

000 001 002 003 004 005 006 007 008 009 010 011 012 UNLOCKING EXPLORATION IN RLVR: UNCERTAINTY- AWARE ADVANTAGE SHAPING FOR DEEPER REASON- ING

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
Anonymous authors
042
043
044
045
046
047
048
049
050
051
052
053
Paper under double-blind review

ABSTRACT

042
043
044
045
046
047
048
049
050
051
052
053
Reinforcement Learning with Verifiable Rewards (RLVR) has shown significant promise for enhancing the reasoning capabilities of large language models (LLMs). However, prevailing algorithms like GRPO broadcast a uniform advantage signal across all tokens in a sequence. This coarse-grained approach overlooks the pivotal role of uncertain, high-stakes decisions during reasoning, leading to inefficient exploration and the well-documented problem of entropy collapse. To address this, we introduce **UnCertainty-aware Advantage Shaping (UCAS)**, a model-free method that refines credit assignment by leveraging the model’s internal uncertainty signals. UCAS operates in two stages: it first modulates the response-level advantage using the model’s overall self-confidence, and then applies a token-level penalty based on raw logit certainty. This dual mechanism encourages exploration of high-uncertainty paths that yield correct answers while penalizing overconfident yet erroneous reasoning, effectively balancing the exploration-exploitation trade-off. Extensive experiments on five mathematical reasoning benchmarks show that UCAS significantly outperforms strong RLVR baselines across multiple model scales, including 1.5B and 7B. Our analysis confirms that UCAS not only achieves higher rewards but also promotes greater reasoning diversity and successfully mitigates entropy collapse.

1 INTRODUCTION

042
043
044
045
046
047
048
049
050
051
052
053
Reinforcement learning (RL) has recently become a cornerstone for enhancing the complex reasoning abilities of Large Language Models (LLMs), moving beyond simple pattern matching toward more robust problem-solving. Among the various RL approaches, Reinforcement Learning with Verifiable Rewards (RLVR) has proven particularly effective. In this paradigm, a policy model explores a vast solution space and receives feedback from verifiable signals, such as the correctness of a final answer in mathematical reasoning. This direct feedback loop has enabled policy optimization algorithms like Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to achieve substantial performance gains, powering state-of-the-art systems such as DeepSeek-R1 (Guo et al., 2025).

042
043
044
045
046
047
048
049
050
051
052
053
However, the success of RLVR reveals a critical underlying tension: the trade-off between precision and diversity. While methods like GRPO excel at increasing the probability of generating correct answers, they often do so at the cost of exploration. Due to the absence of a critic model, the learning signal in GRPO, which applies a single uniform advantage across all tokens, provides an indiscriminate and overly coarse form of credit assignment. It rewards all steps of a correct path equally and penalizes all steps of an incorrect one, failing to distinguish crucial reasoning leaps from trivial ones. This coarse-grained feedback drives the policy to converge prematurely on a small set of “safe” high-reward trajectories. A common side effect is **entropy collapse** (Cui et al., 2025b), where the output distribution contracts, reducing solution diversity and impairing performance on complex problems that demand novel reasoning strategies.

042
043
044
045
046
047
048
049
050
051
052
053
Previous studies (Wang et al., 2023; Lightman et al., 2024; Chen et al., 2024; Zhang et al., 2024; Sun et al., 2025a) have attempted to employ process-level reward models to deliver more fine-grained signals. However, as DeepSeek (Guo et al., 2025) points out, training fine-grained reward models is

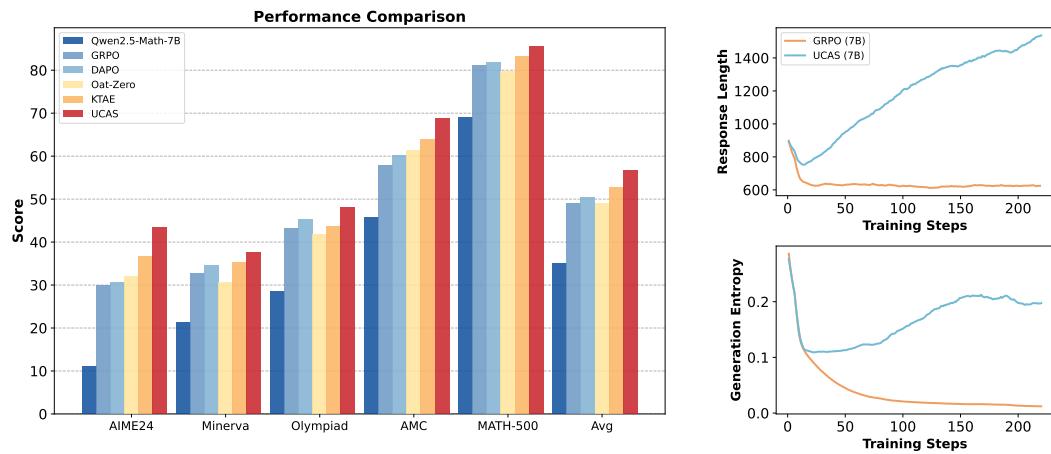


Figure 1: *Left*: Benchmark results across five math reasoning datasets, where our UCAS consistently outperforms RLVR baselines trained on models of the same parameter scale. *Right*: Training trajectories of UCAS and GRPO on Qwen2.5-Math-7B, showing that UCAS experiences an initial decline but subsequently rises in response length and generation entropy as training progresses. In contrast, GRPO exhibits a continual downward trend in entropy, reflecting the phenomenon of entropy collapse.

costly, difficult to scale, limited in its ability to provide accurate signals, and vulnerable to reward hacking. Some recent efforts (Chen et al., 2025; Cheng et al., 2025; Wang et al., 2025a) have tried to incorporate entropy-based feedback to enhance advantages, such as integrating semantic entropy or policy entropy related to the response into advantage calculations. Yet, most studies either pursue low entropy to improve accuracy or encourage high entropy to maintain exploration, lacking fine-grained modeling of the relationship between responses and their policy entropy.

To address the above challenges, we propose an **UnCertainty-aware Advantage Shaping (UCAS)**, a model-free method that refines credit assignment in RLVR by leveraging the model’s intrinsic uncertainty. UCAS is designed to resolve the precision–diversity dilemma by reshaping the advantage signal at two complementary levels. At the **response level**, UCAS modulates the sequence-level advantage using the model’s overall self-confidence, amplifying rewards for correct-but-uncertain responses and penalties for incorrect-but-confident ones. At the **token level**, it further introduces a certainty-based penalty derived directly from raw logits, discouraging local overconfidence while preserving diversity in reasoning. Collectively, these mechanisms promote exploration of uncertain but potentially fruitful reasoning paths, while efficiently suppressing confidently wrong solutions. Extensive experiments on five mathematical reasoning benchmarks demonstrate that UCAS consistently outperforms strong RLVR baselines at both the 1.5B and 7B model scales. Beyond reward improvements, UCAS fosters greater reasoning diversity and substantially mitigates entropy collapse, confirming the effectiveness of uncertainty as a fine-grained learning signal.

Our contributions can be summarized as follows:

- We propose UCAS, an extra-model-free fine-grained advantage shaping mechanism based on internal confidence signals, which performs uncertainty-aware advantage adjustment at both response and token levels.
- We provide a novel mechanism to adaptively calibrate advantages based on uncertainty, enabling steady reward gains, longer reasoning chains, and entropy recovery, thus preventing entropy collapse in RLVR and improving reasoning accuracy.
- Extensive experiments on multiple mathematical reasoning benchmarks demonstrate that UCAS significantly improves model reasoning performance, validating its effectiveness in enhancing exploration diversity and optimization outcomes.

108 2 BACKGROUND: REINFORCEMENT LEARNING WITH VERIFIABLE
 109 REWARDS
 110

111 In the training of large language models, early mainstream reinforcement learning alignment meth-
 112 ods primarily relied on PPO. By introducing a clipping ratio into the objective function, PPO sta-
 113 bilizes training by constraining the magnitude of policy updates. This method has been widely
 114 adopted in Reinforcement Learning from Human Feedback (RLHF), where reward models provide
 115 preference-based scores to gradually shape model behavior. However, PPO exhibits key limitations:
 116 it depends on critic-based value estimation and requires large-scale preference annotation, both of
 117 which are costly and prone to noise accumulation.

118 To overcome these limitations, recent research has introduced RLVR. RLVR converts open-ended
 119 outputs into programmatically checkable signals, such as numerical consistency in mathematics,
 120 unit-test pass rates in code generation, or formal constraint satisfaction (Su et al., 2025; Wang et al.,
 121 2025b), thereby avoiding the noise and cost of preference models. By forming a closed loop of
 122 model–environment–verifier, RLVR enables policies to be updated directly from binary or graded
 123 correctness signals, improving both sample efficiency and reproducibility in structured reasoning
 124 tasks.

125 In the concrete implementation of RLVR, GRPO (Shao et al., 2024) emerges as a representative
 126 algorithm. Unlike PPO, which relies on critic-based value estimation, GRPO computes advantages
 127 by normalizing group-level verifiable rewards and updates the policy directly.

128 Formally, the objective is given by:

$$130 \quad \mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{q \sim \mathcal{D}, o \sim \pi_{\theta_{\text{old}}}} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \min(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t}) - \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right] \quad (1)$$

135 where

$$136 \quad r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t} \mid q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid q, o_{i,<t})}, \quad (2)$$

138 denotes the probability ratio between the new and old policies for token $o_{i,t}$, and the advantage $\hat{A}_{i,t}$
 139 is estimated from group rewards as:

$$141 \quad \hat{A}_{i,t} = \frac{R_i - \mu(R)}{\sigma(R) + \epsilon}, \quad (3)$$

143 with R_i the cumulative verifiable reward of trajectory o_i , $\mu(R)$ and $\sigma(R)$ the mean and standard
 144 deviation across the sampled group, and ϵ a small constant for numerical stability.

145 By eliminating dependency on value models and instead exploiting group-normalized verifiable
 146 rewards, GRPO achieves stable and cost-efficient training.

148 Building on GRPO, Decouple Clip and Dynamic Sampling Policy Optimization (DAPO) (Yu et al.,
 149 2025) is proposed to further improve stability and exploration. DAPO integrates four key techniques:
 150 Clip-Higher, Dynamic Sampling, Token-Level Policy Gradient Loss, and Overlong Reward Shap-
 151 ing. Similar to GRPO, DAPO samples multiple responses per prompt and optimizes the following
 152 objective:

$$153 \quad \mathcal{J}_{\text{DAPO}}(\theta) = \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot \mid q)} \left[\frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \min \left(r_t^i(\theta) \hat{A}_t^i, \text{clip}(r_t^i(\theta), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}}) \hat{A}_t^i \right) \right], \quad (4)$$

158 s.t. $0 < |\{i \in \{1, \dots, G\} \mid \text{is_equivalent}(o^i, a)\}| < G$

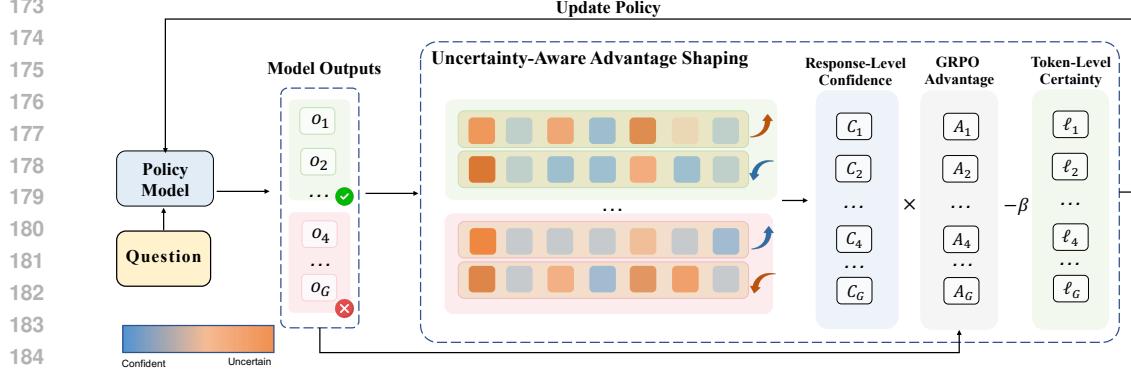
160 where ϵ_{low} and ϵ_{high} denote the lower and upper bounds of the clipping range. Compared to GRPO,
 161 DAPO explicitly decouples the clipping bounds, incorporates adaptive sampling strategies, thereby
 alleviating entropy collapse and improving the generalizability of RLVR-trained models.

162 3 METHOD

163

164 To address the coarse credit assignment problem in RLVR, we introduce **Uncertainty-aware Ad-**
 165 **vantage Shaping (UCAS)**, a method designed to replace the blunt instrument of uniform advantage
 166 with a more nuanced, two-stage mechanism. The central idea is to reshape the learning signal by
 167 considering uncertainty at two distinct granularities: the entire reasoning path (response-level) and
 168 the individual generative steps within it (token-level). This hierarchical approach first sets a *strategic*
 169 learning objective by evaluating the value of the overall trajectory, and then *locally* refines the policy
 170 update to encourage robust exploration and prevent the premature convergence that leads to entropy
 171 collapse.

172



186 **Figure 2: Overview of the UCAS Advantage Shaping Mechanism.** UCAS refines the uniform
 187 GRPO advantage through a two-stage process. **Stage 1 (Macro-level):** It applies Response-Level
 188 Modulation using the trajectory’s overall self-confidence to determine its strategic value for explo-
 189 ration vs. exploitation. **Stage 2 (Micro-level):** It introduces a Token-Level Certainty Penalty using
 190 raw logits to discourage local overconfidence. The final shaped advantage $\hat{A}_{i,t}^{\text{UCAS}}$ guides a more
 191 nuanced policy update.

192

193

194 3.1 UNCERTAINTY SIGNALS: FROM CONFIDENCE TO LOGITS

195

196

197

To perform this hierarchical shaping, UCAS requires signals that capture the model’s epistemic state at both macro and micro levels. We extract these directly from the model’s intrinsic generative process, avoiding the need for auxiliary networks.

198

199

200

Response-Level Confidence. For a high-level assessment of a full reasoning trajectory, we use the model’s **self-confidence**. As defined in Kang et al. (2025), this is the average KL-divergence between the model’s predictive distribution and a uniform distribution over the vocabulary \mathcal{V} . We denote this as $\mathcal{C}(o_i|q)$:

203

204

205

$$\mathcal{C}(o_i|q) := \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \text{KL}(U(\mathcal{V}) \parallel p_{\pi_\theta}(\cdot | q, o_{i,<t})) \quad (5)$$

206

207

A higher value of $\mathcal{C}(o_i|q)$ signifies higher overall confidence (low uncertainty) in the generated sequence, suggesting the model is following a well-trodden path.

208

209

210

211

212

213

214

215

Token-Level Certainty. While self-confidence is effective at the sequence level, it is derived from post-softmax probabilities, which can suffer from poor calibration (Liu et al., 2025a; Ma et al., 2025). This can cause the model to appear equally confident in different choices, masking subtle but important variations in uncertainty. To capture a more direct and sensitive signal at the token level, we use the model’s raw **logit** value for the chosen token $o_{i,t}$ as a proxy for certainty. Let $\ell_{i,t}$ be the logit corresponding to token $o_{i,t}$ at step t . A higher logit value indicates greater model certainty in its choice, prior to softmax normalization.

216 3.2 UCAS: TWO-STAGE ADVANTAGE SHAPING
217218 Given a group of G responses $\{o_1, \dots, o_G\}$ to a prompt q , UCAS reshapes the original GRPO
219 advantage \hat{A}_i into a fine-grained, token-specific advantage $\hat{A}_{i,t}^{\text{UCAS}}$. This process unfolds in two
220 complementary stages.
221222 **Stage 1: Response-Level Advantage Modulation.** This stage adjusts the advantage for an entire
223 response to encourage exploration of novel correct paths and suppress confident, well-trodden in-
224 correct paths. First, we compute the self-confidence $\mathcal{C}(o_i|q)$ for each response o_i in the group. To
225 assess confidence relative to other responses in the same group, we apply z-score normalization:
226

227
$$\hat{\mathcal{C}}_i = \frac{\mathcal{C}(o_i|q) - \mu_{\mathcal{C}}}{\sigma_{\mathcal{C}} + \epsilon}, \quad (6)$$

228

229 where $\mu_{\mathcal{C}}$ and $\sigma_{\mathcal{C}}$ are the mean and standard deviation of confidence scores across the group.
230231 We then compute a modulation weight $W(\hat{\mathcal{C}}_i)$ based on the sign of the original advantage \hat{A}_i , which
232 directly encodes the correctness of the answer. Theoretically, we select an exponential form to
233 act as a non-linear filter. This addresses the compressed variance often found in group-normalized
234 scores, where linear rescaling fails to sufficiently distinguish “novel” exploration from “routine”
235 exploitation.
236

237
$$W(\hat{\mathcal{C}}_i) = \begin{cases} \exp(-\alpha \cdot \hat{\mathcal{C}}_i) & \text{if } \hat{A}_i > 0 \quad (\text{Correct response}) \\ \exp(\alpha \cdot \hat{\mathcal{C}}_i) & \text{if } \hat{A}_i < 0 \quad (\text{Incorrect response}) \end{cases} \quad (7)$$

238

239 where $\alpha > 0$ is a hyperparameter controlling the shaping intensity. This formulation ensures that for
240 correct responses, lower confidence (negative $\hat{\mathcal{C}}_i$) results in a larger weight, amplifying the reward.
241 For incorrect responses, higher confidence (positive $\hat{\mathcal{C}}_i$) results in a larger weight, amplifying the
242 penalty. The resulting modulated advantage is $\hat{A}_i^{\text{mod}} = W(\hat{\mathcal{C}}_i) \cdot \hat{A}_i$.
243244 **Stage 2: Token-Level Certainty Penalty.** Response-level modulation sets a global learning ob-
245 jective for each trajectory, but this modulated advantage, \hat{A}_i^{mod} , is still a uniform signal broadcast
246 to all tokens within that sequence. This alone is insufficient to prevent the model from developing
247 localized overconfidence—a key driver of entropy collapse. The second stage therefore introduces a
248 token-specific penalty to directly address this. By penalizing high-certainty choices at each step, we
249 encourage the model to maintain a degree of epistemic humility, which preserves local exploration.
250251 We use the raw logit $\ell_{i,t}$ as our certainty measure and apply Min-Max normalization within each
252 sequence to create a standardized penalty score $\hat{\ell}_{i,t} \in [0, 1]$:
253

254
$$\hat{\ell}_{i,t} = \frac{\ell_{i,t} - \min_k(\ell_{i,k})}{\max_k(\ell_{i,k}) - \min_k(\ell_{i,k})} \quad (8)$$

255

256 A value of $\hat{\ell}_{i,t}$ close to 1 indicates high relative certainty for that token choice. This penalty acts as
257 a regularizer, complementing the directional guidance from Stage 1.
258259 **Final Advantage Shaping Formula.** By combining these two stages, UCAS creates a composite
260 advantage signal that is both globally informed and locally sensitive. The final shaped advantage for
261 each token is:
262

263
$$\hat{A}_{i,t}^{\text{UCAS}} = \underbrace{\hat{A}_i^{\text{mod}}}_{\text{Global Direction}} - \underbrace{\beta \cdot \hat{\ell}_{i,t}}_{\text{Local Penalty}} \quad (9)$$

264

265 where $\beta > 0$ is a hyperparameter controlling the penalty strength. This composite structure steers
266 the model toward novel correct solutions (via \hat{A}_i^{mod}) while ensuring it traverses reasoning paths with
267 a healthy degree of caution (via the penalty term), thereby mitigating entropy collapse and fostering
268 more robust problem-solving abilities. This final advantage term then replaces the original advantage
269

270 in the RL objective:

$$\begin{aligned}
 272 \quad \mathcal{J}_{\text{UCAS}}(\theta) &= \mathbb{E}_{(q,a) \sim \mathcal{D}, \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot|q)} \\
 273 \quad &\left[\frac{1}{\sum_{i=1}^G |o_i|} \sum_{i=1}^G \sum_{t=1}^{|o_i|} \min \left(r_t^i(\theta) \hat{A}_{i,t}^{\text{UCAS}}, \text{clip}(r_t^i(\theta), 1 - \epsilon_{\text{low}}, 1 + \epsilon_{\text{high}}) \hat{A}_{i,t}^{\text{UCAS}} \right) \right], \\
 274 \quad \text{s.t. } 0 &< |\{i \in \{1, \dots, G\} \mid \text{is_equivalent}(o^i, a)\}| < G
 \end{aligned} \tag{10}$$

275 The complete implementation process of UCAS is shown in Algorithm 1.

Algorithm 1 Uncertainty-aware Advantage Shaping (UCAS)

280 **Input:** A group of G responses $\{o_i\}_{i=1}^G$ sampled from policy π_θ , their rule-based rewards $\{R_i\}_{i=1}^G$.
281 Hyperparameters α and β .
282 1: Compute standard group-normalized advantages $\{\hat{A}_i\}_{i=1}^G$ from $\{R_i\}_{i=1}^G$.
283 2: Compute response-level self-confidence $\{\mathcal{C}(o_i|q)\}_{i=1}^G$ for all responses.
284 3: Normalize confidences across the group to get $\{\hat{\mathcal{C}}_i\}_{i=1}^G$.
285 4: **for** $i = 1$ to G **do** ▷ Stage 1: Response-Level Advantage Modulation
286 5: Compute modulation weight $W(\hat{\mathcal{C}}_i)$ based on \hat{A}_i using Eq. 7.
287 6: Modulate the advantage: $\hat{A}_i^{\text{mod}} \leftarrow W(\hat{\mathcal{C}}_i) \cdot \hat{A}_i$.
288 7: ▷ Stage 2: Token-Level Certainty Penalty
289 8: Get the sequence of logits $\{\ell_{i,t}\}_{t=1}^{|o_i|}$ for the generated tokens in o_i .
290 9: Normalize logits within the sequence to get $\{\hat{\ell}_{i,t}\}_{t=1}^{|o_i|}$.
291 10: **for** $t = 1$ to $|o_i|$ **do**
292 11: Compute the final UCAS advantage:
293 12: $\hat{A}_{i,t}^{\text{UCAS}} \leftarrow \hat{A}_i^{\text{mod}} - \beta \cdot \hat{\ell}_{i,t}$.
294 13: **end for**
295 14: **end for**
296 **Output:** The set of token-level UCAS advantages $\{\hat{A}_{i,t}^{\text{UCAS}}\}$.

300 **4 EXPERIMENTS**

301 4.1 EXPERIMENTAL SETUP

302 **Training Data and Benchmarks.** During the training phase, we utilize the widely-used MATH
303 dataset as our training set. To maintain consistency with prior research, we only use the more
304 challenging subset of this dataset for training, specifically problems from levels 3 to 5. To com-
305 prehensively evaluate the reasoning capabilities of the model trained with our method, we select
306 five widely recognized benchmarks in the mathematical reasoning domain for testing: AIME24 (LI
307 et al., 2024), MATH-500 (Hendrycks et al., 2021), AMC (LI et al., 2024), Minerva (Lewkowycz
308 et al., 2022), and OlympiadBench (Huang et al., 2024), which collectively contain 1,560 problems.

309 **Models and Baselines.** We employ two variants of the Qwen2.5-Math (Yang et al., 2024) series as
310 our foundation models: Qwen2.5-Math-1.5B and Qwen2.5-Math-7B. First, to quantify the per-
311 formance improvement introduced by our method, we select the widely used GRPO and DAPO algo-
312 rithms as comparison baselines. Furthermore, to benchmark against existing reinforcement learning
313 techniques, we also select the following representative methods for comparison:

- 314 • **Simple-RL-Zoo (Zeng et al., 2025):** Based on Qwen2.5-Math-7B, trained on the math-
315 level3-5 dataset using the standard GRPO algorithm with rule-based rewards.
- 316 • **PRIME-Zero (Cui et al., 2025a):** An online PRM update approach that leverages implicit
317 process rewards from rollouts and outcome labels without requiring explicit annotations.
- 318 • **OpenReasonerZero (Hu et al., 2025):** A zero-RL baseline on Qwen2.5-7B employing the
319 standard PPO algorithm for policy optimization.

- 324 • **Oat-Zero (Liu et al., 2025b)**: Built on Qwen2.5-Math-7B, trained with rule-based rewards
325 using a modified Dr.GRPO algorithm that removes variance terms and applies token-level
326 normalization in the policy loss.
- 327 • **GRPO with Entropy Adv. (Cheng et al., 2025)**: Extends RLVR training by incorpo-
328 rating a clipped, gradient-detached entropy term into the advantage function to encourage
329 exploration.
- 330 • **KTAE (Sun et al., 2025b)**: A token-level advantage estimation method trained with
331 DAPO, quantifying key-token contributions via statistical association tests and combining
332 them with rollout-level advantages.

334 These baselines cover applications of fundamental RL algorithms, process-reward-based methods,
335 and algorithms improved for specific tasks like mathematical reasoning, aiming to evaluate the ef-
336 fectiveness and novelty of our method from multiple perspectives.

337 **Implementation Details.** We adopt the VERL framework (Sheng et al., 2024) and train our model
338 using the optimization objective defined in Eq. 10. During training, the model’s maximum context
339 length is set to 4096, with a maximum prompt length of 1024 and a maximum response length of
340 3072. The learning rate is fixed at 1×10^{-6} , and the training batch size is set to 512. For each
341 prompt, we sample 16 rollouts with a sampling temperature of 1.0. For the DAPO baseline, we use
342 clipping thresholds of $\epsilon_{\text{low}} = 0.2$ and $\epsilon_{\text{high}} = 0.28$. The KL penalty loss and entropy regularization
343 loss are omitted from the objective function. The hyperparameters for our UCAS method, α and
344 β , are set to 0.25 and 0.01, respectively. All experiments are conducted on 2 compute nodes, each
345 equipped with 8 NVIDIA A800 80GB GPUs.

347 4.2 MAIN RESULTS

348 The greedy pass@1 performance comparison between 1.5B and 7B models across five mathematical
349 reasoning benchmarks is presented in Table 1. We can clearly find that **the UCAS model achieved**
350 **the highest performance across all five math reasoning benchmarks on both the 1.5B and 7B**
351 **parameter scales**. Compared with the DAPO baseline, UCAS improves the average accuracy from
352 41.2 to 47.3 (+6.1) on Qwen2.5-Math-1.5B and from 50.5 to 56.7 (+6.2) on Qwen2.5-Math-7B. Be-
353 yond DAPO, UCAS also surpasses strong baselines such as KTAE and Oat-Zero, with pronounced
354 gains on challenging benchmarks including AIME24, AMC, and OlympiadBench. These results
355 highlight the robustness and scalability of uncertainty-aware advantage shaping, demonstrating con-
356 sistent benefits across model sizes and diverse reasoning tasks.

358 4.3 ANALYSIS

359 4.3.1 ABLATION STUDY

360 The ablation comparison between response-level and token-level uncertainty modeling is presented
361 in Table 2. We can clearly observe that **both response-level and token-level uncertainty bring**
362 **consistent gains over the DAPO baseline**. Compared with the model trained with DAPO, incor-
363 porating response-level confidence increases the average score on Qwen2.5-Math-1.5B from 41.2
364 to 44.7 (+3.5%), while token-level uncertainty further raises it to 45.1 (+3.9%). A similar trend
365 holds on the 7B model, where both variants surpass the DAPO baseline. Their integration in UCAS
366 achieves the best performance, confirming that both signals are individually useful and jointly nec-
367 essary.

368 4.3.2 TRAINING DYNAMICS

369 The training process highlights several key performance trends, as shown in Figure 3. Compared to
370 vanilla GRPO, UCAS demonstrates a consistent increase in the inference reward on the MATH500
371 benchmark. Regarding the average response length, the inclusion of UCAS enables the model to
372 generate longer reasoning chains, reflecting more comprehensive problem-solving (Guo et al., 2025;
373 Cheng et al., 2025), while simultaneously improving accuracy. For generation entropy, UCAS shows
374 an early decline but later recovers and stabilizes at a higher level, effectively avoiding the entropy
375 collapse reported in prior work (Cui et al., 2025b). Notably, the model’s reward continues to rise

Models	AIME24	MATH-500	AMC	Minerva	Olympiad	Avg
<i>Qwen2.5-Math-1.5B</i>						
Base Model	7.3	61.8	43.4	15.1	28.4	31.2
GRPO	15.6	76.0	51.8	22.1	36.3	40.4
DAPO	16.7	77.6	47.0	25.7	39.0	41.2
Oat-Zero(Liu et al., 2025b)	20.0	74.4	50.6	23.9	37.0	41.2
KTAE(Sun et al., 2025b)	20.0	77.6	50.6	29.0	40.0	43.4
SEED-GRPO(Chen et al., 2025)	23.3	75.4	50.6	26.8	41.3	43.5
UCAS	23.3	80.6	59.0	31.6	42.1	47.3
<i>Qwen2.5-Math-7B</i>						
Base Model	11.0	69.0	45.8	21.3	28.4	35.1
GRPO	30.0	81.0	57.8	32.7	43.2	48.9
DAPO	30.5	81.8	60.2	34.5	45.3	50.5
PRIME-Zero (Cui et al., 2025a)	23.3	82.2	57.8	36.0	39.9	47.8
OpenReasonerZero (Hu et al., 2025)	17.9	78.4	45.8	27.9	45.0	43.0
Oat-Zero(Liu et al., 2025b)	32.1	79.8	61.4	30.5	41.8	49.1
Simple RL-Zero(Zeng et al., 2025)	26.7	78.6	59.0	33.8	43.4	48.3
GRPO with Entropy Adv. (Cheng et al., 2025) [†]	33.7	83.1	69.8	-	-	-
KTAE(Sun et al., 2025b)	36.7	83.2	63.9	35.3	43.7	52.6
SEED-GRPO(Chen et al., 2025)	43.3	82.2	64.7	35.0	45.2	54.7
UCAS	43.3	85.6	68.7	37.6	48.0	56.7

Table 1: The greedy pass@1 performance of 1.5B and 7B models across five math reasoning benchmarks. [†]: results from Cheng et al. (2025). Our method UCAS consistently surpasses all baselines in both parameter scales.

Models	AIME24	MATH-500	AMC	Minerva	Olympiad	Avg
<i>Qwen2.5-Math-1.5B</i>						
Base Model	7.3	61.8	43.4	15.1	28.4	31.2
w/ DAPO	16.7	77.6	47.0	25.7	39.0	41.2
w/ DAPO + Response-Level Confidence	23.3	79.6	51.8	27.6	41.0	44.7
w/ DAPO + Token-Level Certainty	20.0	80.2	55.4	29.7	40.1	45.1
w/ DAPO + UCAS (Ours)	23.3	80.6	59.0	31.6	42.1	47.3
<i>Qwen2.5-Math-7B</i>						
Base Model	11.0	69.0	45.8	21.3	28.4	35.2
w/ DAPO	30.5	81.8	60.2	34.5	45.3	50.5
w/ DAPO + Response-Level Confidence	40.0	85.0	63.9	36.7	47.4	54.6
w/ DAPO + Token-Level Certainty	36.7	84.6	65.0	29.7	47.7	52.7
w/ DAPO + UCAS (Ours)	43.3	85.6	68.7	37.6	48.0	56.7

Table 2: Ablation study of uncertainty modeling. Both sentence-level and token-level uncertainty bring consistent gains over the DAPO baseline.

even as the entropy increases, which indicates a stable and effective training dynamic where exploration and optimization are well-balanced.

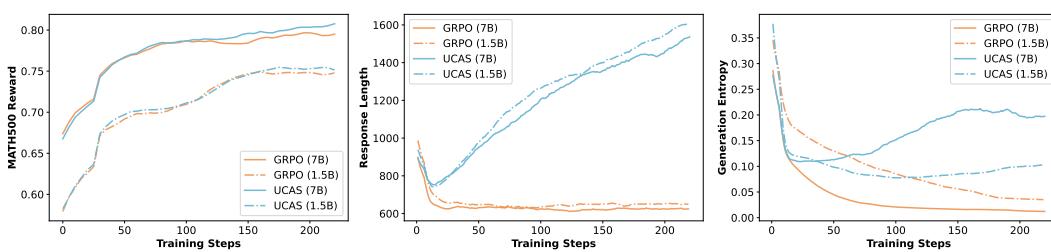


Figure 3: Training dynamics of UCAS compared with GRPO across both 7B and 1.5B models. **Left:** Reward; **Middle:** Response Length; **Right:** Generation Entropy.

432 4.3.3 CROSS-DOMAIN GENERALIZATION

434 To assess the generality of UCAS beyond mathematical reasoning, we conduct additional experiments on three diverse benchmarks: **LeetCode** (Guo et al., 2024) (code generation), **Live-435 CodeBench** (Jain et al., 2024) (competitive programming), and **MMLU** (Hendrycks et al., 2020) 436 (general task reasoning). As shown in Table 3, despite being trained solely on mathematical reasoning 437 data, UCAS consistently outperforms the strong DAPO baseline across all non-math tasks. 438

439 440 441 Method	442 443 444 LeetCode (Pass@1)	442 443 444 LiveCodeBench (Pass@1)	442 443 444 MMLU (Acc)	442 443 444 Avg
442 443 444 Base Model	442 443 444 11.7	442 443 444 5.7	442 443 444 65.7	442 443 444 27.7
442 443 444 DAPO	442 443 444 18.3	442 443 444 9.2	442 443 444 67.3	442 443 444 31.6
442 443 444 UCAS (Ours)	442 443 444 23.6 (+5.3)	442 443 444 14.8 (+5.6)	442 443 444 70.8 (+3.5)	442 443 444 36.4 (+4.8)

445 Table 3: Generalizing UCAS from math-only training to evaluations on non-math tasks.

446 This strong transferability suggests that our uncertainty-aware exploration mechanism is a broadly 447 applicable principle. By unlocking exploration for high-uncertainty paths, UCAS improves performance 448 not just in calculation but also in the multi-step logical planning required for programming 449 and general reasoning, demonstrating gains beyond the mathematical domain. 450

451 4.3.4 PASS@K EVALUATION

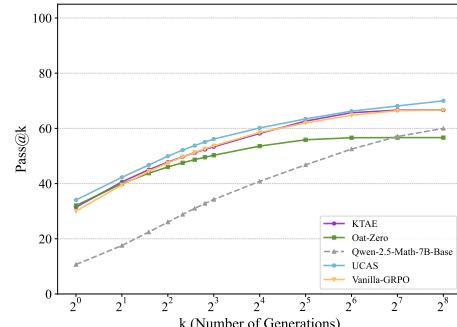
452 Prior studies (Wang et al., 2022; Wu et al., 2024) 453 have shown that with a limited number of rollouts, 454 models often struggle to solve certain tasks. In 455 contrast, when the rollout budget is sufficiently large, 456 the probability of sampling effective solutions 457 increases considerably. This observation suggests that 458 pass@k accuracy with a large k provides a more 459 reliable estimate of a model’s potential performance 460 (Yue et al., 2025). Under this evaluation protocol, a 461 problem is considered solved if any of the k sampled 462 reasoning trajectories yield the correct answer. 463

464 Figure 4 reports pass@k results on the AIME24 465 benchmark. The results indicate that UCAS achieves 466 more consistent improvements as k grows. In 467 contrast, Vanilla-GRPO and its enhanced variants show 468 slower growth, consistent with findings from 469 Yue et al. (2025). The stronger performance of UCAS 470 under the pass@k metric highlights its effectiveness, 471 which can be attributed to differences in exploration 472 strategies. Unlike Vanilla-GRPO, which often suffers 473 from exploration stagnation, where the model 474 repeatedly samples low-diversity rollouts, UCAS leverages 475 uncertainty-aware advantage shaping to sustain diverse 476 exploration and escape local optima.

477 5 RELATED WORK

478 5.1 RL FOR LLM REASONING

479 Recent advances in reinforcement learning have transformed the training of large language models 480 for reasoning tasks. Process reward models (PRMs) (Lightman et al., 2023) have emerged as a key 481 innovation, providing step-level supervision that improves both efficiency and accuracy compared to 482 outcome-only rewards. Approaches such as PRIME (Cui et al., 2025a) eliminate costly human annotation 483 by deriving implicit process feedback, while OmegaPRM (Luo et al., 2024) leverages Monte Carlo Tree 484 Search (MCTS) to automatically identify reasoning errors. Alongside this, DeepSeek-R1 485 (Guo et al., 2025) demonstrates that sophisticated reasoning can emerge purely from RL without supervised 486 fine-tuning, enabled by GRPO, which replaces value functions with group-based baselines.



487 Figure 4: Comparison of pass@k results on 488 the AIME24 Benchmark.

486 These advances redefine alignment and reasoning in LLMs, positioning reinforcement learning with
 487 verifiable or process-level rewards as a scalable and principled alternative to preference-model-based
 488 RLHF.

490 **5.2 REINFORCEMENT LEARNING FROM VERIFIABLE REWARDS**

492 RLVR has emerged as a scalable alternative to preference-based alignment by converting open-
 493 ended outputs into checkable signals such as mathematical correctness or unit-test pass rates (Guo
 494 et al., 2025; Yue et al., 2025). While early implementations demonstrated strong gains in pass@1
 495 accuracy, subsequent studies revealed a consistent *policy entropy collapse*: models rapidly concen-
 496 trate probability mass on a narrow set of high-reward trajectories, diminishing output diversity and
 497 limiting exploration (Cui et al., 2025b). Empirical analyses show that RLVR-trained models often
 498 underperform base models on pass@ k (Shao et al., 2024; Yue et al., 2025), highlighting a precision-
 499 diversity trade-off (Wu et al., 2025; Dong et al., 2025).

500 Algorithmic responses to entropy collapse vary. Standard entropy or KL penalties provide partial
 501 remedies, though their effectiveness often depends heavily on the divergence form (Li et al., 2025).
 502 **More recent uncertainty-aware approaches have sought to refine the learning signal, though with**
 503 **differing philosophies.** For instance, SEED-GRPO (Chen et al., 2025) leverages semantic entropy to
 504 downscale updates for uncertain queries, adopting a conservative risk-mitigation strategy. In stark
 505 contrast, UCAS adopts an exploratory philosophy: we explicitly amplify rewards for correct-but-
 506 uncertain trajectories to incentivize venturing into novel reasoning domains, rather than inhibiting
 507 learning from uncertainty.

508 Similarly, while entropy-based shaping methods (Cheng et al., 2025) introduce indiscriminate en-
 509 tropy bonuses to encourage diversity, UCAS implements a *conditional*, two-stage mechanism. By
 510 combining response-level confidence with token-level raw logits which we find to be a more sen-
 511 sitive proxy for local overconfidence than post-softmax entropy, UCAS distinguishes between pro-
 512 ductive exploration and blind guessing. Unlike pure entropy-based frameworks, UCAS introduces
 513 correctness-contingent modulation, amplifying penalties for confident errors while guiding explo-
 514 ration through uncertainty, offering a more fine-grained solution to the entropy collapse problem
 515 than global regularization or token-level covariance control (Cui et al., 2025b).

516 **6 CONCLUSION**

518 In this work, we introduced UnCertainty-aware Advantage Shaping (UCAS), a fine-grained advan-
 519 tage estimation framework that leverages internal confidence signals without requiring additional re-
 520 ward models. By jointly modeling uncertainty at both the response and token levels, UCAS reshapes
 521 advantages to highlight critical uncertain reasoning steps and suppress overconfident yet erroneous
 522 segments. Experimental results on major mathematical reasoning benchmarks show that UCAS
 523 achieves substantial performance improvements over GRPO and its enhanced variants. Analysis of
 524 the training dynamics further reveals that, as training progresses, UCAS demonstrates steadily in-
 525 creasing rewards, longer reasoning chains, and an entropy trajectory that first declines and then rises,
 526 reflecting stronger exploratory capability. These findings indicate that uncertainty-aware advantage
 527 shaping offers an effective pathway toward more robust reinforcement learning for large language
 528 models.

540 REFERENCES
541

542 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. Step-level value preference optimization
543 for mathematical reasoning. *arXiv preprint arXiv:2406.10858*, 2024.

544 Minghan Chen, Guikun Chen, Wenguan Wang, and Yi Yang. Seed-grpo: Semantic entropy enhanced
545 grpo for uncertainty-aware policy optimization. *arXiv preprint arXiv:2505.12346*, 2025.

546 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and
547 Furu Wei. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*,
548 2025.

549 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
550 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint
551 arXiv:2502.01456*, 2025a.

552 Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen
553 Fan, Huayu Chen, Weize Chen, et al. The entropy mechanism of reinforcement learning for
554 reasoning language models. *arXiv preprint arXiv:2505.22617*, 2025b.

555 Yihong Dong, Xue Jiang, Yongding Tao, Huanyu Liu, Kechi Zhang, Lili Mou, Rongyu Cao, Ying-
556 wei Ma, Jue Chen, Binhua Li, et al. RL-plus: Countering capability boundary collapse of llms in
557 reinforcement learning with hybrid-policy optimization. *arXiv preprint arXiv:2508.00222*, 2025.

558 Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao
559 Bi, Yu Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming—
560 the rise of code intelligence. *arXiv preprint arXiv:2401.14196*, 2024.

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

564 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
565 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint
566 arXiv:2009.03300*, 2020.

567 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
568 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv
569 preprint arXiv:2103.03874*, 2021.

570 Jingcheng Hu, Yinmin Zhang, Qi Han, Dixin Jiang, Xiangyu Zhang, and Heung-Yeung Shum.
571 Open-reasoner-zero: An open source approach to scaling up reinforcement learning on the base
572 model. *arXiv preprint arXiv:2503.24290*, 2025.

573 Zhen Huang, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, Lyu-
574 manshan Ye, Ethan Chern, Yixin Ye, et al. Olympicarena: Benchmarking multi-discipline cog-
575 nitive reasoning for superintelligent ai. *Advances in Neural Information Processing Systems*, 37:
576 19209–19253, 2024.

577 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
578 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
579 evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

580 Zhewei Kang, Xuandong Zhao, and Dawn Song. Scalable best-of-n selection for large language
581 models via self-certainty. *arXiv preprint arXiv:2502.18581*, 2025.

582 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-
583 masesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative
584 reasoning problems with language models. *Advances in neural information processing systems*,
585 35:3843–3857, 2022.

594 Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang,
 595 Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann
 596 Fleureau, Guillaume Lample, and Stanislas Polu. Numinamath. [<https://huggingface.co/AI-MO/NuminaMath-CoT>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.

597

598

599 Long Li, Jiaran Hao, Jason Klein Liu, Zhijian Zhou, Xiaoyu Tan, Wei Chu, Zhe Wang, Shirui Pan,
 600 Chao Qu, and Yuan Qi. The choice of divergence: A neglected key to mitigating diversity collapse
 601 in reinforcement learning with verifiable reward. *arXiv preprint arXiv:2509.07430*, 2025.

602

603 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 604 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth*
 605 *International Conference on Learning Representations*, 2023.

606

607 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
 608 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *The Twelfth*
 609 *International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=v8L0pN6EOi>.

610

611 Xiaou Liu, Tiejin Chen, Longchao Da, Chacha Chen, Zhen Lin, and Hua Wei. Uncertainty quan-
 612 tification and confidence calibration in large language models: A survey. In *Proceedings of the*
 613 *31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining* V. 2, pp. 6107–6117,
 614 2025a.

615

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

619

620 Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li,
 621 Lei Shu, Yun Zhu, Lei Meng, et al. Improve mathematical reasoning in language models by
 622 automated process supervision. *arXiv preprint arXiv:2406.06592*, 2024.

623

624 Huan Ma, Jingdong Chen, Guangyu Wang, and Changqing Zhang. Estimating llm uncertainty with
 625 logits. *arXiv e-prints*, pp. arXiv–2502, 2025.

626

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

630

631 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
 632 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. *arXiv preprint*
 633 *arXiv: 2409.19256*, 2024.

634

635 Yi Su, Dian Yu, Linfeng Song, Juntao Li, Haitao Mi, Zhaopeng Tu, Min Zhang, and Dong Yu.
 636 Crossing the reward bridge: Expanding rl with verifiable rewards across diverse domains. *arXiv*
 637 *preprint arXiv:2503.23829*, 2025.

638

639 Wei Sun, Qianlong Du, Fuwei Cui, and Jiajun Zhang. An efficient and precise training data construc-
 640 tion framework for process-supervised reward model in mathematical reasoning. *arXiv preprint*
 641 *arXiv:2503.02382*, 2025a.

642

643 Wei Sun, Wen Yang, Pu Jian, Qianlong Du, Fuwei Cui, Shuo Ren, and Jiajun Zhang. Ktae: A model-
 644 free algorithm to key-tokens advantage estimation in mathematical reasoning. *arXiv preprint*
 645 *arXiv:2505.16826*, 2025b.

646

647 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang
 648 Sui. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv*
 649 *preprint arXiv:2312.08935*, 2023.

650

651 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,
 652 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive
 653 effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025a.

648 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 649 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 650 *arXiv preprint arXiv:2203.11171*, 2022.

651

652 Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Liyuan Liu, Baolin Peng, Hao Cheng, Xuehai
 653 He, Kuan Wang, Jianfeng Gao, et al. Reinforcement learning for reasoning in large language
 654 models with one training example. *arXiv preprint arXiv:2504.20571*, 2025b.

655 Fang Wu, Weihao Xuan, Ximing Lu, Zaid Harchaoui, and Yejin Choi. The invisible leash: Why rlv
 656 may not escape its origin. *arXiv preprint arXiv:2507.14843*, 2025.

657

658 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws:
 659 An empirical analysis of compute-optimal inference for problem-solving with language models.
 660 *arXiv preprint arXiv:2408.00724*, 2024.

661 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-
 662 hong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical
 663 expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

664

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

668 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Shiji Song, and Gao Huang. Does re-
 669 enforcement learning really incentivize reasoning capacity in llms beyond the base model? *arXiv
 670 preprint arXiv:2504.13837*, 2025.

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

674

675 Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh
 676 Agarwal. Generative verifiers: Reward modeling as next-token prediction. *arXiv preprint
 677 arXiv:2408.15240*, 2024.

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A LIMITATION AND FUTURE WORK**
703704 While our work demonstrates the effectiveness of UCAS in the domain of mathematical reasoning,
705 we acknowledge several limitations that present valuable opportunities for future research.
706707 First, our experiments are exclusively focused on mathematical tasks, which benefit from clear,
708 binary verifiable rewards (i.e., the answer is either correct or incorrect). The direct applicability of
709 UCAS to domains with more nuanced, subjective, or dense reward signals, such as creative writing,
710 summarization, or open-domain dialogue remains an open question. Adapting the uncertainty-aware
711 shaping mechanism to these softer reward landscapes would be a crucial next step.
712713 Second, our method relies on self-confidence and raw logits as proxies for model uncertainty. While
714 these internal signals are computationally efficient and effective, future work could explore alter-
715 native or complementary uncertainty metrics. Techniques such as Monte Carlo dropout, model
716 ensembles, or semantic entropy could potentially capture different facets of model uncertainty and
717 lead to even more refined and robust advantage shaping. Investigating these areas will be essential
718 for understanding the broader generalizability of our approach.
719720 **B IMPACT OF DIFFERENT HYPERPARAMETER WEIGHTS**
721722 To evaluate the robustness of UCAS and ensure the reliability of our findings, we conduct a sensi-
723 tivity analysis on the two key hyperparameters: the response-level modulation coefficient α and the
724 token-level penalty coefficient β . For each experiment, we vary one hyperparameter while fixing the
725 other at its empirically optimal value, and report the post-training results of the Qwen-Math-1.5B
726 model on the Math-500 benchmark.727 Table 4: Sensitivity analysis on Math-500 when varying the token-level penalty β (upper block) and
728 the response-level modulation coefficient α (lower block).
729

Response-Level α	Token-Level β	Math-500
<i>Varying Token-Level Penalty (β), Fixed $\alpha = 0.2$</i>		
0.2	0.005	79.6
0.2	0.01	80.4
0.2	0.05	78.8
0.2	0.1	78.4
<i>Varying Response Modulation (α), Fixed $\beta = 0.01$</i>		
0.1	0.01	80.0
0.2	0.01	80.4
0.4	0.01	78.8

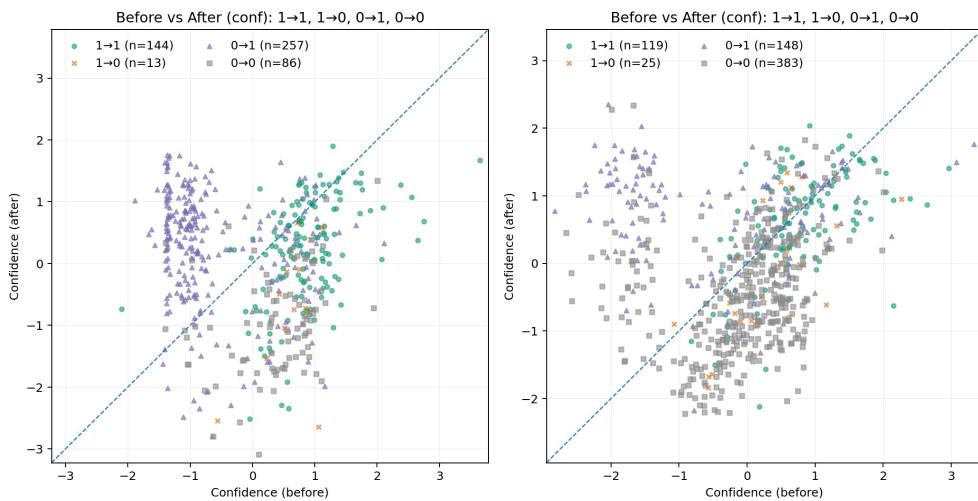
743 **Impact of Token-Level Penalty (β).** The upper block of Table 4 elucidates the critical trade-off
744 controlled by the token-level penalty β . We observe that a moderate penalty ($\beta = 0.01$) yields opti-
745 mal performance (80.4), suggesting it successfully balances the exploration-exploitation dynamic.
746 Notably, when β is too low (0.005), performance dips to 79.6, indicating that the penalty is in-
747 sufficient to counteract the model’s natural tendency toward overconfidence and entropy collapse.
748 Conversely, as β increases beyond 0.05, performance degrades further to 78.4. This suggests that
749 excessive penalization over-regularizes the policy, forcing the model to artificially flatten its distribu-
750 tion even for necessary, high-certainty reasoning steps, thereby hindering the generation of coherent
751 solution chains.
752753 **Impact of Response-Level Modulation (α).** The lower block of Table 4 examines the sensitiv-
754 ity to the response-level modulation coefficient α . The method exhibits stability within the range
755 $\alpha \in [0.1, 0.2]$, where the modulation appropriately emphasizes uncertain-but-correct trajectories.
756 However, increasing α to 0.4 results in a noticeable performance drop to 78.8. This deterioration

756 implies a "signal dominance" issue: when the modulation is overly aggressive, the confidence-based
 757 scaling factor begins to overshadow the fundamental correctness signal (the original advantage).
 758 This introduces high variance into the reward landscape, causing the policy optimization to drift
 759 away from the primary objective of mathematical accuracy in favor of gaming the uncertainty met-
 760 ric.

762 C FURTHER ANALYSIS

764 To further analyze the effect of UCAS training, we compute the response-level confidence scores
 765 of model outputs according to Eq. 5, measured before and after training on Qwen2.5-Math-1.5B
 766 across MATH and Olympiad datasets. We focus on the MATH and Olympiad datasets because they
 767 contain more samples and a larger number of responses whose correctness changes after training,
 768 which makes them well suited for detailed analysis. For comparability, the confidence values are
 769 normalized by subtracting the mean and dividing by the standard deviation.

770 Based on the correctness of the responses before and after training, the samples are categorized into
 771 three groups: (i) consistently correct ($1 \rightarrow 1$), (ii) correct before but incorrect after ($1 \rightarrow 0$), (iii) incor-
 772 rect before but correct after ($0 \rightarrow 1$), and (iv) incorrect before and incorrect after ($0 \rightarrow 0$). Figure 5
 773 illustrates the distribution of these categories, where each point represents model's response to a
 774 given problem.



792 Figure 5: Confidence dynamics before and after UCAS training on the MATH and Olympiad
 793 datasets.

794 From Figure 5, we observe that for many problems correctly solved only after UCAS training ($0 \rightarrow 1$),
 795 the model's confidence notably increases. In contrast, for problems that remain unsolved before
 796 and after training ($0 \rightarrow 0$), the model tends to reduce its confidence, suggesting a more calibrated
 797 estimation of its own uncertainty.

801 D LLMs USAGE STATEMENT

803 We employed a Large Language Model (LLM) to assist exclusively in the editorial stage of
 804 manuscript preparation. Its role was limited to refining phrasing, correcting grammar, and enhanc-
 805 ing clarity and readability across different sections. The LLM had no involvement in formulating
 806 research ideas, designing experiments, or conducting analyses. All scientific contributions and find-
 807 ings are entirely the work of the authors. The authors have ensured that the use of the LLM complies
 808 with ethical standards, avoiding plagiarism and scientific misconduct.