

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 HINT: HELPING INEFFECTIVE ROLLOUTS NAVIGATE TOWARDS EFFECTIVENESS

Anonymous authors

Paper under double-blind review

ABSTRACT

Reinforcement Learning (RL) has become a key driver for enhancing the long chain-of-thought (CoT) reasoning capabilities of Large Language Models (LLMs). However, prevalent methods like GRPO often fail when task difficulty exceeds the model’s capacity, leading to reward sparsity and inefficient training. While prior work attempts to mitigate this using off-policy data, such as mixing RL with Supervised Fine-Tuning (SFT) or using hints, they often misguide policy updates. In this work, we identify a core issue underlying these failures, which we term low training affinity. This condition arises from a large distributional mismatch between external guidance and the model’s policy. To diagnose this, we introduce *Affinity*, the first quantitative metric for monitoring exploration efficiency and training stability. To improve *Affinity*, we propose HINT: **H**elping **I**nffective rollouts **N**avigate **T**owards effectiveness, an adaptive hinting framework. Instead of providing direct answers, HINT supplies heuristic hints that **guide the model to discover solutions on its own**, preserving its autonomous reasoning capabilities. Extensive experiments on mathematical reasoning tasks show that HINT consistently outperforms existing methods, achieving state-of-the-art results with models of various scales, while also demonstrating significantly more stable learning and greater data efficiency. Code is available on Github¹.

1 INTRODUCTION

RL methods, particularly GRPO (Shao et al., 2024), play a pivotal role in advancing long CoT reasoning (Wei et al., 2022). By avoiding the instability and overhead of training a separate value model, GRPO leverages group-based reward aggregation to deliver stable and efficient learning signals. Such RL approaches (Ahmadian et al., 2024; Shao et al., 2024; Hu, 2025; Yu et al., 2025) have become a key driver of progress in reasoning ability, enabling models to explore solution paths on verifiable problems. Building on these advances, recent reasoning models such as DeepSeek-R1 (Guo et al., 2025), OpenAI-o1 (Jaech et al., 2024), and Kimi-1.5 (Team et al., 2025) have achieved remarkable performance on complex tasks like mathematical problem solving (Shao et al., 2024) and programming (Jiang et al., 2024).

A critical challenge for GRPO, despite its strong empirical performance, is its tendency to generate sample groups consisting entirely of incorrect answers on tasks whose difficulty exceeds the policy model’s evolving capacity (Zhao et al., 2025; Yue et al., 2025). In such cases, the learning process suffers from reward sparsity, where the feedback becomes uniform and uninformative (Yu et al., 2025), ultimately reducing training efficiency and wasting valuable data.

Leveraging external, off-policy data is a key method for addressing this issue. This method has been implemented in prior work through two main lines of remedies. (I) **Mixed-policy** (Yan et al., 2025; Zhang et al., 2025a; Fu et al., 2025b): Mixed-policy involves interleaving RL with SFT in a hybrid scheme to stabilize training by leveraging off-policy data. (II) **Using hints** (Li et al., 2025; Liu et al., 2025b; Zhang et al., 2025b): To mitigate reward sparsity and ensure continuous training updates, another common approach is to leverage prompts derived from the ground truth during the rollout phase, guiding the model’s exploration along correct trajectories.

¹<https://anonymous.4open.science/r/HINT-9DD9/>




Figure 1: Comparison of Hint Mechanisms and Their Impact on Learning. The answer-level hint provides an explicit partial solution. The model can achieve a reward by simply completing this pre-defined path, which encourages learning a superficial shortcut rather than genuine reasoning. In contrast, our heuristic hint offers a high-level conceptual prompt, **compelling the model to develop its own solution path independently**.

Despite their potential benefits, both of these approaches introduce a significant drawback rooted in a **substantial distributional mismatch**. In mixed-policy training, this mismatch arises between the off-policy SFT data and the on-policy updates, which lead to conflicting gradients and training instability (Yan et al., 2025). Similarly, answer-level hints create a severe mismatch between the distribution of the ground truth and the distribution of the current policy. This results in a deceptive learning signal that, while inflating training rewards, ultimately misguides policy updates toward non-generalizable or spurious solution paths (See Figure 2).

Fundamentally, the aforementioned drawbacks stem from a **lack of what we term training affinity**. This core issue that arises from an **over-reliance** on off-policy sources, such as SFT data or answer-level hints, which inevitably creates a significant distributional mismatch with the model’s current policy (Fu et al., 2025a). This mismatch, in turn, leads to excessively high variance in the importance sampling ratios, destabilizing the entire training process. This instability is such a core challenge that prominent algorithms like PPO introduce mechanisms such as clipping to manage it (Schulman et al., 2017), the behavior of which itself provides a signal of training dynamics. To leverage this insight and create a quantitative diagnostic, we define *Affinity* metric in terms of training stability, considering both the frequency of clipping and the variance of the importance sampling ratios.

To leverage off-policy data for enhancing model capability while preserving training affinity, the guiding principle must be to **help the model articulate the solution on its own, rather than being directly told the answer**. To this end, we propose HINT: Helping Ineffective rollouts Navigate Towards effectiveness, an adaptive hinting framework. As illustrated in Figure 1, HINT implements this principle by providing heuristic hints instead of partial ground-truth answers. These hints serve as high-level guidance, helping the model navigate challenging problems without disclosing solutions. This dynamic is akin to the Socratic method in teaching, where guiding a student with thoughtful prompts, rather than supplying answers, is crucial for developing robust and generalizable reasoning skills.

Our contributions can be summarized as follows:

- We introduce the first formal definition of low training affinity, a key failure mode in RL methods that incorporate off-policy data. Building on this formalization, we propose *Affinity*, a quantitative metric that enables the continuous monitoring of these critical training dynamics.
- To effectively enhance the model’s reasoning capabilities while preserving high *Affinity*, we propose HINT, a framework that adaptively providing heuristic hints. HINT guides the model towards successful trajectories without compromising its autonomous exploration and reasoning capabilities.
- Extensive experiments validate our approach. HINT consistently outperforms methods based on mixed-policy and answer-level hints, achieving state-of-the-art results with models of various scales across multiple datasets. Furthermore, our method demonstrates robustness and superior generalization.

108

109

2 METHODS

111

Following common practice in recent work, our experiments build on the GRPO algorithm (Guo et al., 2025) while omitting the KL penalty term, as also done in (Yu et al., 2025; Yan et al., 2025). Mathematically, GRPO optimizes the model's behavior through the following objective function:

115
116
117
118
119

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{(q, a) \sim D, \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}}(\cdot | q) \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left(\min \left(\frac{\pi_{\theta}(o_{i,t} \mid o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid o_{i,<t})} A_{i,t}, \text{clip} \left(\frac{\pi_{\theta}(o_{i,t} \mid o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} \mid o_{i,<t})}, 1 \pm \epsilon \right) A_{i,t} \right) \right) \right], \quad (1)$$

120
121
122

for each prompt, GRPO draws a group of G rollouts and computes a group-normalized advantage for every token. Let $\{R_i\}_{i=1}^G$ denote the sequence-level rewards assigned to these rollouts. The token-level advantages $A_{i,t}$ are computed by normalizing each trajectory's reward within the group:

123

$$A_{i,t} = \frac{R_i - \text{mean}(\{R_j\}_{j=1}^G)}{\text{std}(\{R_j\}_{j=1}^G) + \varepsilon}. \quad (2)$$

126
127
128
129
130

When all rollouts in a group are assigned identical rewards, $R_i - \text{mean}(\{R_j\}_{j=1}^G)$ becomes zero for every i , causing every advantage $A_{i,t}$ to collapse to zero. Such prompts therefore provide no learning signal during training. Conversely, prompts that produce non-identical rewards across the group yield non-zero advantages and therefore generate meaningful gradients.

131

2.2 THE ILLUSION OF HIGH REWARD

133
134
135
136
137
138
139
140
141
142
143

A central challenge in RL is discovering successful trajectories under a limited sampling budget. Although most approaches rely on the reward signal during training to evaluate learning quality, this signal is not always reliable or accurate. To demonstrate this, we conduct a simple experiment where we train Qwen2.5-7B (Team, 2024) on the DAPO-Math-170K (Yu et al., 2025), with periodic evaluation on MATH-500 (Hendrycks et al., 2021) test set. During the training phase, if all of its rollouts for a problem are incorrect, we will give an answer-level hint to the model.

143
144
145
146
147
148
149
150
151
152

Figure 2 shows the outcome of this experiments. Answer-level hint rapidly boosts rewards, creating the illusion of faster convergence. However, the plot on the bottom reveals a different story, as this apparent improvement does not translate into better generalization, with test accuracy stagnating at a low level. Furthermore, providing more detailed hints does not necessarily yield better outcomes, since excessive bias may cause the model’s behavior to deviate substantially from its current policy and potentially destabilize training.

153
154
155
156
157
158
159
160
161

The discrepancy between high training rewards and stagnant test accuracy raises a critical question: **why does an apparently strong learning signal fail to produce a generalizable policy?** Our analysis reveals that this problem originates from the severe answer leakage caused by answer-level hints. At a mechanistic level, these hints encourage large deviations from the current policy, generating updates with high importance ratios. These updates are then frequently clipped, which nullifies much of the potential learning signal. While this points to the importance of clipping, we find that its frequency alone is an incomplete indicator of training quality. The stability and diversity of the updates that survive clipping are also crucial for effective learning. To properly




Figure 2: A comparison of training rewards (top) and test accuracy (bottom). High rewards during training do not necessarily lead to high test accuracy, indicating that reward signals may be misleading indicators of model generalization.

162 diagnose these dynamics, we must quantify both how much of the learning signal survives clipping
 163 and the variability of those surviving updates. This motivates our proposal of a new set of metrics
 164 to evaluate exploration efficiency and quality.
 165

166 2.3 QUANTIFYING EXPLORATION EFFICIENCY AND QUALITY 167

168 The foundation for our new metrics is a direct analysis of the clipping mechanism, which constrains
 169 policy updates within a trust region (Schulman et al., 2015). While clipping improves stability, it
 170 also suppresses part of the original learning signal, making it difficult to evaluate how effectively
 171 the model leverages sampled trajectories. To quantitatively assess this, we focus on two factors that
 172 critically influence training quality: (I) the frequency with which policy updates are clipped, and (II)
 173 the variability of importance ratios. The first determines how much of the learning signal survives
 174 clipping, while the second reflects how stably the surviving updates are distributed. To capture these
 175 two aspects in a principled way, we introduce two complementary metrics: Effective Update Ratio
 176 (EUR), which measures how much of the learning signal survives clipping, and Update Consistency
 177 (UC), which characterizes the stability of the surviving updates. As we will show, these metrics
 178 further motivate a unified measure, *Affinity*, which combines both dimensions into a single indicator
 179 of exploration efficiency and update quality.

180 **Effective Update Ratio (EUR).** We use the EUR to quantify **how many token-level updates re-**
 181 **main unclipped** under the clipped objective introduced in PPO (Schulman et al., 2017), while stay-
 182 ing within the trust-region regime that supports TRPO’s monotonic improvement guarantee (Schul-
 183 man et al., 2015).

184 Consider a sampled trajectory $(s_1, a_1), \dots, (s_T, a_T)$. For each token step i , let s_i denote the prefix
 185 tokens before generation step i , and let a_i be the token generated at that step. The policy $\pi_\theta(a_i | s_i)$
 186 represents the current model, while $\pi_{\theta_{\text{old}}}(a_i | s_i)$ denotes the behavior policy used to collect the
 187 trajectory. We define the importance ratio as $r_i = \frac{\pi_\theta(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}$ and its log form $\ell_i = \log r_i$, which
 188 quantifies the local divergence between the updated policy and the behavior policy. A token is
 189 considered to remain within a trust region if its log ratio satisfies $|\ell_i| \leq \delta$, equivalently $e^{-\delta} \leq r_i \leq$
 190 e^δ . The term A_i denotes the token-level advantage, obtained by distributing the trajectory advantage
 191 across tokens. Given these definitions, we introduce the EUR:

$$192 \text{EUR} = \frac{\sum_i w_i \mathbf{1}|\ell_i| \leq \delta}{\sum_i w_i}, \quad w_i = |A_i|, \quad \ell_i = \log \frac{\pi_\theta(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}. \quad (3)$$

195 EUR measures the advantage-weighted fraction of token-level updates whose probability ratios re-
 196 main inside the trust region and therefore behave like unclipped PPO updates. This quantity is cru-
 197 cial because unclipped updates preserve the true policy gradient direction, whereas clipped updates
 198 either attenuate or nullify it, leading to ineffective learning even when reward appears high.

199 Importantly, we show that EUR provides (I) a principled estimate of the proportion of gradient
 200 contributions that remain unclipped under the PPO surrogate, and (II) a proxy for controlling the
 201 upper bound of policy divergence in the sense of the TRPO improvement guarantee. These two
 202 facts together imply that a high EUR indicates stable and meaningful policy improvement, while
 203 a low EUR signals that most gradient contributions are suppressed and the optimizer is effectively
 204 operating with a near-zero learning rate. We present the full derivations and theoretical justification
 205 in Appendix A.1.

206 **Update Consistency (UC).** While EUR measures how many token-level updates remain usable,
 207 it does not capture how consistent these effective updates are. In practice, even if a large propor-
 208 tion of updates fall within the trust region, their magnitudes may vary substantially: some updates
 209 correspond to very small log-ratios (i.e., conservative steps), while others lie near the trust-region
 210 boundary (i.e., aggressive steps). To distinguish stable updates from unstable ones, we introduce the
 211 UC metric.

212 Recall that $\ell_i = \log \frac{\pi_\theta(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}$ denotes the log-importance ratio, and A_i the token-level advantage.
 213 We focus on the subset of token steps whose updates remain within the trust region, $\mathcal{I} = \{i : |\ell_i| \leq$
 214 $\delta\}$, which correspond exactly to the unclipped updates in the PPO objective. Within this set, we
 215 define the weighted mean log-ratio as $\mu_\ell = \frac{\sum_{i \in \mathcal{I}} |A_i| \ell_i}{\sum_{i \in \mathcal{I}} |A_i|}$. With these quantities in place, we define

216 the UC as the advantage-weighted standard deviation of the log-importance ratios:
 217

$$218 \quad 219 \quad 220 \quad UC = \sqrt{\frac{\sum_{i \in \mathcal{I}} |A_i| (\ell_i - \mu_\ell)^2}{\sum_{i \in \mathcal{I}} |A_i|}}. \quad (4)$$

221 A low UC indicates that the effective updates exhibit small variability in their log-ratios and thus
 222 form a stable and coherent update direction. Conversely, a high UC indicates that the supposedly
 223 valid updates differ significantly in magnitude, with many lying near the trust-region boundary,
 224 which in turn suggests unstable or oscillatory learning dynamics. In other words, while EUR cap-
 225 tures the quantity of effective updates, UC captures their quality by measuring whether these updates
 226 move the policy in a consistent direction.

227 We further show in Appendix A.2 that UC is closely related to the variance of the local KL diver-
 228 gence and therefore reflects the stability of the policy update within the trust region. This connection
 229 provides the theoretical motivation for using UC alongside EUR to characterize the reliability of
 230 gradient-based policy improvement.

231 **Affinity.** While EUR quantifies how many token-level updates remain effective and UC measures
 232 how consistent those effective updates are, neither metric alone is sufficient to characterize the qual-
 233 ity of policy improvement. A high EUR may still correspond to unstable learning if the valid updates
 234 exhibit large variability (i.e., high UC), indicating that many of them lie near the trust-region bound-
 235 ary and pull the policy in conflicting directions. Conversely, a low UC provides little value when
 236 EUR is small, as almost all gradients are clipped and the policy barely changes despite being “con-
 237 sistent”.

238 A desirable training process therefore requires both a substantial number of effective updates (high
 239 EUR) and stable, coherent update magnitudes (low UC). To capture this joint requirement in a single
 240 measure, we introduce the unified metric *Affinity*. Let δ denote the log-ratio trust-region threshold
 241 used to define the unclipped set in EUR and UC, and let $\tau = \delta/2$ be a temperature parameter
 242 controlling the sensitivity of UC. We define:

$$243 \quad 244 \quad 245 \quad Affinity = EUR \cdot \exp\left(-\frac{UC}{\tau}\right). \quad (5)$$

246 This multiplicative formulation ensures that *Affinity* is high only when both conditions hold sim-
 247 taneously: EUR must be large enough to provide meaningful learning signal, and UC must be small
 248 enough to guarantee that those updates move the policy in a stable direction. The exponential term
 249 modulates the influence of UC, yielding a smooth but decisive penalty on inconsistent updates.

250 As a result, *Affinity* serves as a holistic indicator of exploration efficiency and training stability in
 251 online RL. It summarizes, in a single scalar quantity, both the amount and the quality of effective
 252 PPO updates. In Appendix A.3, we further discuss the theoretical motivation for this formulation
 253 and its relationship to trust-region optimization principles.

254 2.4 HINT: HELPING INEFFECTIVE ROLLOUTS NAVIGATE TOWARDS EFFECTIVENESS

255 The preceding analysis shows that excessively strong or answer-level hints used in prior work tend
 256 to degrade training quality by causing frequent clipping (low EUR) or unstable update magnitudes
 257 (high UC). **To improve *Affinity*, guidance should therefore avoid providing partial or complete
 258 solution steps and instead operate at an abstract, conceptual level that encourages the model
 259 to generate the reasoning autonomously.**

260 We operationalize this principle through the design of HINT. As illustrated in Figure 3, HINT is an
 261 adaptive mechanism that **guides the model toward productive reasoning trajectories using hints
 262 that are deliberately constrained to the abstract and conceptual level**. These hints avoid reveal-
 263 ing answers or intermediate steps and instead provide high-level reasoning cues that activate the
 264 model’s own problem-solving process, thereby preserving the high-*Affinity* update regime required
 265 for stable and effective GRPO training.

266 Formally, the HINT framework operates as a two-stage process. The first stage mirrors a standard
 267 GRPO update cycle. On the rollout stage, for a given problem q , the model begins by sampling
 268 a set of trajectories $\{o_1, o_2, \dots, o_G\}$ using its current policy. These trajectories are then evaluated




Figure 3: The HINT Framework: An Adaptive Two-Stage Rollout Process. HINT operates in two stages. **(I) Standard Rollout:** The model first samples trajectories from the original problem. If the rewards are non-sparse (at least one is correct), the process follows the standard GRPO update path. **(II) Hint-Augmented Rollout:** If, however, the rewards are sparse (all trajectories are incorrect), the hint mechanism is activated. The model then re-rolls out conditioned on a heuristic hint from a “teacher model”. This stage is designed to produce non-sparse rewards, turning a failed sample into a valuable learning opportunity.

by a reward model or predefined rules to obtain a set of rewards $\{r_1, r_2, \dots, r_G\}$. If these rewards are not sparse (i.e., at least one trajectory is correct), the process proceeds identically to the GRPO algorithm. The non-sparse rewards are used to compute advantages and perform a normal policy update.

The second stage, the hint-augmented rollout, is activated **only if the initial rewards from the first stage are sparse** (i.e., all trajectories are incorrect). In this scenario, where GRPO would stall due to a lack of learning signal, HINT intervenes. A pre-defined hint h is used to construct a hint-augmented query q_h . The model is then prompted to resample a new set of trajectories $\{o_1^h, o_2^h, \dots, o_G^h\}$, this time conditioned on q_h . These new, hinted trajectories are re-evaluated to produce a new set of rewards $\{r_1^h, r_2^h, \dots, r_G^h\}$. This rescue mechanism thus turns a failed rollout into a valuable learning opportunity. By providing a heuristic hint, it is intended to enable a meaningful gradient update, which enhances training efficiency. This is accomplished while the hint itself is carefully constructed to avoid degrading training *Affinity*.

A key design in our method is to decouple the prompts used for trajectory generation from those used for policy optimization. We refer to the input provided to the model during sampling as the rollout prompt, and the input used when updating the policy as the policy prompt.

When HINT is triggered, the rollout prompt may include the hint-augmented version of the problem in order to guide exploration and increase the likelihood of generating successful rollouts. However, the policy prompt is always kept strictly identical to the original problem without any hint. This separation ensures that hints influence only the sampling distribution—not the optimization objective—thereby preventing the model from implicitly learning to rely on hints. In other words, hints are used solely as an exploration aid, while the policy is optimized on the original tasks, preserving the correct training–inference alignment.

3 EXPERIMENTS

3.1 SETUP

Experimental Setup. Our experiments are conducted using Qwen2.5-7B and Qwen2.5-3B (Team, 2024) as backbone models. To ensure a fair and controlled comparison, we constructed a high-quality training set derived from the DAPO-Math-170K dataset (Yu et al., 2025). This process involved using Qwen2.5-72B-Instruct (Team, 2024) to generate four distinct reasoning trajectories for each problem. These outputs were then validated for correctness with Math Verify², from which we retained 30k fully correct samples to form our final training data. For baseline methods that

²<https://github.com/huggingface/Math-Verify>

324 require a ground-truth reference solution, we designated the shortest of the four correct trajectories
 325 for each problem.

326 **Benchmarks.** We evaluate the generalization ability of HINT on seven datasets, covering both in-
 327 distribution and out-of-distribution scenarios, without using any hint during evaluation. For mathe-
 328 matical reasoning, we adopt AIME24³, MATH-500 (Hendrycks et al., 2021), OlympiadBench (He
 329 et al., 2024), and Minerva (Lewkowycz et al., 2022), which are widely used benchmarks. Since the
 330 test sets of AIME24 are relatively small, we report avg@32, while for the other datasets we use
 331 pass@1. To assess complex reasoning and out-of-distribution generalization, we further evaluate on
 332 ARC-Challenge (Clark et al., 2018), GPQA-Diamond (Rein et al., 2024), and MMLU-Pro (Wang
 333 et al., 2024). To demonstrate HINT effectiveness, we conduct systematic experiments across multi-
 334 ple benchmarks.

335 **Baselines.** We compare HINT against several existing methods designed to improve rollout accu-
 336 racy rate or rollout efficiency in GRPO. The baselines include: (1)**LUFFY** (Yan et al., 2025): A
 337 hybrid approach that combines on-policy and off-policy training, ensuring that each sampled batch
 338 contains at least one correct trajectory. (2)**CHORD** (Zhang et al., 2025a): A method dynamically in-
 339 tegrating SFT as a weighted objective within on-policy RL. (3)**GHPO** (Liu et al., 2025b): A method
 340 that adaptively adjusts the hint length based on the ground-truth solution. If a shorter hint fails to
 341 solve the problem, the hint length is progressively increased until the correct answer is obtained.
 342 (4)**QuestA** (Li et al., 2025): A method constructs the hint by using the initial 50% of a reasoning
 343 trajectory generated by a larger, more capable model. (5)**BREAD** (Zhang et al., 2025b): A binary
 344 search-based method that identifies a hint length such that the model’s rollouts are neither all correct
 345 nor all incorrect, and uses this balanced point as the hint for training.

346 A comprehensive overview of our experimental configuration, including detailed prompts, hyper-
 347 parameters, and implementation settings for all methods, can be found in the Appendix B for full
 348 reproducibility.

350 3.2 MAIN RESULTS

351 We benchmarked our proposed method against several mainstream approaches, including both
 352 mixed-policy strategies and other hint-based methods. These experiments were conducted on two
 353 scales of backbone models: Qwen2.5-7B and Qwen2.5-3B. We report our results in Table 3. Our
 354 analysis reveals the following key findings:

355 **HINT enhances In-Distribution reasoning and teaches problem-solving skills.** HINT signif-
 356 icantly enhances the reasoning capabilities of models, achieving state-of-the-art performance on
 357 multiple in-distribution benchmarks. Models trained with HINT demonstrate substantial gains, with
 358 Qwen2.5-7B and Qwen2.5-3B showing average improvements of 9.0% and 6.8%, respectively, un-
 359 der scoring the effectiveness of our approach. We also observed an interesting emergent behavior
 360 during training: when a model encountered two similar, challenging problems, it would often rely
 361 on a hint for the first but then solve the second independently by applying the same reasoning pat-
 362 tern. This observation provides strong evidence that our heuristic and minimal hints teach the model
 363 how to reason about a class of problems, rather than simply encouraging it to memorize a solution
 364 path for a single instance.

365 **HINT generalizes to Out-of-Distribution problems by optimizing reasoning paths.** HINT also
 366 demonstrates strong generalization, enhancing the model’s ability to tackle novel problems. Even on
 367 out-of-distribution (OOD) test sets, models trained with HINT showed marked improvements. On
 368 the OOD test sets, models trained with HINT demonstrated strong generalization, with Qwen2.5-7B
 369 and Qwen2.5-3B achieving average performance gains of 7.4% and 1.6%, respectively, highlight-
 370 ing the method’s robust ability to generalize. This strong OOD performance is explained by a deeper
 371 phenomenon observed in our case studies. We found that the model successfully reappplies high-level
 372 reasoning methods from our hints, such as Proof by Contradiction to solve new OOD problems. This
 373 demonstrates that our method operates on a conceptual level, effectively teaching the model trans-
 374 ferable problem-solving paradigms rather than just answers. It is this acquisition of new, abstract
 375 reasoning skills that drives the model’s robust generalization.

376
 377 ³<https://huggingface.co/datasets/math-ai/aime24>

378
 379
 380
 381
 382
 383 Table 1: Main Performance Comparison of HINT against Baselines. HINT demonstrates significant
 384 performance gains on in-distribution datasets, improving the Qwen2.5-7B and Qwen2.5-3B models
 385 by **13.5%** and **6.8%**, respectively. The method also **shows strong generalization capabilities on**
 386 **out-of-distribution data**.

383 Methods	384 In-Distribution				385 Avg	386 Out-of-Distribution			387 Avg
	388 AIME	389 Math	390 Olympiad	391 Minerva		392 ARC	393 GPQA	394 MMLU	
Qwen2.5-7B									
Vanilla	9.8	50.2	34.0	19.5	28.4	85.3	25.6	46.0	52.3
GRPO	12.8	75.2	40.8	<u>31.2</u>	40.0	87.3	30.8	<u>56.6</u>	<u>58.2</u>
SFT	13.0	77.8	42.4	32.0	41.3	77.7	25.8	44.4	49.3
CHORD	<u>13.2</u>	74.4	40.0	31.2	39.7	86.6	30.1	51.2	56.0
LUFFY	12.6	70.2	38.6	30.8	38.1	87.2	32.2	46.8	55.4
GHPO	13.1	75.6	42.2	30.0	40.2	87.0	32.0	50.0	56.3
QuestA	13.1	73.6	38.8	28.6	38.5	<u>88.0</u>	26.6	53.2	55.9
BREAD	11.7	72.8	41.8	29.2	38.9	85.0	29.4	48.8	54.4
HINT	13.3	79.6	43.6	31.0	41.9	88.8	<u>31.8</u>	58.4	59.7
Qwen2.5-3B									
Vanilla	2.9	39.8	12.0	9.8	16.1	44.8	11.4	28.8	28.3
GRPO	4.3	44.0	18.2	12.2	19.7	45.0	11.8	28.0	28.3
SFT	5.0	48.0	20.8	14.0	22.0	20.4	7.6	20.2	16.1
CHORD	4.5	46.6	20.2	13.0	21.1	40.0	11.0	26.4	25.8
LUFFY	3.3	40.0	18.0	13.2	18.6	40.8	11.2	24.0	25.3
GHPO	4.0	42.2	19.6	12.8	19.7	<u>45.5</u>	12.0	28.2	28.6
QuestA	3.9	42.0	19.6	12.4	19.5	44.8	12.0	29.0	28.6
BREAD	4.1	44.4	<u>20.4</u>	13.4	20.6	<u>45.5</u>	<u>11.8</u>	<u>29.2</u>	<u>28.8</u>
HINT	<u>4.9</u>	48.6	<u>20.2</u>	<u>13.4</u>	<u>21.8</u>	48.8	<u>11.8</u>	30.2	29.9

403
 404
 405 The effectiveness of HINT scales with model size. Our results show that the benefits of HINT
 406 are more pronounced in larger models, with the performance gains for Qwen2.5-7B consistently
 407 outpacing those for Qwen2.5-3B across all evaluations. To understand the mechanism behind this
 408 trend, we analyzed the training rollouts and found a clear difference in how effectively each model
 409 leveraged the provided hints. A quantitative analysis confirmed that out of 100 randomly sampled
 410 rollouts where hints were provided to each model, Qwen2.5-7B produced a successful trajectory
 411 following the hint 34.0% more often than Qwen2.5-3B did. This superior efficacy in converting
 412 hints into successful outcomes directly explains the more pronounced performance gains, indicating
 413 that the greater capacity of larger models allows them to better capitalize on the abstract guidance
 414 offered by HINT.

415 3.3 TRAINING DYNAMICS

416 To investigate the impact of various off-policy strategies, we tracked the EUR, UC, and *Affinity*
 417 metrics for our method alongside several key baselines which detailed in Section 3.1, with the full
 418 training dynamics plotted in Figure 4. This analysis led to the following key observations.

419 In the early stages of training, the model shows strong resistance to off-policy data. As illustrated
 420 in the left plot of Figure 4, all three off-policy methods exhibit a sharp drop in EUR, indicating
 421 that clipping occurs very frequently at this stage. We call this initial period the "EUR Collapse
 422 Stage", where the model is highly resistant to the off-policy data and the clipping frequency is
 423 consequently high. With more training steps, the model gradually adapts, leading to reduced clipping
 424 frequency and eventual stabilization. Notably, compared to GHPO and LUFFY, HINT achieves
 425 a higher steady-state EUR, demonstrating its superior ability to help the model accommodate and
 426 leverage off-policy data.

427 Over-reliance on off-policy data often prevents the model from converging. As shown in the
 428 middle plot of Figure 4, both GHPO and LUFFY quickly reach high UC values at the beginning of
 429 training and remain at that level. This indicates persistently large variance in importance sampling,
 430 which results in unstable model updates and hampers convergence. In contrast, the UC of HINT




Figure 4: We record the EUR, UC, and *Affinity* metrics across different training processes to investigate the impact of various off-policy strategies on training. **Left:** EUR during training; **Middle:** UC during training; **Right:** *Affinity* during training. Overall, HINT most effectively alleviates the EUR collapse, avoids persistently high UC, and achieves higher *Affinity*, thereby enabling more stable and efficient training.

does not spike early on but instead indicates that our heuristic hints avoid causing large distributional shifts, allowing the policy updates to remain centered around a stable learning direction.

HINT enables the model to genuinely absorb the knowledge provided by hints. As presented in the right plot of Figure 4, the *Affinity* of HINT gradually approaches that of GRPO as training progresses. This implies that the model becomes increasingly capable of identifying which hints are truly useful. In other words, HINT enhances training efficiency and sample utilization in the early stages, while maintaining convergence trends consistent with GRPO in the later stages, thereby balancing early gains with eventual stability.

3.4 IN-DEPTH ANALYSIS

Does hinting truly enhance training effectiveness? We measured the number of valid samples (i.e., rollouts that are not entirely incorrect) generated by GRPO and HINT under an equal computational budget (8 hours of training). As shown in the top of Figure 5, although HINT produced slightly fewer total samples than GRPO, it yielded a greater number of valid samples. This indicates that HINT achieves higher training efficiency under the same time constraints, suggesting that hints guide the model toward more productive exploration trajectories rather than wasting updates on implausible rollouts.

From a broader perspective of the entire training process, the proportion of valid samples with HINT is higher than that of GRPO by 18.9%, further confirming that hinting improves the signal-to-noise ratio of training data. In other words, the gradient updates induced by HINT are more likely to be based on partially correct reasoning chains, thereby amplifying useful supervision signals and mitigating the destabilizing effects of noisy rollouts.

The dominance of valid rollouts under HINT suggests that hints not only improve rollout quality but also reshape the global optimization landscape by steering policy learning toward regions where correct reasoning is more likely to occur. This mechanism explains why HINT can achieve sustained improvements even without relying on answer leakage, ultimately leading to more robust and generalizable training outcomes.




Figure 5: Sampling Efficiency of HINT and GRPO at Different Training Stages. Under an equal budget, HINT yields **1,485 more valid samples** (top) and achieves a **18.9% higher final proportion of valid samples** (bottom).

486 **Does hinting affect the diversity of model’s outputs?**
 487 Entropy serves as a key metric for measuring generation
 488 diversity (Cheng et al., 2025; Zheng et al., 2025). Building
 489 on the training processes for HINT and the GHPO
 490 baseline detailed in Section 3.1, we further compared
 491 their dynamics by analyzing the average entropy of rea-
 492 soning trajectories throughout the training period. For
 493 each method, we separately computed the mean entropy
 494 on samples with and without hints.

495 As illustrated in Table 2, on the subset requiring hints, the
 496 entropy of HINT is notably higher than GHPO, which is
 497 answer-level hints. This is because answer-level hints of-
 498 ten provide a “half-completed” reasoning trajectory, forc-
 499 ing the model to follow a predetermined path with limited exploration. In contrast, ours do not
 500 disclose specific solution steps, leaving the reasoning process entirely up to the model and thereby
 501 encouraging broader exploration across different trajectories.

502 Even more surprisingly, we find that on samples where no hints are needed, GHPO still yield the
 503 lowest entropy compared to both GRPO and HINT. This suggests that long-term exposure to answer-
 504 level hints suppresses diversity at a deeper level: even when no hints are provided, the model’s ability
 505 to generate diverse reasoning paths is diminished.

506 4 RELATED WORK

507 **Reinforcement Learning for Large Language Model Reasoning.** Recent advances in RL ap-
 508 proaches have significantly enhanced the reasoning capabilities of LLMs. Large reasoning Mod-
 509 els (LRMs) such as OpenAI-o1 (Jaech et al., 2024), DeepSeek-R1 (Guo et al., 2025), and Kimi-
 510 1.5 (Team et al., 2025) achieve state-of-the-art performance on complex reasoning tasks (e.g., math-
 511 ematics, coding, scientific problem solving) by leveraging Reinforcement Learning from Verifiable
 512 Rewards (RLVR) (Liu et al., 2025a; Hu et al., 2025; Cui et al., 2025), where automatically checkable
 513 rules provide supervision signals. Compared to earlier methods like SFT or reinforcement learning
 514 from human feedback (RLHF), RLVR has shown superior generalization and robustness (Chu et al.,
 515 2025; Snell et al., 2025). Building on this paradigm, subsequent studies have proposed improved
 516 optimization strategies and structured prompting techniques that further strengthen reasoning capa-
 517 bilities (Schulman et al., 2017; Wang et al., 2020). Despite this progress, a critical failure mode
 518 for existing RL methods is reward sparsity, which occurs when all rollouts in a sample fail. Over-
 519 coming this challenge is essential for enhancing the stability and sample efficiency of training large
 520 reasoning models.

521 **Improving Rollout Efficiency in RL for LLMs.** A well-known challenge in methods such as
 522 GRPO is the vanishing gradient issue. This problem occurs when all trajectories in a sample group
 523 are incorrect, as the group advantage collapses to zero, yielding no gradient for policy updates (Shao
 524 et al., 2024; Guo et al., 2025). To mitigate this, some works have focused on injecting external,
 525 off-policy data to improve training efficiency and stability. This has been explored through two
 526 main strategies. Some methods use mixed-policy, replacing a portion of on-policy rollouts with
 527 complete, high-quality trajectories from off-policy datasets (Yan et al., 2025; Lin et al., 2025; Xu
 528 et al., 2025; Wang et al., 2025). Others employ partial supervision, providing segments of a ground
 529 truth to rescue failed rollouts (Li et al., 2025; Liu et al., 2025b; Zhang et al., 2025b). While these
 530 approaches effectively improve rollout efficiency, their over-reliance on off-policy data can misguide
 531 policy updates, steering the model toward non-generalizable or spurious solution paths.

532 5 CONCLUSION

533 In this work, we identify the problem of low training affinity caused by an over-reliance on off-policy
 534 data and propose HINT, an adaptive framework to resolve this trade-off. HINT significantly outper-
 535 forms strong baselines on competitive math benchmarks and demonstrates robust out-of-distribution
 536 generalization. Our work showcases a scalable and principled path toward more capable, self-
 537 improving reasoning models, with future work pointing towards extending HINT to new domains
 538 and modalities.

539 Table 2: We compare the average entropy for different methods on samples both with and without hints. The results consistently show that **HINT promotes higher entropy than answer-level hints across both scenarios.**

	w/ hint	w/o hint	All
GRPO	–	0.143	0.143
GHPO	0.123	0.141	0.129
HINT	0.188	0.198	0.193

540 **6 ETHICS STATEMENT**
541542 This work adheres to the ICLR Code of Ethics. In this study, no human subjects or animal ex-
543 perimentation was involved. All datasets used were sourced in compliance with relevant usage
544 guidelines, ensuring no violation of privacy. We have taken care to avoid any biases or discrimi-
545 natory outcomes in our research process. No personally identifiable information was used, and no
546 experiments were conducted that could raise privacy or security concerns. We are committed to
547 maintaining transparency and integrity throughout the research process.
548549 **7 REPRODUCIBILITY STATEMENT**
550551 We have made every effort to ensure that the results presented in this paper are reproducible. All
552 code and datasets have been made publicly available in an anonymous repository to facilitate repli-
553 cation and verification. The experimental setup, including training steps, model configurations, and
554 hardware details, is described in detail in the paper.
555556 Additionally, All datasets are publicly available, ensuring consistent and reproducible evaluation
557 results.
558559 We believe these measures will enable other researchers to reproduce our work and further advance
560 the field.
561562 **REFERENCES**
563564 Arash Ahmadian, Chris Cremer, Matthias Gallé, Marzieh Fadaee, Julia Kreutzer, Olivier Pietquin,
565 Ahmet Üstün, and Sara Hooker. Back to basics: Revisiting reinforce style optimization for learn-
566 ing from human feedback in llms. *arXiv preprint arXiv:2402.14740*, 2024.567 Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Wayne Xin Zhao, Zhenliang Zhang, and
568 Furu Wei. Reasoning with exploration: An entropy perspective. *arXiv preprint arXiv:2506.14758*,
569 2025.570 Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V
571 Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation
572 model post-training. *arXiv preprint arXiv:2501.17161*, 2025.573 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
574 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
575 *arXiv preprint arXiv:1803.05457*, 2018.
576577 Ganqu Cui, Lifan Yuan, Zefan Wang, Hanbin Wang, Wendi Li, Bingxiang He, Yuchen Fan, Tianyu
578 Yu, Qixin Xu, Weize Chen, et al. Process reinforcement through implicit rewards. *arXiv preprint*
579 *arXiv:2502.01456*, 2025.580 Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun
581 Mei, Jiashu Wang, et al. Areal: A large-scale asynchronous reinforcement learning system for
582 language reasoning. *arXiv preprint arXiv:2505.24298*, 2025a.
583584 Yuqian Fu, Tinghong Chen, Jiajun Chai, Xihuai Wang, Songjun Tu, Guojun Yin, Wei Lin, Qichao
585 Zhang, Yuanheng Zhu, and Dongbin Zhao. Srft: A single-stage method with supervised and
586 reinforcement fine-tuning for reasoning. *arXiv preprint arXiv:2506.19767*, 2025b.
587588 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
589 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
590 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.591 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,
592 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for
593 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*
594 *arXiv:2402.14008*, 2024.

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

597

598 Jian Hu. Reinforce++: A simple and efficient approach for aligning large language models. *arXiv*
 599 *preprint arXiv:2501.03262*, 2025.

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

603

604 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 605 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv*
 606 *preprint arXiv:2412.16720*, 2024.

607

608 Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. A survey on large language
 609 models for code generation. *arXiv preprint arXiv:2406.00515*, 2024.

610

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

614

615 Jiazheng Li, Hong Lu, Kaiyue Wen, Zaiwen Yang, Jiaxuan Gao, Hongzhou Lin, Yi Wu, and
 616 Jingzhao Zhang. Questa: Expanding reasoning capacity in llms via question augmentation. *arXiv*
 617 *preprint arXiv:2507.13266*, 2025.

618

619 Zhihang Lin, Mingbao Lin, Yuan Xie, and Rongrong Ji. Cppo: Accelerating the training of group
 620 relative policy optimization-based reasoning models. *arXiv preprint arXiv:2503.22342*, 2025.

621

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

625

626 Ziru Liu, Cheng Gong, Xinyu Fu, Yaofang Liu, Ran Chen, Shoubo Hu, Suiyun Zhang, Rui Liu,
 627 Qingfu Zhang, and Dandan Tu. Ghpo: Adaptive guidance for stable and efficient llm reinforce-
 628 ment learning. *arXiv preprint arXiv:2507.10628*, 2025b.

629

630 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-
 631 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-
 632 mark. In *First Conference on Language Modeling*, 2024.

633

634 John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region
 635 policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR,
 636 2015.

637

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

640

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

644

645 Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute
 646 optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth Inter-
 647 national Conference on Learning Representations*, 2025.

648

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

652

653 Qwen Team. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

648 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming
 649 Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-
 650 task language understanding benchmark. *Advances in Neural Information Processing Systems*,
 651 37:95266–95290, 2024.

652 Yuhui Wang, Hao He, and Xiaoyang Tan. Truly proximal policy optimization. In *Uncertainty in
 653 artificial intelligence*, pages 113–122. PMLR, 2020.

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

657 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 658 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 659 neural information processing systems*, 35:24824–24837, 2022.

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

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

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

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

670 Wenhao Zhang, Yuexiang Xie, Yuchang Sun, Yanxi Chen, Guoyin Wang, Yaliang Li, Bolin Ding,
 671 and Jingren Zhou. On-policy rl meets off-policy experts: Harmonizing supervised fine-tuning and
 672 reinforcement learning via dynamic weighting. *arXiv preprint arXiv:2508.11408*, 2025a.

673 Xuechen Zhang, Zijian Huang, Yingcong Li, Chenshun Ni, Jiasi Chen, and Samet Oymak.
 674 Bread: Branched rollouts from expert anchors bridge sft & rl for reasoning. *arXiv preprint
 675 arXiv:2506.17211*, 2025b.

676 Rosie Zhao, Alexandru Meterez, Sham Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach.
 677 Echo chamber: Rl post-training amplifies behaviors learned in pretraining. *arXiv preprint
 678 arXiv:2504.07912*, 2025.

679 Tianyu Zheng, Tianshun Xing, Qingshui Gu, Taoran Liang, Xingwei Qu, Xin Zhou, Yizhi Li, Zhou-
 680 futu Wen, Chenghua Lin, Wenhao Huang, et al. First return, entropy-eliciting explore. *arXiv
 681 preprint arXiv:2507.07017*, 2025.

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

APPENDIX

A THEORETICAL FOUNDATIONS OF EUR, UC, AND AFFINITY

A.1 PROOFS FOR EUR

In this section, we provide the theoretical justification for the two main claims made in the main paper regarding the EUR: (I) EUR estimates the fraction of unclipped PPO gradient contributions (Schulman et al., 2017); (II) EUR serves as a proxy for bounding policy divergence in the sense of TRPO’s monotonic improvement guarantee (Schulman et al., 2015).

A.1.1 PRELIMINARIES

For each token step i , let

$$r_i = \frac{\pi_\theta(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}, \quad \ell_i = \log r_i.$$

PPO optimizes a clipped surrogate objective (Schulman et al., 2017), defined as

$$L_{\text{CLIP}}(\theta) = \hat{\mathbb{E}}_i \left[\min (r_i A_i, \text{clip}(r_i, 1 - \varepsilon, 1 + \varepsilon) A_i) \right], \quad (6)$$

and then maximizes $L_{\text{CLIP}}(\theta)$ with respect to θ .

Let $\mathcal{I} = \{i : |r_i - 1| \leq \varepsilon\}$ denote the set of unclipped updates and \mathcal{C} the clipped ones. The gradient of equation 6 decomposes as:

$$\nabla_\theta L_{\text{CLIP}} = \mathbb{E}[\nabla_\theta(r_i A_i) \mathbf{1}(i \in \mathcal{I})] + \mathbb{E}[\nabla_\theta(r_i^{\text{clip}} A_i) \mathbf{1}(i \in \mathcal{C})]. \quad (7)$$

As noted in Schulman et al. (2017), gradients from clipped terms either vanish or are directionally distorted, while terms in \mathcal{I} preserve the correct policy gradient direction.

The Effective Update Ratio is defined in the main paper as:

$$\text{EUR} = \frac{\sum_i |A_i| \mathbf{1}(|\ell_i| \leq \delta)}{\sum_i |A_i|}. \quad (8)$$

A.1.2 PROOF OF CLAIM (I): EUR ESTIMATES THE FRACTION OF UNCLIPPED PPO GRADIENT CONTRIBUTIONS

We show that EUR provides a principled empirical estimate of the proportion of gradient contributions arising from unclipped PPO updates. Recall that, for token-level PPO, the unclipped surrogate gradient at position i is

$$g_i = \nabla_\theta(r_i A_i) = A_i r_i \nabla_\theta \log \pi_\theta(a_i | s_i),$$

where $r_i = \frac{\pi_\theta(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}$. For updates that fall inside the trust region (i.e., $i \in \mathcal{I}$ with $|\ell_i| \leq \delta$), we have $r_i = e^{\ell_i} \approx 1$ because ℓ_i is small. Thus the gradient magnitude simplifies to

$$\|g_i\| \approx |A_i| \|\nabla_\theta \log \pi_\theta(a_i | s_i)\|,$$

and variations in $\|g_i\|$ across token steps are dominated by variations in $|A_i|$. Since $\|\nabla_\theta \log \pi_\theta(a_i | s_i)\|$ is locally bounded and does not change substantially across nearby policy iterates, the total contribution of unclipped updates to the overall gradient is proportional to

$$\mathbb{E}[|A_i| \mathbf{1}(i \in \mathcal{I})].$$

Similarly, the total gradient magnitude (including both clipped and unclipped updates) is proportional to $\mathbb{E}[|A_i|]$. Therefore, the fraction of gradient contributions that originate from unclipped updates is

$$\frac{\mathbb{E}[|A_i| \mathbf{1}(i \in \mathcal{I})]}{\mathbb{E}[|A_i|]}.$$

By construction, this is exactly the EUR. Consequently, EUR serves as an effective estimator for the fraction of gradient contributions that are not suppressed by clipping.

756 A.1.3 PROOF OF CLAIM (II): EUR CONTROLS POLICY DIVERGENCE IN THE TRPO SENSE
757758 TRPO (Schulman et al., 2015) establishes a monotonic improvement lower bound dependent on the
759 KL divergence:

760
$$\eta(\theta) \geq L_{\theta_{\text{old}}}(\theta) - C \cdot D_{\text{KL}}^{\max}(\pi_{\theta_{\text{old}}}, \pi_{\theta}), \quad (9)$$

761 where C is a constant depending on γ and ϵ . The token-level empirical KL divergence can be
762 approximated by the expectation of log-ratios:

763
$$D_{\text{KL}}(\pi_{\theta_{\text{old}}} \| \pi_{\theta}) \approx \mathbb{E}_{s, a \sim \pi_{\text{old}}} [|\ell_i|].$$

764

765 Recall that EUR is the advantage-weighted fraction of updates within the trust region ($|\ell_i| \leq \delta$).
766 Let $\mathcal{C} = \{i : |\ell_i| > \delta\}$ denote the set of clipped updates. The relationship between EUR and the
767 probability mass of \mathcal{C} depends on the distribution of advantages.768 **Assumption 1.** *The expected magnitude of advantages for clipped updates is lower bounded by a
769 factor of the global expected magnitude, i.e., $\mathbb{E}[|A_i| \mid i \in \mathcal{C}] \geq \alpha \mathbb{E}[|A_i|]$ for some $\alpha > 0$.*770 Under this mild assumption, we can relate EUR to the probability of clipping $P(\mathcal{C})$:

772
$$1 - \text{EUR} = \frac{\sum_{i \in \mathcal{C}} |A_i|}{\sum_{\text{all}} |A_i|} \approx \frac{P(\mathcal{C}) \cdot \mathbb{E}[|A_i| \mid \mathcal{C}]}{\mathbb{E}[|A_i|]} \geq \alpha P(\mathcal{C}).$$

773

774 This implies $P(\mathcal{C}) \leq \frac{1 - \text{EUR}}{\alpha}$. Conversely, the contribution to the KL divergence from clipped
775 samples is lower bounded:

776
$$D_{\text{KL}} \geq P(\mathcal{C}) \cdot \min_{i \in \mathcal{C}} |\ell_i| > P(\mathcal{C}) \cdot \delta.$$

777

778 If EUR is low (close to 0), the advantage mass is concentrated in \mathcal{C} . Unless the advantages in \mathcal{C}
779 are negligibly small (contradicting meaningful exploration), a low EUR implies a significant $P(\mathcal{C})$,
780 which forces D_{KL} to exceed the trust region boundary δ . Therefore, maintaining a high EUR is a
781 necessary proxy for constraining D_{KL} and preserving the validity of the TRPO bound.782 A.1.4 SUMMARY
783784 Taken together, the results above show that EUR simultaneously quantifies the fraction of gradi-
785 ent mass preserved by the unclipped PPO surrogate and provides a practical handle on the policy
786 divergence term appearing in TRPO’s monotonic improvement bound. Consequently, a high EUR
787 indicates that most updates lie within a stable trust-region regime where policy gradients remain
788 informative, whereas a low EUR reveals that clipped updates dominate the optimization process,
789 leading to vanishing effective gradients and ineffective learning.790 A.2 PROOFS FOR UC
791792 In this section, we provide the theoretical justification for the UC metric introduced in the main
793 paper. We show that UC can be interpreted as (I) an advantage-weighted measure of variability in
794 local log-importance ratios among unclipped updates, and (II) a proxy for the variance of the local
795 KL divergence, which is closely tied to the stability of policy updates.796 A.2.1 PRELIMINARIES
797798 Recall that for each token step i , we define

800
$$r_i = \frac{\pi_{\theta}(a_i \mid s_i)}{\pi_{\theta_{\text{old}}}(a_i \mid s_i)}, \quad \ell_i = \log r_i,$$

801

802 and the trust-region condition $|\ell_i| \leq \delta$ identifies the set of unclipped updates:

803
$$\mathcal{I} = \{i : |\ell_i| \leq \delta\}.$$

804 The token-level advantages are denoted by A_i , and we use the absolute values $|A_i|$ as importance
805 weights on the contribution of each token.806 Within the set \mathcal{I} , we define the advantage-weighted mean log-ratio:

807
$$\mu_{\ell} = \frac{\sum_{i \in \mathcal{I}} |A_i| \ell_i}{\sum_{i \in \mathcal{I}} |A_i|}, \quad (10)$$

808

810 and the UC is given by the advantage-weighted standard deviation:
 811

$$812 \quad 813 \quad 814 \quad \text{UC} = \sqrt{\frac{\sum_{i \in \mathcal{I}} |A_i| (\ell_i - \mu_\ell)^2}{\sum_{i \in \mathcal{I}} |A_i|}}. \quad (11)$$

815 **A.2.2 UC AS A MEASURE OF VARIABILITY AMONG EFFECTIVE UPDATES**
 816

817 As shown in equation 11, UC is precisely the standard deviation of the log-importance ratios ℓ_i
 818 over the set of effective updates \mathcal{I} . A small UC indicates that the ℓ_i values within \mathcal{I} are tightly
 819 concentrated around their weighted mean μ_ℓ , implying that the magnitudes of the effective updates
 820 are consistent and that the resulting policy changes are approximately uniform across token positions.
 821 In contrast, a large UC reflects substantial variability among the ℓ_i values: some effective
 822 updates correspond to very small log-ratios (i.e., conservative steps), while others lie close to the
 823 trust-region boundary (i.e., aggressive steps). Such heterogeneity results in uneven and potentially
 824 unstable policy updates.
 825

Formally, define the normalized weights

$$826 \quad 827 \quad 828 \quad \tilde{w}_i = \frac{|A_i|}{\sum_{j \in \mathcal{I}} |A_j|}, \quad i \in \mathcal{I}.$$

829 Then equation 11 can be rewritten as

$$830 \quad 831 \quad 832 \quad \text{UC}^2 = \sum_{i \in \mathcal{I}} \tilde{w}_i (\ell_i - \mu_\ell)^2,$$

833 which is the weighted variance of ℓ_i under the empirical distribution induced by the advantages $|A_i|$.
 834 Thus UC quantifies how “spread out” the log-ratios are among those updates that are not clipped.

835 **A.2.3 RELATION BETWEEN UC AND GRADIENT VARIANCE**
 836

837 We now connect UC to the variance of the policy gradient updates. Consider the gradient contribu-
 838 tion scale for a single token i within the trust region ($i \in \mathcal{I}$), defined as $X_i = A_i r_i \approx A_i (1 + \ell_i)$.
 839 The stability of training depends on the variance of this update scale. Assuming that the advantage
 840 A_i and the log-ratio ℓ_i are uncorrelated within the local trust region, we can apply the variance
 841 decomposition formula $\text{Var}(XY) \approx \mathbb{E}[X]^2 \text{Var}(Y) + \mathbb{E}[Y]^2 \text{Var}(X) + \text{Var}(X)\text{Var}(Y)$.
 842

Evaluating $\text{Var}(X_i)$ where $X_i \approx A_i + A_i \ell_i$:

$$843 \quad \text{Var}(g_i) \propto \text{Var}(A_i(1 + \ell_i)) \approx \text{Var}(A_i) + \text{Var}(A_i \ell_i). \quad (12)$$

844 The first term $\text{Var}(A_i)$ represents the inherent variance of the reward structure (baseline variance),
 845 which is irreducible by policy constraint. The second term captures the variance introduced by the
 846 policy shift. Applying the decomposition to $A_i \ell_i$:

$$847 \quad \text{Var}(A_i \ell_i) \approx \mathbb{E}[A_i^2] \text{Var}(\ell_i) + \mathbb{E}[\ell_i]^2 \text{Var}(A_i). \quad (13)$$

848 Inside the trust region, ℓ_i is centered near 0, making $\mathbb{E}[\ell_i]^2$ small. Thus, the dominant component of
 849 the induced variance is:

$$850 \quad \text{Var}_{\text{induced}} \approx \mathbb{E}[A_i^2] \cdot \text{Var}(\ell_i).$$

851 Recall that UC^2 is defined as the advantage-weighted variance of ℓ_i . Although strictly distinct from
 852 the unweighted $\text{Var}(\ell_i)$, they are empirically aligned. As shown in equation 13, UC acts as a mul-
 853 tiplicative gain on the gradient variance. A high UC amplifies the gradient noise proportional to
 854 the squared advantages $\mathbb{E}[A_i^2]$, destabilizing the update direction. Thus, minimizing UC is theoreti-
 855 cally justified to dampen the variance of policy updates specifically arising from diverse impor-
 856 tance ratios.
 857

858 **A.2.4 RELATION BETWEEN UC AND LOCAL KL VARIABILITY**
 859

860 We next relate UC to the variability in local KL divergence. The per-state KL divergence between
 861 the old and new policy can be expressed as
 862

$$863 \quad D_{\text{KL}}(\pi_{\theta_{\text{old}}}(\cdot | s) \| \pi_{\theta}(\cdot | s)) = \mathbb{E}_{a \sim \pi_{\theta_{\text{old}}}(\cdot | s)} [\log r(a, s)].$$

864 At the token level, the empirical KL is estimated by averaging ℓ_i over samples from $\pi_{\theta_{\text{old}}}$. Thus, the
 865 variability of ℓ_i within \mathcal{I} directly reflects how much the local per-state KL can fluctuate around its
 866 mean.

867 Since TRPO’s monotonic improvement bound (Schulman et al., 2015) depends on controlling KL,
 868 large fluctuations in ℓ_i (i.e., a high UC) indicate that some states experience near-boundary pol-
 869 icy changes even when the average KL remains small. This effectively weakens the trust-region
 870 assumption and can cause oscillatory learning dynamics. By contrast, a low UC ensures that the
 871 per-token KL changes are not only small on average but also uniformly bounded, leading to more
 872 reliable surrogate optimization.

873

874 A.2.5 SUMMARY

875

876 In summary, UC captures the internal stability of policy updates within the trust region by measuring
 877 the advantage-weighted variance of log-importance ratios among unclipped samples. A low UC
 878 implies that effective updates move the policy in a coherent and conservative manner, while a high
 879 UC reveals that updates, though nominally “valid,” are heterogeneous and prone to inducing unstable
 880 or oscillatory behavior. Together with EUR, UC thus provides a complementary view of both the
 881 quantity and the quality of effective policy updates during training.

882

883 A.3 THEORETICAL DISCUSSION OF AFFINITY

884

885 In this section, we provide the theoretical motivation for combining the EUR and UC into the unified
 886 *Affinity* metric introduced in the main paper. We show that *Affinity* captures the joint requirements for
 887 effective and stable policy updates in PPO-style RL, and we relate its form to principles underlying
 888 trust-region optimization.

889

890 A.3.1 PRELIMINARIES

891

Recall the definitions of EUR and UC from the main paper. Let

892
 893
 894

$$\ell_i = \log \frac{\pi_{\theta}(a_i | s_i)}{\pi_{\theta_{\text{old}}}(a_i | s_i)}$$

895

denote the log-importance ratio at token step i , and let $\mathcal{I} = \{i : |\ell_i| \leq \delta\}$ be the set of unclipped
 896 updates under the PPO objective. EUR measures the fraction of effective updates:

897

898
 899

$$\text{EUR} = \frac{\sum_i |A_i| \mathbf{1}(i \in \mathcal{I})}{\sum_i |A_i|},$$

900

901

while UC quantifies the internal variability of those updates:

902

903

904

905

906 A.3.2 RATIONALE FOR COMBINING EUR AND UC

907

908

909

910

911

As shown in Appendix A.1, EUR provides an unbiased estimate of the proportion of gradient mass
 908 preserved by the unclipped PPO surrogate. Hence, a high EUR indicates that most updates meaning-
 909 fully contribute to the policy gradient. However, EUR alone cannot ensure stability: if the log-ratios
 910 within \mathcal{I} vary widely (high UC), many of those “effective” updates may be close to the trust-region
 911 boundary and induce oscillatory policy shifts.

912

913

914

915

916

Appendix A.2 further shows that UC approximates the variance of token-level policy divergence
 913 and characterizes the consistency of unclipped gradients. Yet UC by itself is also insufficient: a
 914 perfectly consistent set of updates (low UC) provides little value when EUR is small, since almost
 915 all gradients are clipped and the policy barely moves.

917

Thus, a high-quality update requires satisfying both conditions simultaneously: a sufficiently large
 918 proportion of effective updates (high EUR) and low variability among them (low UC).

918 A.3.3 AFFINITY AS A JOINT STABILITY-EFFICIENCY INDICATOR
919920 To encode this joint requirement in a single quantity, we define the *Affinity* metric:
921

922
$$\text{Affinity} = \text{EUR} \cdot \exp\left(-\frac{\text{UC}}{\tau}\right), \quad \tau = \delta/2. \quad (14)$$

923

924 This multiplicative formulation has two motivations:
925926 **Logical conjunction.** The product structure ensures that a failure in either condition (low EUR
927 or high UC) produces a proportionally low *Affinity*. This structure captures the fact that effective
928 PPO-style updates require both conditions to be satisfied simultaneously, rather than individually.
929930 **Exponential penalty on inconsistency.** Since UC measures weighted variance in log-ratios, an
931 exponential term $\exp(-\text{UC}/\tau)$ acts analogously to an inverse smoothness regularizer, sharply pe-
932 nalizing updates near the trust-region boundary. The temperature term $\tau = \delta/2$ stabilizes the scaling
933 and ensures that the penalty becomes substantial when UC approaches the trust-region limit.
934935 A.3.4 RELATIONSHIP TO TRUST-REGION OPTIMIZATION
936937 Trust-region methods (including TRPO) rely on bounding the KL divergence to guarantee mono-
938 tonic policy improvement. While EUR controls the fraction of updates that satisfy the trust-region
939 condition and thus reflects the mean KL contribution, UC characterizes the variability of the local KL
940 divergence within that region by capturing the variance of the log-importance ratios. Consequently,
941 *Affinity* integrates both aspects of policy divergence: high *Affinity* indicates that the empirical KL
942 remains not only small (as ensured by high EUR) but also stable across updates (as ensured by low
943 UC), aligning with the conditions under which trust-region guarantees are most effective.
944945 A.3.5 SUMMARY
946947 *Affinity* synthesizes two complementary perspectives on PPO update quality: **(I) how many updates
948 remain effective (EUR), and (II) how consistent those updates are (UC).** The multiplicative for-
949 mulation in equation 14 captures the synergy required for reliable policy improvement and provides
950 a practical scalar diagnostic for monitoring exploration efficiency and training stability.
951952 B EXPERIMENTAL DETAILS
953954 B.1 DETAILED SETUP
955956 **Platform.** All of our experiments are conducted on workstations equipped with 8 NVIDIA A100
957 PCIe GPUs with 80GB memory.
958959 **Training Data.** The training was performed using a carefully selected subset of the DAPO-Math-
960 170K dataset (Yu et al., 2025). As the original dataset lacks ground-truth solutions, we curated our
961 own by first using Qwen2.5-72B-Instruct to generate four reasoning trajectories for each problem.
962 After validating the final answers with *Math-verify*, we compiled a high-quality training set of 30k
963 problems for which all four generated trajectories were correct. For baselines requiring a ground
964 truth, the most token-efficient of these four correct trajectories was designated as the ground truth.
965 For our methods, we pre-generated the required heuristic hints for the entire 30k-sample training set
966 using Qwen2.5-72B-Instruct. The prompts used in the above process will be detailed in Section B.2.
967968 **Important Parameters of HINT.** HINT is implemented based on the open-source RL framework
969 lsrl⁴. The RL algorithm employs the GRPO advantage estimator with no KL penalty (kl_coef is set
970 to 0.0). The clipping parameter ϵ is set to 0.2. For each group, 8 answers are generated, and the
971 training batch size is set to 2. Distributed training utilizes the DeepSpeed library with the AdamW
972 optimizer and a learning rate of 1e-6. The *train batch size* is set to 8, *gen batch size* is set to 32,
973 *accum steps* is set to 64, *gen update steps* is set to 128, *temperature* is set to 0.9, *max response* is
974 set to 4096. Mixed-precision training with BF16 is enabled. Memory optimization employs ZeRO
975 Stage 2, with optimizer state offloading to CPU.
976977
978 ⁴<https://github.com/lsdefine/lsrl>

972 **Important Parameters of Other Baselines.** For baselines with publicly available code repositories,
 973 we utilized their official implementations and the parameters specified in their respective publica-
 974 tions. For methods without public code, such as BREAD(Zhang et al., 2025b) and QuestA(Li et al.,
 975 2025), we reproduced their results using the lsrl framework, strictly adhering to the experimental
 976 parameters detailed in their papers.

977 **Reward Setup.** For our experiments, we employ a sparse, binary reward function. The reward
 978 is determined exclusively by the correctness of the final answer in a model’s generated trajectory.
 979 We use the *Math-Verify* tool for automatic verification, assigning a reward of **+1** for a correct final
 980 answer and **0** for an incorrect one.
 981

982 B.2 PROMPT LIST
 983

984 Prompt Template for GRPO
 985

986 **System:** You are a helpful AI assistant. A conversation takes place between the User and
 987 the Assistant. The User asks a question, and the Assistant solves it. Please help me solve
 988 this question. Wrap only the final answer in $\boxed{\text{ }}$.
 989

990 **Question:** [Question]
 991

992 **User:**
 993

994 Prompt Template for HINT
 995

996 **System:** You are a helpful AI assistant. A conversation takes place between the User and
 997 the Assistant. The User asks a question, and the Assistant solves it. Please help me solve
 998 this question. Wrap only the final answer in $\boxed{\text{ }}$.
 999

1000 **Hint:** Here are some key information provided to assist you in solving the problem: [Hint]
 1001

1002 **Question:** [Question]
 1003

1004 **User:**
 1005

1006 Prompt Template for Generating hints
 1007

1008 **System:**
 1009

1010 * Role and Goal
 1011

1012 You are a top-tier problem-solving expert and a master educator. Your goal is not to solve
 1013 the problem, but to distill the single most critical ”Core Insight” or ”Aha! Moment” required
 1014 to find the solution.
 1015

1016 * Core Task
 1017

1018 You will be given a [Question] and its final [Answer]. Your sole job is to reverse-engineer the
 1019 most likely solution path and identify the crucial ”mental bridge”—the non-obvious insight,
 1020 change in perspective, or core principle—that unlocks the problem.
 1021

1022 * Thinking Framework
 1023

1024 Analyze the Gap: First, understand the [Question] and look at the [Answer]. The core
 1025 difficulty lies in the conceptual space between them. What makes bridging this gap non-

1026
 1027 trivial? Reconstruct the "Hidden" Step: Mentally construct the most elegant solution path.
 1028 In that path, what is the single most pivotal, non-obvious leap of logic or application of a
 1029 principle that a student is most likely to miss? Distill the Insight: Condense this pivotal leap
 1030 into an extremely short, potent, and core-focused sentence. This sentence is the key that
 1031 unlocks the door, not the map of the room.
 1032
 1033 * Constraints
 1034 Absolute Brevity: The insight must be a single sentence, ideally under 20 words. No Spoil-
 1035 ers: The insight must not reveal any part of the [Answer] or the specific numbers used to
 1036 calculate it. Inspirational, Not Instructional: It should inspire thought ("heuristic"), not pro-
 1037 vide a step-by-step recipe ("algorithmic"). Target the Crux: It must address the most critical
 1038 linchpin that makes the entire solution possible.
 1039
 1040 * Output Format
 1041 Directly output the single, distilled "Core Insight". Do not include any other explanations,
 1042 headings, or conversational text.
 1043
 1044 **User:**
 1045
 1046 **### Question:**
 1047 [Question]
 1048
 1049 **### Answer:**
 1050 [Answer]

1051
 1052 Prompt Template for Generating Ground Truth
 1053
 1054 **System:** You are a helpful AI assistant. A conversation takes place between the User and
 1055 the Assistant. The User asks a question, and the Assistant solves it. Please help me solve
 1056 this question. Wrap only the final answer in $\boxed{\text{ }}$.
 1057
 1058 **Question:** [Question]
 1059
 1060 **User:**

1062
 1063 Prompt Template for Evaluation
 1064
 1065 **System:** You are a helpful AI assistant. A conversation takes place between the User and
 1066 the Assistant. The User asks a question, and the Assistant solves it. Please help me solve
 1067 this question. Wrap only the final answer in $\boxed{\text{ }}$.
 1068
 1069 **Question:** [Question]
 1070
 1071 **User:**

1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079

1080
1081

C FURTHER ANALYSIS

1082
1083

C.1 DETAILS OF HINT’S ENTROPY

1084

HINT Encourages Sustained Exploration. The entropy of the generation distribution serves as a key indicator of exploration diversity. As illustrated in Figure 6, HINT avoids the rapid entropy collapse observed in GRPO during the early stages of training. Instead, HINT maintains a consistently high level of entropy, indicating that the model actively explores when first introduced to the hints. This period of high exploration corresponds directly to the “EUR collapse” phase (discussed in Section 3.3), explaining that while the model initially resists the off-policy guidance, it is nevertheless engaged in a productive and diverse search of the solution space.

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

During the middle stages of training, HINT’s entropy does not decrease monotonically. It exhibits periodic increases. We attribute this to the model encountering novel types of hints and adapting its exploratory behavior to learn how to utilize them. Crucially, even after the policy stabilizes in the later stages, HINT maintains a significantly higher entropy level than GRPO. This provides strong evidence that HINT’s heuristic guidance successfully fosters more continuous and diverse exploration, preventing the policy from prematurely converging to a deterministic state.

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

1124

1125

C.2 DETAILS OF HINT’S ACCURACY

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1122

1123

Our results reveal an interesting trade-off: while the off-policy guidance from HINT may initially slow the rate of convergence, it ultimately enables the model to achieve a higher performance ceiling. As shown in Figure 7, HINT initially exhibits a slower rate of accuracy improvement compared to GRPO. This initial lag is consistent with the early training stages where the model shows resistance to the heuristic hints and has not yet learned to leverage them effectively. However, as training progresses, the model begins to adapt and utilize the guidance. This leads to an accelerated learning rate after approximately 100 steps, with HINT’s accuracy eventually surpassing GRPO’s and reaching a higher final value. This dynamic suggests that the model requires an adaptation period to master the use of heuristic hints, but once learned, this skill allows it to develop stronger and more robust capabilities.

1124

1125

C.3 QWEN3-4B OUTCOME

1126

1127

1128

1129

1130

1131

1132

1133

To verify the effectiveness of HINT on other models, we have supplemented our experiments with results on Qwen3-4B. These results underscore the distinct advantages of HINT in balancing effectiveness and generalization. Compared to the GRPO baseline, HINT delivers comprehensive gains, raising the average scores by 4.1 points on in-domain tasks and 5.6 points on out-of-domain benchmarks, thereby validating the efficacy of our framework. Furthermore, when contrasted with SFT, HINT demonstrates superior robustness; although HINT trails SFT marginally by 0.2 points on in-domain tasks, which is an expected outcome of supervised fitting, it significantly outperforms SFT on out-of-domain benchmarks with a substantial lead of 14.9 points. This stark contrast confirms that while SFT tends to overfit to the domain, HINT cultivates transferable reasoning skills.




Figure 6: HINT Prevents Entropy Collapse and Encourages Sustained Exploration. HINT maintains a high entropy level, especially in the early stages, and stabilizes at a significantly higher value. This demonstrates that HINT’s heuristic guidance fosters more continuous and diverse exploration, preventing premature policy convergence.




Figure 7: Accuracy of Different Methods. HINT Achieves Higher Final Accuracy Despite Slower Initial Convergence.

1134 Table 3: Main Performance Comparison of HINT against Baselines on Qwen3-4B. HINT
 1135 consistently outperforms the GRPO baseline, achieving average improvements of **4.1 points** on in-
 1136 distribution tasks and **5.6 points** on out-of-distribution benchmarks, validating the method’s robust-
 1137 ness across different model generations.

Methods	In-Distribution				Avg	Out-of-Distribution			Avg
	AIME	Math	Olympiad	Minerva		ARC	GPQA	MMLU	
Qwen3-4B									
Vanilla	53.3	82.0	61.8	32.2	57.3	82.5	28.0	60.4	57.0
GRPO	58.8	84.8	63.9	36.4	61.0	87.0	28.4	61.2	58.9
SFT	<u>68.6</u>	<u>87.6</u>	<u>66.1</u>	<u>39.0</u>	<u>65.3</u>	72.0	22.4	54.4	49.6
LUFFY	59.6	85.2	64.4	37.1	61.6	83.0	20.1	62.2	55.1
CHORD	59.1	88.0	64.0	36.2	61.8	85.0	<u>33.2</u>	66.6	61.6
GHPO	62.3	87.0	66.6	34.2	62.5	87.5	29.6	<u>68.4</u>	<u>61.8</u>
QuestA	60.3	84.4	63.0	36.6	61.1	88.4	29.6	62.3	60.1
BREAD	65.1	85.0	63.4	35.4	62.2	86.0	28.4	68.2	60.9
HINT	68.9	87.2	65.0	39.3	<u>65.1</u>	<u>88.1</u>	35.4	70.1	64.5

C.4 CASE STUDY

Case Study

Question Let a and b be positive integers such that all but 2009 positive integers are expressible in the form $ma + nb$, where m and n are nonnegative integers. If 1776 is one of the numbers that is not expressible, find $a + b$.

Hint: How many positive integers cannot be expressed in the form $ma + nb$ with nonnegative integers m, n ?

Before the Hint:

By the two-coin Frobenius result, for coprime positive integers a, b the number of positive integers not representable as $ma + nb$ equals $\frac{(a-1)(b-1)}{2}$.

So

$$\frac{(a-1)(b-1)}{2} = 2009$$

$$(a-1)(b-1) = 4018$$

Factor $4018 = 41 \cdot 98$. Take $a = 42, b = 99$.

Hence $a + b = 141$.

Answer: 141 

After the Hint:

By the two-coin Frobenius result, for coprime positive integers a, b the number of positive integers not representable as $ma + nb$ equals $\frac{(a-1)(b-1)}{2}$.

So

$$\frac{(a-1)(b-1)}{2} = 2009$$

$$(a-1)(b-1) = 4018$$

Candidates:

$(2, 4019), (3, 2010), (8, 575),$

$(15, 288), (42, 99), (50, 83)$.

Require $\text{gcd}(a, b) = 1$,

leaving $(2, 4019), (8, 575), (50, 83)$.

Check 1776: representable for first two, not for $(50, 83)$.

Thus $a + b = 133$.

Answer: 133 

1188

1189

1190

Why the hints works in this question?

1191

1192

1193

1194

1195

1196

The key to solving this problem is identifying it as an instance of the **Frobenius Coin Problem**. This theorem states that for two coprime positive integers, a and b , the number of positive integers not expressible in the form $ma + nb$ is precisely $\frac{(a-1)(b-1)}{2}$. Without guidance, the model correctly sets up the equation $(a - 1)(b - 1) = 4018$ but may hastily select a factor pair, such as $(42, 99)$, while neglecting the crucial precondition that a and b must be coprime. As shown in the "Before the Hint" example, this oversight leads to an incorrect answer.

1197

1198

1199

1200

1201

1202

1203

1204

The provided hint is designed to steer the model away from this pitfall. By asking about the *number* of non-representable integers, the hint explicitly directs the model's attention toward the Frobenius formula. This encourages a more rigorous, systematic approach: first, finding all possible integer pairs for (a, b) ; second, filtering these candidates by checking the essential coprimality condition $(\text{gcd}(a, b) = 1)$; and finally, verifying which of the remaining valid pairs satisfies the constraint that 1776 is non-representable. This structured reasoning process, prompted by the hint, is effective because it signals the specific theoretical framework needed to solve the problem, thereby preventing common errors and guiding the model to the correct solution.

1205

1206

C.5 ALGORITHM DETAILS

1207

Algorithm 1 HINT: Helping Ineffective rollouts Navigate Towards effectiveness

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

```

1: Input: initial policy model  $\pi_{\theta_{\text{init}}}$ ; reward models  $r_\phi$ ; task prompts  $\mathcal{D}$ ; hints  $\mathcal{H}$ ; hyperparameters  $\epsilon, \beta, \mu$ 
2: policy model  $\pi_\theta \leftarrow \pi_{\theta_{\text{init}}}$ 
3: for iteration = 1, ..., I do
4:   reference model  $\pi_{\text{ref}} \leftarrow \pi_\theta$ 
5:   for step = 1, ..., M do
6:     Sample a batch  $\mathcal{D}_b$  from  $\mathcal{D}$ 
7:     Update the old policy model  $\pi_{\theta_{\text{old}}} \leftarrow \pi_\theta$ 
8:     Sample  $G$  outputs  $\{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)$  for each  $q \in \mathcal{D}_b$ 
9:     Compute rewards  $\{r_{ij}\}_{i=1}^G$  for each  $o_i$  by running  $r_\phi$ 
10:    if all rewards  $\{r_{ij}\}$  are sparse (e.g., zero) then
11:      Get hint  $h \in \mathcal{H}$  for problem  $q$ 
12:      Construct hint-augmented query  $q_h$ 
13:      Resample  $G$  new outputs  $\{o_i^h\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q_h)$ 
14:      Compute new rewards  $\{r_{ij}^h\}_{i=1}^G$ 
15:      Let  $\{o_i\} \leftarrow \{o_i^h\}$ ,  $\{r_{ij}\} \leftarrow \{r_{ij}^h\}$ 
16:    end if
17:    Compute  $\hat{A}_{i,t}$  for each token  $t$  of  $o_i$  using final rewards
18:    for HINT iteration = 1, ...,  $\mu$  do
19:      Update  $\pi_\theta$  by maximizing GRPO objective
20:    end for
21:    Update  $r_\phi$  via replay training
22:  end for
23: end for
24: Output:  $\pi_\theta$ 

```

1238

D LLM USAGE

1239

1240

1241

Large Language Models (LLMs) were used to aid in the writing and polishing of the manuscript. Specifically, we used an LLM to assist in refining the language, improving readability, and ensuring clarity in various sections of the paper. The model helped with tasks such as sentence rephrasing, grammar checking, and enhancing the overall flow of the text.

1242 It is important to note that the LLM was not involved in the ideation, research methodology, or
1243 experimental design. All research concepts, ideas, and analyses were developed and conducted by
1244 the authors. The contributions of the LLM were solely focused on improving the linguistic quality
1245 of the paper, with no involvement in the scientific content or data analysis.

1246 The authors take full responsibility for the content of the manuscript, including any text generated
1247 or polished by the LLM. We have ensured that the LLM-generated text adheres to ethical guidelines
1248 and does not contribute to plagiarism or scientific misconduct.

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295