the

⁵<https://huggingface.co/datasets/BytedTsinghua-SIA/DAPO-Math-17k>

324 Table 1: Main performance comparison of VCRL against other RL baselines on Qwen3 models.
325

326 Method	327 AIME-2024	327 AIME-2025	327 MATH500	327 OlympiadBench	327 AMC23	327 Avg.
<i>Starting from Qwen3-4B-Base</i>						
328 Base Model	329 9.58	329 4.79	329 56.69	329 27.27	329 35.09	329 26.68
330 + GRPO	330 15.63	330 12.92	330 80.78	330 45.39	330 54.07	330 41.76
331 + DAPO	331 14.79	331 12.29	331 79.86	331 44.23	331 51.81	331 40.60
332 + GSPO	332 14.58	332 10.42	332 79.90	332 44.38	332 51.13	332 40.08
333 + VCRL	333 23.96	333 22.71	333 86.48	333 53.24	333 60.77	333 49.43
<i>Starting from Qwen3-8B-Base</i>						
334 Base Model	335 10.83	335 10.00	335 68.75	335 34.10	335 41.11	335 32.96
336 + GRPO	336 23.13	336 21.88	336 86.94	336 54.02	336 65.29	336 50.25
337 + DAPO	337 22.08	337 20.42	337 87.14	337 53.52	337 64.01	337 49.43
338 + GSPO	338 27.29	338 22.92	338 89.23	338 56.75	338 69.28	338 53.09
339 + VCRL	339 34.38	339 27.08	339 91.99	339 60.21	339 75.15	339 57.76

340 scalability and generalizability of our method. The results unequivocally demonstrate the superiority
341 of VCRL. Across all five benchmarks and on both model sizes, VCRL consistently achieves state-of-
342 the-art performance, outperforming all baseline methods, including GRPO, DAPO, and GSPO. This
343 consistent dominance, indicated by the bolded scores, highlights the robustness and effectiveness of
344 our proposed methodology.

345 A deeper analysis reveals the substantial performance gains enabled by VCRL. For instance, on the
346 *Qwen3-8B-Base* model, VCRL achieves an average score of 57.76, a significant margin of over 4.67
347 points above the strongest baseline, GSPO (53.09), and a remarkable 24.8 points improvement over
348 the base model. This trend holds for the *Qwen3-4B-Base* model, where VCRL elevates the average
349 performance from 26.68 (Base Model) to 49.43, far surpassing the gains from other RL techniques.
350 Notably, the performance leap is particularly pronounced on highly challenging, competition-level
351 datasets such as AIME-2024 and AIME-2025, suggesting that VCRL is exceptionally proficient
352 at unlocking the complex, multi-step reasoning capabilities essential for advanced mathematical
353 problem-solving. These empirical findings strongly validate VCRL as a superior alignment strategy
354 for enhancing the mathematical reasoning prowess of LLMs.

355 4.3 PERFORMANCE TREND

356 During RL training, the LLM starts with low ability and steadily improves, showing an upward
357 trend on benchmark tests. To illustrate how VCRL compares to baseline methods during training,
358 we show how model performance changes with training steps on each benchmark, as seen in Figure
359 2 for *Qwen3-4B-Base* and Figure 3 for *Qwen3-8B-Base*.

360 For performance trend, the results clearly demonstrate that VCRL consistently and significantly
361 outperforms all other baseline methods across all benchmarks. In terms of the speed of performance
362 improvement, VCRL also has a considerable advantage. In the first 100 training steps, VCRL's
363 performance increases quickly, with its curve staying above the other methods. This is likely due
364 to VCRL's control of high- p training samples in the early stages of model training, which improves
365 training efficiency. Later in training, the performance of all methods generally converges, but VCRL
366 still achieves significantly better final results than the RL baselines. This demonstrates VCRL's
367 strong competitiveness. More training dynamics are in Appendix Section A.

370 4.4 ABLATION STUDY

371 To verify the effectiveness of the two core components of our proposed VCRL, we conduct the
372 ablation study, as shown in Table 2. Starting from the *Qwen3-4B-Base*, our Naive GRPO base-
373 line improves the average score from 26.68 to 41.76. The integration of Variance-based Dynamic
374 Sampling further pushes this score to 44.73. Finally, the inclusion of Replay Learning achieves the
375 best performance of 49.43, showing the largest marginal gain. This consistent trend on the larger
376 *Qwen3-8B-Base* model robustly validates the positive impact of each component within our VCRL
377 framework.

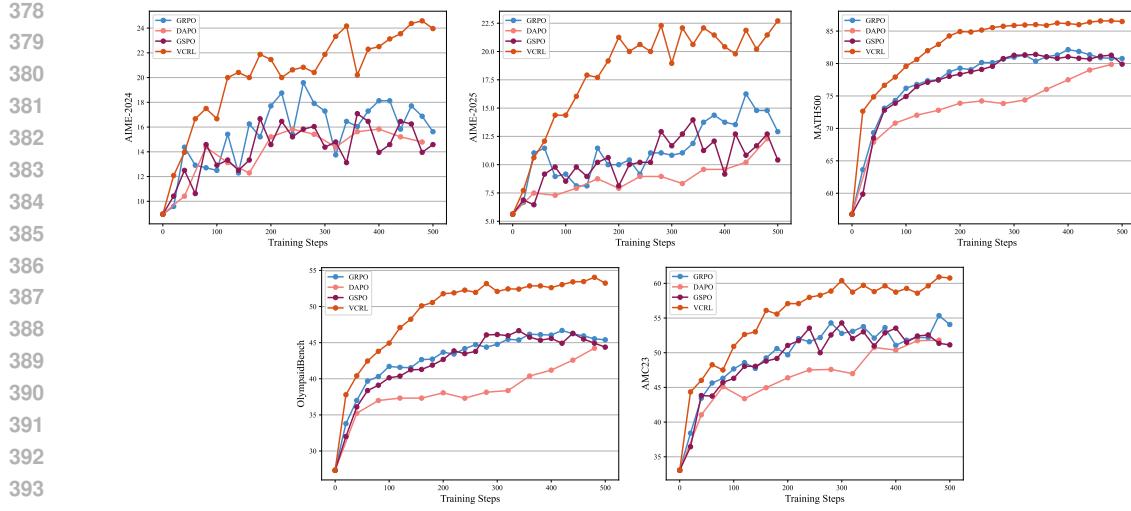


Figure 2: The performance curve of *Qwen3-4B-Base* on the five benchmarks using various RL methods over training steps.

5 RELATED WORK

Recent work on using RL methods with LLMs has greatly improved their ability to handle complex tasks. DeepSeek-R1 (Guo et al., 2025) introduces a zero RL training framework, which directly trains the base LLM using a simple rule-based reward model. Many RL methods have built on this idea to further boost LLM performance.

Some approaches use novel RL mechanisms to make training more efficient and stable. DAPO (Yu et al., 2025) analyzes GRPO’s training and applies four main techniques to improve RL efficiency. Dr. GRPO (Liu et al., 2025b) removes the output length and standard deviation terms from GRPO’s relative advantage, which increases token efficiency without hurting reasoning performance. SimpleRL-Zoo (Zeng et al., 2025) runs experiments on different base models and sizes to map out behavioral patterns and suggest future improvements. LUFFY (Yan et al., 2025) enhances RLVR with off-policy reasoning traces, helping to balance imitation and exploration by combining off-policy demonstrations with on-policy rollouts. VAPO (Yue et al., 2025) introduces the first value-model-based RL training framework built on PPO, with seven new techniques to improve

432 Table 2: Ablation study of the key components of our proposed method VCRL. Starting from a Naive
 433 GRPO baseline, we incrementally add **Variance-based Dynamic Sampling** and **Replay Learning**.
 434 The results on both *Qwen3-4B-Base* and *Qwen3-8B-Base* models show that each component con-
 435 tributes positively to the final performance, validating their effectiveness.

437 Model	438 Avg.
439 <i>Qwen3-4B-Base</i>	438 26.68
440 w/ Naive GRPO	41.76
441 w/ Variance-based Dynamic Sampling	44.73
442 w/ Replay Learning	49.43
443 <i>Qwen3-8B-Base</i>	32.96
444 w/ Naive GRPO	50.25
445 w/ Variance-based Dynamic Sampling	52.67
446 w/ Replay Learning	57.76

447 training stability and performance. Yeo et al. (2025) investigates how RL helps models create longer
 448 reasoning chains, showing which factors matter most for extended CoT reasoning. PVPO (Feng
 449 et al., 2025c) presents an efficient reinforcement learning method enhanced by an advantage refer-
 450 ence anchor and data pre-sampling.

451 Other work explores curriculum learning in LLM training for better results. Hammoud et al. (2025)
 452 improve GRPO with a reward function that balances task correctness (via verifier feedback), length
 453 efficiency, and formatting (using structural tags), leading to higher accuracy and better token effi-
 454 ciency. Feng et al. (2025a) propose a self-adaptive curriculum that picks fine-tuning examples based
 455 on difficulty scores predicted by pre-trained models. Shen et al. (2025) introduce TTI (Test-Time
 456 Interaction), an online RL method that adapts rollout lengths using a curriculum approach. Parashar
 457 et al. (2025) provide convergence guarantees for easy-to-hard training within an approximate pol-
 458 icy iteration framework. RAGEN (Wang et al., 2025c) introduces uncertainty-based filtering to
 459 maintain high training efficiency based on active learning (Settles, 2009). PODS (Xu et al., 2025b)
 460 generates numerous rollouts in parallel but updating only on informative subset. Curr-ReFT (Deng
 461 et al., 2025) explores the Out-of-Distribution generalization on small-scale Vision Language Mod-
 462 els based on the curriculum learning framework. Xi et al. (2024) introduce a novel method that
 463 employs only outcome supervision to achieve the benefits of process supervision for large language
 464 models with a step-wise curriculum. More recently, several studies have focused on difficulty-based
 465 filtering and distribution priors. ADCL (Zhang et al., 2025a) employs adaptive difficulty curriculum
 466 via periodic data reordering, while other approaches (Chen et al., 2025; Wang et al., 2025b) rely on
 467 self-evolving schedules or distribution-level priors which often require preset difficulty hierarchies.
 468 Regarding data efficiency, Bae et al. (2025) propose ODF with a pass-based selection strategy simi-
 469 lar to DAPO, and Tzannetos et al. (2023) explore the theoretical foundations of proximal curriculum.
 470 Distinct from these methods, our VCRL operates dynamically by filtering based on group variance
 471 within the current batch and utilizes replay learning to enhance efficiency, without assuming future
 472 data streams or requiring strict difficulty distribution priors.

473 RLVR (Mroueh, 2025) is a promising method for boosting reasoning in LLMs, especially in areas
 474 like math and programming (Jiang et al., 2025b). Gandhi et al. (2025) show that reasoning behav-
 475 iors—not just correct answers—drive RL performance gains. Li et al. (2025) find that the structure
 476 of long chains of thought is key for learning, while the details of each reasoning step matter less.
 477 Vassoyan et al. (2025) identify critical tokens in CoTs, which are decision points where models often
 478 make mistakes, and suggest increasing exploration around these tokens by changing the KL penalty.
 479 Lin et al. (2024) also find tokens that lead to errors and show that changing them can shift model
 behavior.

480 6 CONCLUSION

481 In this paper, we propose VCRL, a curriculum reinforcement learning framework that dynamically
 482 controls the difficulty of training samples based on the variance of group rewards. By introducing
 483 Dynamic Variance Sampling, VCRL can filter out samples in the training batch that are moderately
 484 difficult for the current training model and remove samples that are too difficult or too easy, thereby

486 improving training efficiency. By introducing Replay Learning, VCRL uses a memory bank to
 487 maintain the high- p samples in the training batch, further improving training stability. By carefully
 488 controlling the difficulty of training samples, VCRL achieves state-of-the-art results on five math
 489 benchmarks compared to LLM RL baselines. Further analysis of training dynamics and ablation
 490 study also confirm VCRL’s effectiveness.

491

492 REFERENCES

493

494 Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K
 495 Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv*
 496 *preprint arXiv:2508.10925*, 2025.

497 Sanghwan Bae, Jiwoo Hong, Min Young Lee, Hanbyul Kim, JeongYeon Nam, and Donghyun
 498 Kwak. Online difficulty filtering for reasoning oriented reinforcement learning. *arXiv preprint*
 499 *arXiv:2504.03380*, 2025.

500 Xiaoyin Chen, Jiarui Lu, Minsu Kim, Dinghuai Zhang, Jian Tang, Alexandre Piché, Nicolas Gontier,
 501 Yoshua Bengio, and Ehsan Kamalloo. Self-evolving curriculum for llm reasoning. *arXiv preprint*
 502 *arXiv:2505.14970*, 2025.

503 Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V
 504 Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation
 505 model post-training. *arXiv preprint arXiv:2501.17161*, 2025.

506 Huilin Deng, Ding Zou, Rui Ma, Hongchen Luo, Yang Cao, and Yu Kang. Boosting the general-
 507 ization and reasoning of vision language models with curriculum reinforcement learning. *arXiv*
 508 *preprint arXiv:2503.07065*, 2025.

509 Qi Feng, Yihong Liu, and Hinrich Schütze. Your pretrained model tells the difficulty itself: A
 510 self-adaptive curriculum learning paradigm for natural language understanding. *arXiv preprint*
 511 *arXiv:2507.09758*, 2025a.

512 Wenfeng Feng, Chuzhan Hao, Yuewei Zhang, Guochao Jiang, Jingyi Song, and Hao Wang. Air-
 513 rag: Autonomous strategic planning and reasoning steer retrieval augmented generation. *arXiv*
 514 *preprint arXiv:2501.10053*, 2025b.

515 Wenfeng Feng, Penghong Zhao, Guochao Jiang, Chuzhan Hao, Yuewei Zhang, and Hao Wang.
 516 Pvpo: Pre-estimated value-based policy optimization for agentic reasoning. *arXiv preprint*
 517 *arXiv:2508.21104*, 2025c.

518 Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cogni-
 519 tive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv*
 520 *preprint arXiv:2503.01307*, 2025.

521 Huan-ang Gao, Jiayi Geng, Wenyue Hua, Mengkang Hu, Xinzhe Juan, Hongzhang Liu, Shilong
 522 Liu, Jiahao Qiu, Xuan Qi, Yiran Wu, et al. A survey of self-evolving agents: On path to artificial
 523 super intelligence. *arXiv preprint arXiv:2507.21046*, 2025.

524 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 525 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 526 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

527 Hasan Abed Al Kader Hammoud, Kumail Alhamoud, Abed Hammoud, Elie Bou-Zeid, Marzyeh
 528 Ghassemi, and Bernard Ghanem. Train long, think short: Curriculum learning for efficient rea-
 529 soning. *arXiv preprint arXiv:2508.08940*, 2025.

530 Chuzhan Hao, Wenfeng Feng, Yuewei Zhang, and Hao Wang. Dynasearcher: Dynamic knowl-
 531 edge graph augmented search agent via multi-reward reinforcement learning. *arXiv preprint*
 532 *arXiv:2507.17365*, 2025.

540 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi
 541 Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun.
 542 Olympiadbench: A challenging benchmark for promoting AGI with olympiad-level bilingual
 543 multimodal scientific problems. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Pro-
 544 ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume
 545 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pp. 3828–3850. Associa-
 546 tion for Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.211. URL
 547 <https://doi.org/10.18653/v1/2024.acl-long.211>.

548 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 549 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 550 evaluation of large language models for code. In *The Thirteenth International Conference on
 551 Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025.
 552 URL <https://openreview.net/forum?id=chfJJYC3iL>.

553 Guochao Jiang, Ziqin Luo, Yuchen Shi, Dixuan Wang, Jiaqing Liang, and Deqing Yang. Toner:
 554 Type-oriented named entity recognition with generative language model. In Nicoletta Calzolari,
 555 Min-Yen Kan, Véronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (eds.), *Pro-
 556 ceedings of the 2024 Joint International Conference on Computational Linguistics, Language Re-
 557 sources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pp. 16251–16262.
 558 ELRA and ICCL, 2024. URL <https://aclanthology.org/2024.lrec-main.1412>.

559 Guochao Jiang, Ziqin Luo, Chengwei Hu, Zepeng Ding, and Deqing Yang. Mitigating out-of-
 560 entity errors in named entity recognition: A sentence-level strategy. In Owen Rambow, Leo
 561 Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert
 562 (eds.), *Proceedings of the 31st International Conference on Computational Linguistics, COLING
 563 2025, Abu Dhabi, UAE, January 19-24, 2025*, pp. 7754–7765. Association for Computational
 564 Linguistics, 2025a. URL <https://aclanthology.org/2025.coling-main.519>.

565 Guochao Jiang, Guofeng Quan, Zepeng Ding, Ziqin Luo, Dixuan Wang, and Zheng Hu. Flashthink:
 566 An early exit method for efficient reasoning. *arXiv preprint arXiv:2505.13949*, 2025b.

567 Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh
 568 Hakhamaneshi, Shishir G Patil, Matei Zaharia, et al. Llms can easily learn to reason from demon-
 569 strations structure, not content, is what matters! *arXiv preprint arXiv:2502.07374*, 2025.

570 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 571 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The
 572 Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria,
 573 May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=v8L0pN6EOi>.

574 Zicheng Lin, Tian Liang, Jiahao Xu, Qiuzhi Lin, Xing Wang, Ruilin Luo, Chufan Shi, Siheng
 575 Li, Yujiu Yang, and Zhaopeng Tu. Critical tokens matter: Token-level contrastive estimation
 576 enhances llm's reasoning capability. *arXiv preprint arXiv:2411.19943*, 2024.

577 Mingjie Liu, Shizhe Diao, Ximing Lu, Jian Hu, Xin Dong, Yejin Choi, Jan Kautz, and Yi Dong.
 578 Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models.
 579 *arXiv preprint arXiv:2505.24864*, 2025a.

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

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

586 Jiarui Lu, Thomas Holleis, Yizhe Zhang, Bernhard Aumayer, Feng Nan, Haoping Bai, Shuang Ma,
 587 Shen Ma, Mengyu Li, Guoli Yin, Zirui Wang, and Ruoming Pang. Toolsandbox: A stateful,
 588 conversational, interactive evaluation benchmark for LLM tool use capabilities. In Luis Chiruzzo,
 589 Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics:*

594 NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025, pp. 1160–1183. Association for Computational Linguistics, 2025. doi: 10.18653/V1/2025.FINDINGS-NAACL.65.
 595 URL <https://doi.org/10.18653/v1/2025.findings-naacl.65>.

596

597 Youssef Mroueh. Reinforcement learning with verifiable rewards: Grpo's effective loss, dynamics, and success amplification. *arXiv preprint arXiv:2503.06639*, 2025.

598

599

600 Shubham Parashar, Shurui Gui, Xiner Li, Hongyi Ling, Sushil Vemuri, Blake Olson, Eric Li, Yu Zhang, James Caverlee, Dileep Kalathil, et al. Curriculum reinforcement learning from easy to hard tasks improves llm reasoning. *arXiv preprint arXiv:2506.06632*, 2025.

601

602

603 John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In Yoshua Bengio and Yann LeCun (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1506.02438>.

604

605

606

607

608 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

609

610

611 Burr Settles. Active learning literature survey. 2009.

612

613 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

614

615

616 Junhong Shen, Hao Bai, Lunjun Zhang, Yifei Zhou, Amrith Setlur, Shengbang Tong, Diego Caples, Nan Jiang, Tong Zhang, Ameet Talwalkar, et al. Thinking vs. doing: Agents that reason by scaling test-time interaction. *arXiv preprint arXiv:2506.07976*, 2025.

617

618

619 Zhuocheng Shen. Llm with tools: A survey. *arXiv preprint arXiv:2409.18807*, 2024.

620

621 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*, pp. 1279–1297. ACM, 2025. doi: 10.1145/3689031.3696075. URL <https://doi.org/10.1145/3689031.3696075>.

622

623

624

625

626 Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *Int. J. Comput. Vis.*, 130(6):1526–1565, 2022. doi: 10.1007/S11263-022-01611-X. URL <https://doi.org/10.1007/s11263-022-01611-x>.

627

628

629

630 Richard S. Sutton and Andrew G. Barto. Reinforcement learning: An introduction. *IEEE Trans. Neural Networks*, 9(5):1054–1054, 1998. doi: 10.1109/TNN.1998.712192. URL <https://doi.org/10.1109/TNN.1998.712192>.

631

632

633 Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*, 2025.

634

635

636 Georgios Tzannetos, Bárbara Gomes Ribeiro, Parameswaran Kamalaruban, and Adish Singla. Proximal curriculum for reinforcement learning agents. *Trans. Mach. Learn. Res.*, 2023, 2023. URL <https://openreview.net/forum?id=8WUyeeMxMH>.

637

638

639

640 Jean Vassoyan, Nathanaël Beau, and Roman Plaud. Ignore the KL penalty! boosting exploration on critical tokens to enhance RL fine-tuning. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, pp. 6108–6118. Association for Computational Linguistics, 2025. doi: 10.18653/V1/2025.FINDINGS-NAACL.340. URL <https://doi.org/10.18653/v1/2025.findings-naacl.340>.

641

642

643

644

645 Peng-Yuan Wang, Tian-Shuo Liu, Chenyang Wang, Yi-Di Wang, Shu Yan, Cheng-Xing Jia, Xu-Hui Liu, Xin-Wei Chen, Jia-Cheng Xu, Ziniu Li, et al. A survey on large language models for mathematical reasoning. *arXiv preprint arXiv:2506.08446*, 2025a.

646

647

648 Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Trans. Pattern
649 Anal. Mach. Intell.*, 44(9):4555–4576, 2022. doi: 10.1109/TPAMI.2021.3069908. URL <https://doi.org/10.1109/TPAMI.2021.3069908>.

650

651 Zhenting Wang, Guofeng Cui, Yu-Jhe Li, Kun Wan, and Wentian Zhao. Dump: Auto-
652 mated distribution-level curriculum learning for rl-based llm post-training. *arXiv preprint
653 arXiv:2504.09710*, 2025b.

654

655 Zihan Wang, Kangrui Wang, Qineng Wang, Pingyue Zhang, Linjie Li, Zhengyuan Yang, Xing Jin,
656 Kefan Yu, Minh Nhat Nguyen, Licheng Liu, et al. Ragen: Understanding self-evolution in llm
657 agents via multi-turn reinforcement learning. *arXiv preprint arXiv:2504.20073*, 2025c.

658

659 Zhiheng Xi, Wenxiang Chen, Boyang Hong, Senjie Jin, Rui Zheng, Wei He, Yiwen Ding, Shichun
660 Liu, Xin Guo, Junzhe Wang, Honglin Guo, Wei Shen, Xiaoran Fan, Yuhao Zhou, Shihan Dou,
661 Xiao Wang, Xinbo Zhang, Peng Sun, Tao Gui, Qi Zhang, and Xuanjing Huang. Training large
662 language models for reasoning through reverse curriculum reinforcement learning. In *Forty-first
663 International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*.
664 OpenReview.net, 2024. URL <https://openreview.net/forum?id=t82Y3fmRtk>.

665 Fengli Xu, Qianyue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan,
666 Jiahui Gong, Tianjian Ouyang, Fanjin Meng, et al. Towards large reasoning models: A survey of
667 reinforced reasoning with large language models. *arXiv preprint arXiv:2501.09686*, 2025a.

668 Yixuan Even Xu, Yash Savani, Fei Fang, and Zico Kolter. Not all rollouts are useful: Down-sampling
669 rollouts in llm reinforcement learning. *arXiv preprint arXiv:2504.13818*, 2025b.

670 Jianhao Yan, Yafu Li, Zican Hu, Zhi Wang, Ganqu Cui, Xiaoye Qu, Yu Cheng, and Yue Zhang.
671 Learning to reason under off-policy guidance. *arXiv preprint arXiv:2504.14945*, 2025.

672

673 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
674 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint
675 arXiv:2505.09388*, 2025a.

676 Dayu Yang, Tianyang Liu, Daoan Zhang, Antoine Simoulin, Xiaoyi Liu, Yuwei Cao, Zhaopu Teng,
677 Xin Qian, Grey Yang, Jiebo Luo, et al. Code to think, think to code: A survey on code-enhanced
678 reasoning and reasoning-driven code intelligence in llms. *arXiv preprint arXiv:2502.19411*,
679 2025b.

680 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. Demystifying long chain-
681 of-thought reasoning in llms. *arXiv preprint arXiv:2502.03373*, 2025.

682

683 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian
684 Fan, Gaohong Liu, Lingjun Liu, et al. Dapo: An open-source llm reinforcement learning system
685 at scale. *arXiv preprint arXiv:2503.14476*, 2025.

686 Yu Yue, Yufeng Yuan, Qiying Yu, Xiaochen Zuo, Ruofei Zhu, Wenyuan Xu, Jiaze Chen, Chengyi
687 Wang, TianTian Fan, Zhengyin Du, et al. Vapo: Efficient and reliable reinforcement learning for
688 advanced reasoning tasks. *arXiv preprint arXiv:2504.05118*, 2025.

689

690 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu, Keqing He, Zejun Ma, and Junxian He. Simplerl-
691 zoo: Investigating and taming zero reinforcement learning for open base models in the wild. *arXiv
692 preprint arXiv:2503.18892*, 2025.

693 Enci Zhang, Xingang Yan, Wei Lin, Tianxiang Zhang, and Lu Qianchun. Learning like humans: Ad-
694 vancing llm reasoning capabilities via adaptive difficulty curriculum learning and expert-guided
695 self-reformulation. *arXiv preprint arXiv:2505.08364*, 2025a.

696

697 Guibin Zhang, Hejia Geng, Xiaohang Yu, Zhenfei Yin, Zaibin Zhang, Zelin Tan, Heng Zhou,
698 Zhongzhi Li, Xiangyuan Xue, Yijiang Li, et al. The landscape of agentic reinforcement learning
699 for llms: A survey. *arXiv preprint arXiv:2509.02547*, 2025b.

700

701 Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang, Weixu Zhang, Wenyue Hua, Haolun Wu, Zhihan
702 Guo, Yufei Wang, Niklas Muennighoff, et al. A survey on test-time scaling in large language
703 models: What, how, where, and how well? *arXiv preprint arXiv:2503.24235*, 2025c.

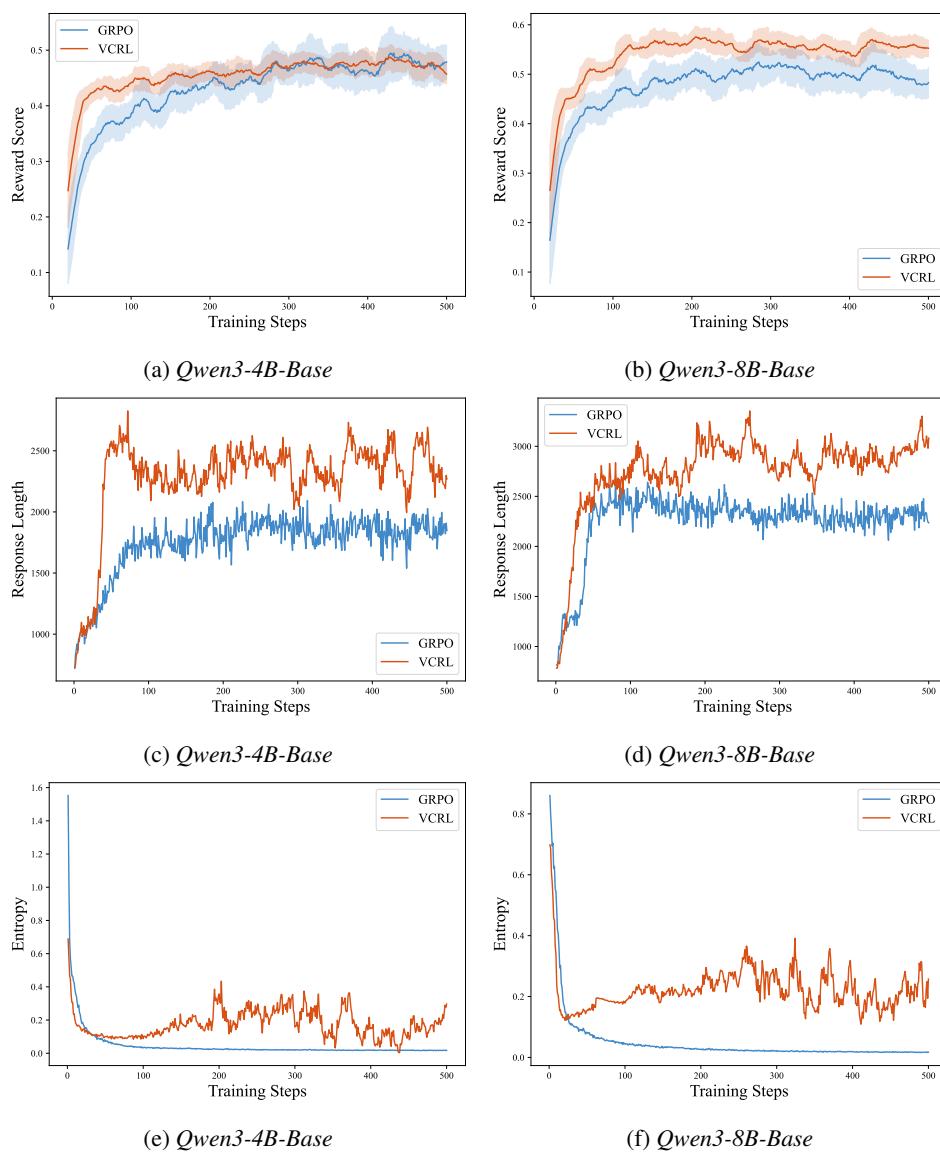


Figure 4: The metric curves of reward score, response length, and entropy of VCRL over GRPO based on *Qwen3-4B-Base* and *Qwen3-8B-Base*, which show the dynamics of RL training and serve as essential monitoring indicators to identify potential issues.

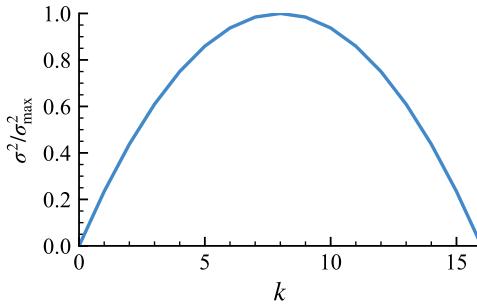
Zikang Zhang, Wangjie You, Tianci Wu, Xinrui Wang, Juntao Li, and Min Zhang. A survey of generative information extraction. In Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (eds.), *Proceedings of the 31st International Conference on Computational Linguistics, COLING 2025, Abu Dhabi, UAE, January 19-24, 2025*, pp. 4840–4870. Association for Computational Linguistics, 2025d. URL <https://aclanthology.org/2025.coling-main.324/>.

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. *arXiv preprint arXiv:2507.18071*, 2025.

756 **A TRAINING DYNAMICS**
757

758 Compared to GRPO, VCRL introduces two main techniques to improve training efficiency. To
759 further understand their effects, we show the training dynamics shown in Figure 4, including reward
760 score, response length, and entropy. For the reward score curve, in order to simultaneously measure
761 their stability in training dynamics, we use moving average and rolling standard deviation with a
762 window size of 20 for visualization.

- 764 • **Reward Score** during training is closely linked to training stability and performance, as
765 shown in Figure 4a and Figure 4b. For both VCRL and GRPO, the reward score rises
766 quickly in the early stages and then slowly improves. For *Qwen3-4B-Base*, before about
767 270 training steps, VCRL’s reward score is much higher than GRPO’s. For *Qwen3-8B-Base*,
768 the reward score of VCRL is significantly higher than that of GRPO throughout the
769 training process. Once the reward score stabilizes, VCRL shows much smaller fluctuations
770 than GRPO, as seen in the shaded areas. This highlights VCRL’s advantage in training
771 stability.
- 772 • **Response Length** relates to how much the model can explore, as shown in Figure 4c and
773 Figure 4d. Longer responses help the model develop more complex reasoning during
774 training and boost performance. In the first 100 steps, VCRL and GRPO both show a rapid
775 increase in response length, then level off and fluctuate. VCRL’s response length grows
776 much faster early on, especially in first 50 steps, due to the training of high- p samples. Af-
777 ter stabilizing, VCRL maintains noticeably longer responses, giving the model more room
778 to explore and optimize its performance.
- 779 • **Entropy** shows how uncertain the model is in its generation ability, as seen in Figure 4e
780 and Figure 4f. For efficient training, entropy should stay at a reasonable level. If entropy is
781 too low, the model becomes too deterministic and loses its ability to explore. For GRPO,
782 entropy quickly drops below 0.1 within 50 steps and stays very low. In contrast, VCRL
783 keeps entropy at a reasonable level throughout training, which encourages the model to
784 keep exploring.

785 **B VARIANCE AS A DIFFICULTY METRIC**
786

799 Figure 5: The curve showing how $p = \frac{\sigma^2}{\sigma_{\max}^2}$ changes with the number of successful rollouts k based
800 on group size $G = 16$.

802 Compared with generation entropy, it is more reasonable to use group variance to measure the diffi-
803 culty of the current sample for the current training model in VCRL or GRPO. Group reward variance
804 is grounded in its unique ability to identify samples at the cusp of the model’s current capabilities.

806 For a binary reward system (correct/incorrect), variance exhibits a non-monotonic, U-shaped rel-
807 ationship with sample difficulty, as shown in Figure 5. Low variance occurs at two extremes. If a
808 sample is too easy, the model consistently succeeds (e.g., all 16 rollouts get a reward of 1), leading
809 to near-zero variance. If a sample is too hard, the model consistently fails (all rewards are 0), also
leading to near-zero variance. Peaks when the model’s success rate is approximately 50% (e.g., 8

810 rollouts succeed and 8 fail). This indicates maximum uncertainty and signifies that the sample is at
 811 the precise frontier of the model’s ability.
 812

813 While related to uncertainty, policy generation entropy measures the diversity of the model’s actions
 814 (tokens). High entropy could mean the model is exploring, but it does not directly map to task-level
 815 success. A model could be highly uncertain (high entropy) while generating non-sensical responses
 816 that all lead to a reward of 0. Variance, on the other hand, is directly related to the final outcome of
 817 the task (the reward), making it a more direct measure of the difficulty relevant to learning. By using
 818 a single indicator of group variance, it is possible to filter samples with high uncertainty results,
 819 while this task is difficult to accomplish based on the generation entropy.
 820

821 C POLICY GRADIENT REDUCTION

823 According to Equation 2 and Equation 11, we give the following theorem:
 824

825 **Theorem 1.** *For policy gradient algorithm GRPO and VCRL, from the policy gradient norm
 826 perspective, the training of VCRL is more stable than that of GRPO in the expectation, that is,
 827 $\mathbb{E}_{VCRL} [\|\nabla_{\theta} \log \pi_{\theta}\|] \leq \mathbb{E}_{GRPO} [\|\nabla_{\theta} \log \pi_{\theta}\|]$.*
 828

829 *Proof.* We first give the gradient form of the GRPO objective function (clipping is omitted for
 830 brevity) with Policy Gradient Theorem (Sutton & Barto, 1998):
 831

$$\begin{aligned} \nabla_{\theta} \mathcal{J}_{GRPO}(\theta) &= \nabla_{\theta} \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{old}(\cdot|x)}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} r_{i,t}(\theta) \hat{A}_{i,t} \right] \\ &= \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{old}(\cdot|x)}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} r_{i,t}(\theta) \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t}) \right], \end{aligned} \quad (12)$$

840 where $r_{i,t}(\theta) = \frac{\pi_{\theta}(y_{i,t}|x, y_{i,<t})}{\pi_{\theta_{old}}(y_{i,t}|x, y_{i,<t})}$ is the importance sampling ratio.
 841

842 We can also derive the gradient of the VCRL objective as follows:
 843

$$\begin{aligned} \nabla_{\theta} \mathcal{J}_{VCRL}(\theta) &= \nabla_{\theta} \mathbb{E}_{x \sim \mathcal{D} \cup \mathcal{M}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{old}(\cdot|x)}} \left[\frac{1}{G} \sum_{i=1}^G \frac{\mathbb{I}(p_i = \frac{\sigma_i^2}{\sigma_{\max}^2} \geq \kappa)}{|y_i|} \sum_{t=1}^{|y_i|} r_{i,t}(\theta) \hat{A}_{i,t} \right] \\ &= \mathbb{E}_{x \sim \mathcal{D} \cup \mathcal{M}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{old}(\cdot|x)}} \left[\frac{1}{G} \sum_{i=1}^G \frac{\mathbb{I}(p_i = \frac{\sigma_i^2}{\sigma_{\max}^2} \geq \kappa)}{|y_i|} \sum_{t=1}^{|y_i|} r_{i,t}(\theta) \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t}) \right]. \end{aligned} \quad (13)$$

853 To align the gradients of the two, we use importance sampling to rewrite the gradient of VCRL to
 854 remove the term of memory bank \mathcal{M} :
 855

$$\begin{aligned} \nabla_{\theta} \mathcal{J}_{VCRL}(\theta) &= \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{old}(\cdot|x)}} \\ &\quad \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} r_{i,t}(\theta) \hat{A}_{i,t} \frac{\mathbb{P}(x \in \mathcal{D} \cup \mathcal{M})}{\mathbb{P}(x \in \mathcal{D})} \mathbb{I}(p_i \geq \kappa) \nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t}) \right]. \end{aligned} \quad (14)$$

862 Note that the blue part in the Equation 14 is the key to affecting the contribution of the policy
 863 gradient term to the overall gradient.

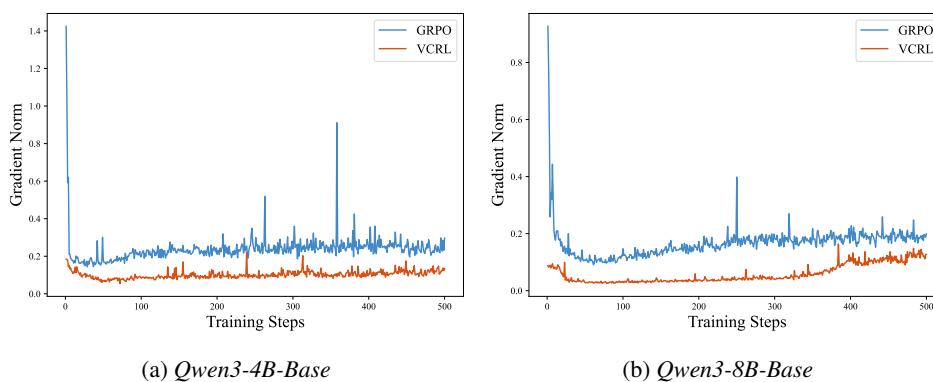


Figure 6: The training dynamics of objective gradient norm $\|\nabla_{\theta}\mathcal{J}(\theta)\|$ of VCRL over GRPO based on *Qwen3-4B-Base* and *Qwen3-8B-Base*.

We simplify the policy gradient terms in Equation 12 and Equation 14 into the following form for comparison:

$$\mathbb{E}_{\text{GRPO}} [\|\nabla_{\theta} \log \pi_{\theta}\|] = \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}(\cdot|x)}} [\|\nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t})\|], \quad (15)$$

$$\begin{aligned} \mathbb{E}_{\text{VCRL}} [\|\nabla_{\theta} \log \pi_{\theta}\|] &= \mathbb{E}_{x \sim \mathcal{D}, \{y_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}(\cdot|x)}} \\ &\quad \left[\left\| \frac{\mathbb{P}(x \in \mathcal{D} \cup \mathcal{M})}{\mathbb{P}(x \in \mathcal{D})} \mathbb{I}(p_i \geq \kappa) \nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t}) \right\| \right]. \end{aligned} \quad (16)$$

For the training sample x , the sampling of the event $x \in \mathcal{D}$ is uniform, so $\mathbb{P}(x \in \mathcal{D}) = \frac{1}{|\mathcal{D}|}$. And according to the nature of sampling probability, we can get $\mathbb{P}(x \in \mathcal{D} \cup \mathcal{M}) \leq \mathbb{P}(x \in \mathcal{D})$. Based on the value range of the indicator function, we can also get $\mathbb{I}(p_i \geq \kappa) \leq 1$. Using the homogeneity of the norm and above results:

$$\begin{aligned} &\left\| \frac{\mathbb{P}(x \in \mathcal{D} \cup \mathcal{M})}{\mathbb{P}(x \in \mathcal{D})} \mathbb{I}(p_i \geq \kappa) \nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t}) \right\| \\ &= \frac{\mathbb{P}(x \in \mathcal{D} \cup \mathcal{M})}{\mathbb{P}(x \in \mathcal{D})} \mathbb{I}(p_i \geq \kappa) \|\nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t})\| \\ &\leq \mathbb{I}(p_i \geq \kappa) \|\nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t})\| \\ &\leq \|\nabla_{\theta} \log \pi_{\theta}(y_{i,t}|x, y_{i,<t})\|, \end{aligned}$$

which completes the proof. \square

Theorem 1 provides a theoretical guarantee for the training stability of VCRL compared to GRPO. To further illustrate the training stability of VCRL from the perspective of gradient norm, we show the training dynamics as shown in the Figure 6.

Figure 6 provides an empirical validation of our proposed VCRL's stability by visualizing the norm of the objective function's gradient, $\|\nabla_{\theta}\mathcal{J}(\theta)\|$, over the training steps. We compare VCRL against the GRPO baseline on two model scales: *Qwen3-4B-Base* and *Qwen3-8B-Base*. The empirical results unequivocally demonstrate the superiority of VCRL in maintaining a well-behaved optimization trajectory. Specifically, VCRL's gradient norm remains consistently confined to a lower and narrower band, indicating that the policy updates are more measured and stable. Furthermore, the VCRL curve is notably smoother, with significantly fewer and less pronounced transient spikes compared to the GRPO baseline. The frequent, high-magnitude oscillations observed in GRPO's gradient norm are indicative of a more challenging optimization landscape, which can lead to inefficient and unstable training. We posit that the demonstrably smaller and more stable gradient norm engendered by VCRL is an important contributor to its enhanced training efficiency and robust performance.

918 Table 3: Performance comparison on mathematical reasoning benchmarks using Qwen2.5-7B.
919

920 Method	921 AIME-2024	922 AIME-2025	923 MATH500	924 AMC23	925 Avg.
926 Base Model	927 16.67	928 16.67	929 69.40	930 32.50	931 33.81
932 + GRPO	933 26.67	934 16.67	935 72.40	936 45.00	937 40.19
938 + ADCL	939 33.33	940 30.00	941 76.20	942 55.00	943 48.63
945 + ADCL & EGSR	946 36.67	947 33.33	948 81.80	949 55.00	950 51.70
953 + VCRL (Ours)	954 35.83	955 32.71	956 79.73	957 61.22	958 52.37

927 **D BASELINE COMPARISON**

930 To further demonstrate the competitiveness of our proposed method VCRL, we provide a supplementary comparison with recent curriculum reinforcement learning methods, specifically focusing
931 on ADCL (Zhang et al., 2025a). ADCL is a difficulty-based curriculum method that employs peri-
932 odic data reordering.

933 To ensure a fair comparison, we align our experimental settings strictly with the configurations
934 reported in ADCL. We utilize the Qwen2.5-7B model as the base policy model. During training, we
935 sample 8 rollouts for each query, with a global batch size of 1,024, a fixed learning rate of 1×10^{-6} , a
936 maximum response length of 4,096, and a temperature parameter of 0.7. We evaluated performance
937 across our four challenging mathematical benchmarks: AIME-2024, AIME-2025, MATH500, and
938 AMC23.

939 As shown in Table 3, VCRL achieves the highest average performance (52.37), surpassing both
940 the standard ADCL (48.63) and its enhanced version ADCL & EGSR (51.70). Notably, VCRL
941 demonstrates a significant advantage on the AMC23 benchmark, outperforming ADCL by over 6
942 points. Beyond numerical improvements, VCRL offers distinct methodological advantages. ADCL
943 relies on difficulty-based reordering of the training data (e.g., reordering every 100 training steps)
944 and operates under the assumption of a known future data stream. In contrast, VCRL is inherently
945 more flexible and dynamic. It does not require pre-processing or assumptions about the global
946 data distribution. Instead, VCRL performs filtering based on the group variance within the current
947 training batch. This allows VCRL to adaptively select high-value samples without the computational
948 overhead of periodic global reordering or the need for pre-computed difficulty metrics.

950 **E CODE BENCHMARK RESULTS**

951 To validate the effectiveness of VCRL on tasks outside of mathematical reasoning, we also conduct
952 experiments on code generation tasks. Specifically, we conduct code-based RL training based on
953 the *ProRL-1.5B-v2* (Liu et al., 2025a) to broaden the generalization of VCRL in terms of model
954 selection. Throughout the model evaluation, we allow the model to execute Python code based on
955 the Sandbox Fusion⁶ environment. We allow the model to generate and execute code to assist in
956 completing tasks, building upon the benchmarks of AIME-2024 and AIME-2025. We also use the
957 LiveCodeBench (Jain et al., 2025), for a total of three benchmarks to comprehensively evaluate the
958 performance of the training methods in code generation. The quantitative results are summarized in
959 Table 4. VCRL achieves state-of-the-art performance across all benchmarks with an average score
960 of 42.98, significantly outperforming the base model (30.99) and surpassing the strongest baseline,
961 GSPO (40.86). Notably, on LiveCodeBench, VCRL reaches 35.20, confirming that our VCRL
962 effectively generalizes to code generation tasks.

963 **F SENSITIVITY ANALYSIS**

964 To better illustrate how the hyperparameters in VCRL affect model performance, we present the
965 results of the sensitivity analysis experiments shown in the Table 5 and 6. Specifically, Table 5
966 shows the sensitivity analysis of the threshold parameter κ in variance-based dynamic sampling

967
968 ⁶<https://bytedance.github.io/SandboxFusion>

972
973
974 Table 4: Performance comparison on code generation benchmarks using *ProRL-1.5B-v2*.
975
976
977
978
979
980

Method	AIME-2024 w Code	AIME-2025 w Code	LiveCodeBench	Avg.
Base Model	40.00	27.50	25.48	30.99
+ GRPO	50.62	33.95	31.11	38.56
+ DAPO	52.50	34.17	30.79	39.15
+ GSPO	53.13	36.88	32.58	40.86
+ VCRL	54.79	38.96	35.20	42.98

981
982 Table 5: Sensitivity analysis based on the threshold parameter κ in variance-based dynamic sampling
983 in VCRL using *Qwen3-8B-Base*.
984
985

VCRL κ	AIME-2024	AIME-2025	MATH500	OlympiadBench	AMC23	Avg.
0.4	30.42	26.88	92.00	60.55	68.98	55.77
0.6	25.00	23.12	90.71	57.55	64.91	52.26
0.7	30.42	27.29	92.78	61.94	69.73	56.43
0.8	34.38	27.08	91.99	60.21	75.15	57.76
0.9	32.08	26.25	92.68	60.92	71.69	56.72
0.95	28.75	26.25	92.36	60.73	70.41	55.70

993
994 of VCRL, and Table 6 shows the sensitivity analysis of the rollout group size G in VCRL. Both
995 experiments are conducted based on *Qwen3-8B-Base*.
996
997998 Table 5 illustrates the impact of the variance-based dynamic sampling threshold κ on VCRL performance.
999 The results exhibit a clear trend where neither excessively small nor large values of κ yield
1000 optimal outcomes. Specifically, when κ is set to a low value, the average performance will decrease,
1001 which may be due to overly lenient sampling, retaining too many samples that may not be of use.
1002 Conversely, setting κ too high also leads to a performance decline, as the mechanism may sample
1003 out too many samples with high variance. The best overall performance is achieved at $\kappa = 0.8$,
1004 which strikes an optimal balance, delivering the highest average score of 57.76 and demonstrating
1005 robust performance across diverse benchmarks, particularly on AMC23.
10061007 Table 6 presents the impact of the rollout group size G in VCRL. In contrast to the threshold parameter κ ,
1008 we observe a positive correlation between the rollout group size G and model performance.
1009 Increasing G consistently leads to better results across all benchmarks in average. Specifically, enlarging
1010 the group size from 4 to 16 yields a substantial performance boost, raising the average score
1011 from 44.17 to 57.76. This significant improvement indicates that a sufficient number of rollouts is
1012 essential for VCRL to accurately estimate the group variance, thereby ensuring the reliability of the
1013 variance-based dynamic sampling. While the best average performance (59.25) is achieved at the
1014 largest setting ($G = 32$), the marginal gains begin to diminish after $G = 16$ (improving by only
1015 1.49 despite doubling the computational cost). This suggests that $G = 16$ already provides a robust
1016 estimation for effective learning, balancing performance with computational efficiency.
10171018
1019

G COMPUTATION COST ANALYSIS

10201021 In VCRL, we introduce two key components: variance-based dynamic sampling and replay learning.
1022 In replay learning, we use a memory bank \mathcal{M} to maintain and supplement high- p training
1023 samples to improve training efficiency, but this also introduces potential additional computational
1024 cost. To compare the computational cost of VCRL with other benchmark methods in terms of actual
1025 computation time, we list the specific GPU computation times for the *Qwen3-8B-Base* experiments
in Table 1, as shown in Table 7.
10261027 As shown in Table 7, the introduction of the replay learning in VCRL incurs a computational over-
1028 head, increasing the total training time for *Qwen3-8B-Base* from 90.72 hours (GRPO) to 112.01
1029 hours—an increase of approximately 23.47%. However, a direct comparison of training steps does
1030 not fully account for this time difference. To evaluate performance under a fixed computational
1031 budget, we align the methods based on wall-clock time: approximately 380 steps of VCRL training
1032

1026 Table 6: Sensitivity analysis based on the rollout group size G in VCRL using *Qwen3-8B-Base*.
1027

1028 VCRL G	1029 AIME-2024	1030 AIME-2025	1031 MATH500	1032 OlympiadBench	1033 AMC23	1034 Avg.
1030 4	18.95	17.29	82.71	47.75	54.14	44.17
1031 8	23.54	18.75	85.21	49.49	58.06	47.01
1032 12	28.54	26.67	87.77	55.48	65.06	52.70
1033 16	34.38	27.08	91.99	60.21	75.15	57.76
1034 24	34.38	27.50	93.74	60.90	77.86	58.88
1035 32	35.83	29.79	93.46	60.88	76.28	59.25

1036 Table 7: Comparison of training time using *Qwen3-8B-Base*. We report the total wall-clock time (in
1037 hours) required for the training process.
1038

1039 Method	1040 Training Time (hours)
1041 GRPO	90.72
1042 DAPO	48.68
1043 GSPO	89.53
1044 VCRL	112.01

1045
1046 correspond to 500 steps of GRPO training. Referring to the training trajectories in Figure 3, we ob-
1047 serve that even at this earlier stage (380 steps), VCRL remains highly competitive and outperforms
1048 GRPO trained for the full 500 steps. This demonstrates that although VCRL has a higher per-step
1049 cost, it achieves greater sample efficiency and utilizes the computational budget more effectively.
10501051 Table 8 details the changes in the effective batch size after variance-based dynamic sampling
1052 throughout the training process. In the early stages (e.g., steps 0–50), the mean batch size is rel-
1053 atively small (48–56), indicating that the variance-based sampling mechanism is actively removing a
1054 significant portion of samples to mitigate instability. As the policy optimizes and training stabilizes,
1055 the effective batch size progressively increases, reaching 128 in the final stages (steps 451–500).
1056 This trend demonstrates that as the training matures, the dynamic sampling mechanism effectively
1057 fades out, allowing the model to utilize the full batch of data when the variance-based selection is
1058 no longer necessary.
10591060 Table 8: Evolution of the effective mean batch size after variance-based dynamic sampling across
1061 training steps in VCRL for *Qwen3-8B-Base*.
1062

1063 Step Interval	1064 Mean Batch Size
1064 0–20	48
1065 21–50	56
1066 51–100	64
1067 101–150	64
1068 151–200	80
1069 201–250	96
1070 251–300	104
1071 301–350	112
1072 351–400	120
1073 451–500	128