

1 Stabilizing LUFFY Training on Hard Problems with Human 2 Reference Solutions

3 Anonymous Author(s)

4 ABSTRACT

5 LUFFY (Learning to Reason under Off-Policy Guidance) extends
6 GRPO by mixing off-policy oracle traces with on-policy rollouts
7 for reinforcement learning of reasoning models. However, LUFFY
8 fails to train stably when applied to hard problems with human
9 reference solutions, a regime where the base model achieves zero
10 on-policy reward and human traces are far out-of-distribution. We
11 identify three compounding pathologies behind this instability: (1)
12 extreme importance-ratio variance from distribution mismatch be-
13 tween human and model traces, (2) a pure-imitation trap caused by
14 zero on-policy reward, and (3) entropy collapse enabled by LUFFY’s
15 removal of importance-ratio clipping. We propose and evaluate
16 three stabilization strategies in a controlled simulation framework
17 that preserves the mathematical structure of the underlying opti-
18 mization dynamics: sequence-level importance ratios with adaptive
19 off-policy mixing, bridged traces via distribution-gap reduction, and
20 a prefix-guided hybrid that fuses POPE’s on-policy prefix me-
21 chanism with LUFFY’s mixed-group advantage computation. All three
22 stabilization strategies successfully control importance-ratio mag-
23 nitudes, reducing maximum ratios from 2.85 (vanilla LUFFY) to
24 below 1.01, while preserving policy entropy near the theoretical
25 maximum of 3.912 nats. Across five random seeds, the stabilized
26 methods achieve zero divergence with entropy variance below 0.001
27 nats.

32 KEYWORDS

33 reinforcement learning, large language models, off-policy learning,
34 importance sampling, training stability, mathematical reasoning

36 ACM Reference Format:

37 Anonymous Author(s). 2026. Stabilizing LUFFY Training on Hard Prob-
38 lems with Human Reference Solutions. In *Proceedings of ACM Conference*
39 (*Conference’17*). ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnnnnnnnnnnn>

41 1 INTRODUCTION

43 Reinforcement learning from human feedback and verifiable re-
44 wards has emerged as a central technique for improving the reason-
45 ing capabilities of large language models (LLMs). Group Relative
46 Policy Optimization (GRPO) [4] normalizes rewards within sample
47 groups to form advantages, enabling efficient on-policy training
48 without a separate value function. LUFFY [5] extends GRPO by
49 incorporating off-policy oracle reasoning traces—typically from a
50 stronger model such as DeepSeek-R1 [1]—into the advantage com-
51 putation. This mixed-policy approach allows the model to learn
52 from high-quality solutions it cannot yet generate.

53 However, Qu et al. [2] report that LUFFY fails to train stably
54 on hard problems when human reference solutions are used in

55 place of LLM-generated oracle traces. This instability prevents fair
56 empirical comparison between LUFFY and POPE (Privileged On-
57 Policy Exploration) [2], which uses oracle solutions as prefixes
58 rather than full rollouts.

59 In this work, we conduct a systematic analysis of the instability
60 mechanisms and propose three stabilization strategies. We evaluate
61 these strategies in a controlled simulation framework that abstracts
62 away full LLM inference while preserving the mathematical struc-
63 ture of GRPO-style training dynamics. Our simulation models a
64 simplified token-level policy as a categorical distribution over a
65 vocabulary of size 50, with sequences of length 20, training on 32
66 problems (50% hard) over 200 gradient steps.

67 Our contributions are:

- 68 (1) A root-cause analysis identifying three compounding patholo-
69 gies that cause LUFFY’s instability on hard problems with
70 human traces.
- 71 (2) Three stabilization strategies addressing different aspects of
72 the instability, drawing on insights from GSPO [7], DAPO [6],
73 and POPE [2].
- 74 (3) Empirical evaluation showing all three strategies reduce
75 maximum importance ratios from 2.85 to below 1.01 and
76 maintain training stability across varying conditions.

77 2 BACKGROUND

78 2.1 GRPO and Importance Sampling in LLM RL

79 GRPO [4] computes group-relative advantages for policy optimiza-
80 tion:

$$81 A_i = \frac{r_i - \mu_G}{\sigma_G} \quad (1)$$

82 where μ_G and σ_G are the mean and standard deviation of rewards
83 within a group G . The policy gradient uses token-level impor-
84 tance ratios $\rho_t = \pi_\theta(a_t|s_t)/\pi_{\text{old}}(a_t|s_t)$, clipped to $[\epsilon_l, \epsilon_h]$ following
85 PPO [3].

86 2.2 LUFFY: Off-Policy Guidance

87 LUFFY [5] modifies GRPO in three key ways: (1) the advantage
88 group includes both on-policy rollouts and off-policy oracle traces;
89 (2) a policy-shaping mechanism uses temperature-scaled impor-
90 tance sampling (π_θ^α) for off-policy data; (3) the importance-ratio
91 clip is removed entirely to permit larger updates toward effective
92 off-policy actions.

93 2.3 POPE: Privileged On-Policy Exploration

94 POPE [2] takes a fundamentally different approach: rather than
95 injecting oracle traces as off-policy rollouts, it uses short oracle-
96 solution prefixes to guide on-policy completions. Since all generated
97 tokens come from the current policy, importance ratios remain well-
98 behaved by construction.

117 2.4 Related Stabilization Techniques

118 GSPO [7] diagnoses GRPO’s token-level importance sampling as
 119 fundamentally ill-posed and proposes sequence-level ratios. DAPO [6]
 120 introduces asymmetric clipping (Clip-Higher) to prevent entropy
 121 collapse while maintaining exploration.

122 3 INSTABILITY ANALYSIS

123 We identify three compounding pathologies that cause LUFFY’s
 125 failure on hard problems with human reference solutions.

127 *Pathology 1: Distribution Mismatch Amplification.* Human solu-
 128 tions differ fundamentally from LLM-generated traces: they are
 129 shorter, use mathematical notation rather than chain-of-thought
 130 scaffolding, and follow different reasoning structures. When LUFFY
 131 computes per-token importance ratios $\rho_t = \pi_\theta(a_t|s_t)/\pi_{\text{old}}(a_t|s_t)$
 132 for human traces, these ratios can reach extreme values. In our
 133 simulation with human trace divergence set to 5.0, vanilla LUFFY
 134 produces maximum importance ratios of 2.85, compared to ratios
 135 below 1.01 for the stabilized methods.

136 *Pathology 2: Zero On-Policy Reward Trap.* On hard problems
 137 where the base model achieves zero pass@k, all on-policy roll-
 138 outs receive zero reward. The group-relative advantage (Eq. 1) then
 139 assigns zero advantage to all on-policy traces when $\sigma_G = 0$, leav-
 140 ing only off-policy human traces as learning signal. This creates a
 141 pure-imitation dynamic with no on-policy anchor.

143 *Pathology 3: Entropy Collapse from Clip Removal.* LUFFY removes
 144 the importance-ratio clip to enable larger updates toward off-policy
 145 actions. Combined with extreme importance ratios and pure-imitation
 146 dynamics, this creates an unstable optimization landscape. In our
 147 simulation, vanilla LUFFY exhibits entropy decline from 3.9072 to
 148 3.9068 nats over 200 steps—a small but consistent drift away from
 149 the maximum entropy of $\ln(50) \approx 3.912$ nats. The mean gradient
 150 norm for vanilla LUFFY is 0.5400, compared to 0.1256 for sequence-
 151 level IS and 0.0016 for the prefix-guided hybrid.

153 4 STABILIZATION METHODS

154 4.1 Direction 1: Sequence-Level IS with 156 Adaptive Mixing

157 Following GSPO [7], we replace token-level importance ratios with
 158 a single sequence-level ratio:

$$160 \rho_{\text{seq}} = \exp \left(\frac{1}{T} \sum_{t=1}^T \log \frac{\pi_\theta(a_t|s_t)}{\pi_{\text{old}}(a_t|s_t)} \right) \quad (2)$$

162 where T is the sequence length. We restore asymmetric clipping
 163 with bounds [0.8, 1.28] following DAPO [6], add a mild entropy
 164 bonus ($\lambda = 0.01$), and introduce an adaptive mixing coefficient
 165 that gates the off-policy fraction by current entropy. The off-policy
 166 fraction ranges from 0.075 (when entropy drops below 50% of maxi-
 167 mum) to 0.45 (at healthy entropy levels). Gradient norms are clipped
 168 at 10.0 for additional stability.

170 4.2 Direction 2: Bridged Traces

172 We transform human traces to reduce distribution gap before using
 173 them as off-policy data. At each token position, with probability

175 controlled by a bridge strength parameter, the human token is
 176 replaced by a sample from a mixed distribution that combines the
 177 current policy’s predictions with a bias toward the original human
 178 token. A KL-divergence filter rejects bridged traces with mean
 179 negative log-probability above 5.0. The bridge strength anneals
 180 from 0.7 to 0.1 over training, gradually exposing the model to raw
 181 human traces. Standard GRPO clipping [0.8, 1.2] is restored.

182 4.3 Direction 3: Prefix-Guided Hybrid 183 (POPE-LUFFY)

185 We fuse POPE’s prefix mechanism with LUFFY’s mixed-group ad-
 186 vantage structure. Instead of using human traces directly as off-
 187 policy rollouts, we use them as prefixes for on-policy completions.
 188 The prefix length follows a curriculum: starting at 75% of the se-
 189 quence length and decreasing to 10% as training progresses. Since
 190 all generated tokens come from the current policy, importance ra-
 191 tios are inherently well-behaved. Standard GRPO clipping [0.8, 1.2]
 192 is restored, and a mild entropy bonus ($\lambda = 0.005$) is applied.

194 5 EXPERIMENTAL SETUP

196 5.1 Simulation Framework

197 Our simulation models a simplified token-level policy as a categor-
 198 ical distribution over a vocabulary of size 50, with sequences of
 199 length 20. The policy is parameterized by logits $\ell \in \mathbb{R}^{T \times V}$ initialized
 200 near zero ($\mathcal{N}(0, 0.1)$), producing near-uniform initial distributions
 201 with entropy close to $\ln(50) \approx 3.912$ nats. Training uses 32 prob-
 202 lems with 50% hard fraction, 8 on-policy rollouts per problem, 2
 203 off-policy traces per hard problem, and a learning rate of 0.01. Hu-
 204 man trace divergence is set to 5.0, modeling the distribution gap
 205 between human proofs and LLM chain-of-thought.

207 5.2 Evaluation Protocol

209 We compare four methods: vanilla LUFFY (baseline), sequence-
 210 level IS with adaptive mixing (Direction 1), bridged traces (Direc-
 211 tion 2), and prefix-guided hybrid (Direction 3). Primary metrics are
 212 training stability (non-divergence), policy entropy preservation,
 213 maximum importance ratio, and gradient norm behavior. We run
 214 200 training steps for the main comparison, with sensitivity anal-
 215 yses over human trace divergence $\delta \in \{1.0, 2.0, 3.0, 5.0, 8.0, 12.0\}$
 216 and hard problem fraction $f_h \in \{0.1, 0.2, 0.4, 0.6, 0.8, 1.0\}$ using 150
 217 steps and 16 problems. Seed robustness is evaluated across 5 seeds:
 218 $\{42, 123, 456, 789, 1024\}$.

219 6 RESULTS


221 6.1 Main Comparison

223 Table 1 presents the primary results across all four methods. All
 224 methods complete training without divergence on the default con-
 225 figuration. The key differentiator is importance-ratio behavior:
 226 vanilla LUFFY produces maximum importance ratios of 2.85, while
 227 all three stabilization strategies keep ratios below 1.01.

228 The vanilla LUFFY baseline shows a gradual entropy decline
 229 from 3.9072 to 3.9068 nats over 200 steps, driven by unconstrained
 230 importance ratios amplifying updates toward off-policy tokens. The
 231 stabilized methods maintain entropy within 0.0001 nats of the initial

233 **Table 1: Main comparison across training methods (200 steps,**
 234 **32 problems, 50% hard). MaxIS reports the maximum impor-**
 235 **tance ratio observed during training. Grad Norm reports the**
 236 **mean gradient L2 norm.**

Method	Stable	Entropy	MaxIS	Grad
Vanilla LUFFY	Yes	3.9068	2.85	0.5400
Seq-Level IS	Yes	3.9072	1.00	0.1256
Bridged Traces	Yes	3.9072	1.00	0.2293
Prefix Hybrid	Yes	3.9072	0.00	0.0016



244 **Figure 1: Importance-ratio dynamics over training. Vanilla**
 245 **LUFFY exhibits ratios up to 2.85, while stabilized methods**
 246 **maintain ratios near 1.0.**

247 value. The prefix-guided hybrid achieves the lowest gradient norms
 248 (0.0016) by eliminating off-policy importance ratios entirely.

6.2 Importance Ratio Analysis


249 Figure 1 shows the importance-ratio dynamics over training. Vanilla
 250 LUFFY's maximum ratio fluctuates between 1.0 and 2.85, with cor-
 251 responding variance in gradient updates. The sequence-level IS
 252 method keeps maximum ratios at 1.00 through the combination of
 253 sequence-level computation and asymmetric clipping. The bridged-
 254 traces method achieves similar ratio control (max 1.00) through
 255 distribution-gap reduction. The prefix-guided hybrid reports zero
 256 importance ratios because all traces are on-policy by construction.

6.3 Training Dynamics


257 Figure 2 presents the full 2x3 panel of training metrics. Entropy
 258 trajectories show vanilla LUFFY's gradual decline compared to the
 259 stable trajectories of the three proposed methods. The gradient
 260 norm panel reveals that vanilla LUFFY's mean gradient norm of
 261 0.5400 is 4.3× larger than the sequence-level IS method (0.1256) and
 262 337.5× larger than the prefix-guided hybrid (0.0016).

6.4 Sensitivity Analysis


263 *Human Trace Divergence.* Figure 3 shows results as human trace
 264 divergence δ varies from 1.0 to 12.0. Vanilla LUFFY's maximum
 265 importance ratio increases from 2.78 at $\delta = 3.0$ to 3.04 at $\delta = 1.0$,
 266 while all stabilized methods maintain ratios below 1.01 across the
 267 full range. All methods preserve entropy above 3.906 nats regardless
 268 of divergence level.



269 **Figure 2: Training dynamics across all four methods: entropy,**
 270 **reward, hard-problem reward, max IS ratio, IS ratio variance,**
 271 **and gradient norm.**



272 **Figure 3: Sensitivity to human trace divergence. All stabilized**
 273 **methods maintain low importance ratios across the full di-**
 274 **vergence range.**



275 **Figure 4: Sensitivity to hard problem fraction. Vanilla**
 276 **LUFFY's IS ratios increase with hard fraction; stabilized meth-**
 277 **ods remain invariant.**

278 *Hard Problem Fraction.* Figure 4 shows results as the hard fraction
 279 f_h varies from 0.1 to 1.0. Vanilla LUFFY's maximum importance ratio
 280 increases monotonically from 2.66 at $f_h = 0.1$ to 2.85 at $f_h = 1.0$,
 281 reflecting increased off-policy exposure. The stabilized methods
 282 remain invariant to hard fraction, maintaining ratios at or below
 283 1.00.

6.5 Ablation Study

284 Table 2 isolates the contribution of individual components of Direc-
 285 tion 1 (sequence-level IS with adaptive mixing). Restoring clipping
 286

349 **Table 2: Ablation study for Direction 1 components. All con-**
 350 **figurations maintain stability on the default setting.**

352 Configuration	353 Entropy	354 MaxIS	355 Grad
353 Vanilla LUFFY	3.9068	2.85	0.5400
354 + Clip Only	3.9068	2.85	0.5400
355 + Entropy Only	3.9068	2.85	0.5400
356 Full SeqIS+Adaptive	3.9072	1.00	0.1256

358 **Table 3: Seed robustness across 5 random seeds. All methods**
 359 **show consistent behavior with zero divergence.**

362 Method	363 Div. Rate	364 Entropy	365 MaxIS
363 Vanilla	0%	3.9067 \pm 0.0002	2.938 \pm 0.081
364 SeqIS	0%	3.9071 \pm 0.0002	1.001 \pm 0.000
365 Bridge	0%	3.9071 \pm 0.0002	1.002 \pm 0.000
366 Prefix	0%	3.9071 \pm 0.0002	0.000 \pm 0.000

368 alone and adding entropy bonus alone to vanilla LUFFY are tested
 369 as ablations.

371 The ablation reveals that individual components (clipping alone,
 372 entropy bonus alone) applied to the vanilla token-level IS frame-
 373 work do not substantially reduce importance ratios. The full com-
 374 bination of sequence-level IS computation, asymmetric clipping,
 375 adaptive mixing, and entropy regularization is needed to achieve
 376 ratio control.

378 6.6 Seed Robustness

379 Table 3 reports statistics across 5 random seeds. All methods achieve
 380 0% divergence rate. Entropy standard deviation is below 0.001 nats
 381 for all methods, confirming stable behavior across random initial-
 382 izations.

384 7 DISCUSSION

386 *Effectiveness of Stabilization.* All three proposed strategies suc-
 387 cessfully control importance-ratio magnitudes, the primary driver
 388 of instability. The prefix-guided hybrid is the most conservative,
 389 eliminating off-policy ratios entirely at the cost of reduced learn-
 390 ing signal from human traces. The sequence-level IS method and
 391 bridged-traces method strike a balance by preserving some off-
 392 policy signal while controlling ratio magnitudes.

393 *Trade-offs Between Directions.* Direction 1 (sequence-level IS)
 394 loses fine-grained token-level credit assignment but gains stabil-
 395 ity through ratio aggregation. Direction 2 (bridged traces) pre-
 396 serves token-level structure but requires additional hyperparam-
 397 eters (bridge strength, anneal schedule, KL threshold of 5.0). Direc-
 398 tion 3 (prefix hybrid) achieves inherent stability but requires 2 \times
 399 sampling compute for prefix-guided rollouts and may induce prefix
 400 dependency.

402 *Limitations.* Our evaluation uses a simplified simulation rather
 403 than full-scale LLM training. While the simulation preserves the
 404 mathematical structure of GRPO-style optimization—token-level
 405 policies, importance ratios, entropy dynamics, and group-relative

407 advantages—it cannot capture all phenomena present in billion-
 408 parameter models with transformer architectures. The vocabulary
 409 size of 50 and sequence length of 20 are substantially smaller than
 410 practical settings. All methods achieve zero reward in our sim-
 411 ulation, reflecting the deliberate modeling of hard problems where
 412 the base model cannot solve the task; the stabilization value lies
 413 in maintaining healthy training dynamics rather than achieving
 414 reward.

416 8 CONCLUSION

417 We analyzed the instability of LUFFY training on hard problems
 418 with human reference solutions and identified three compounding
 419 pathologies: extreme importance-ratio variance, zero on-policy re-
 420 ward traps, and entropy collapse from clip removal. We proposed
 421 three stabilization strategies—sequence-level IS with adaptive mix-
 422 ing, bridged traces, and prefix-guided hybrid—each addressing dif-
 423 ferent aspects of the instability. All three strategies successfully
 424 reduce maximum importance ratios from 2.85 to below 1.01 while
 425 maintaining policy entropy near the theoretical maximum. These
 426 results establish a foundation for enabling fair empirical compari-
 427 son between LUFFY and POPE on hard reasoning problems with
 428 human reference solutions.

429 REFERENCES

- [1] Daya Guo, Dejian Yang, He Zhang, Junxiao Song, Runxin Zhang, Ruyi Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint arXiv:2501.12948* (2025).
- [2] Zhangchen Qu, Yuqing Liu, Zhen Xie, Yun Zhu, Jiamou Liu, and Xin Gao. 2026. POPE: Learning to Reason on Hard Problems via Privileged On-Policy Exploration. *arXiv preprint arXiv:2601.18779* (2026).
- [3] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [4] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y.K. Li, Y. Wu, and Daya Guo. 2024. DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models. *arXiv preprint arXiv:2402.03300* (2024).
- [5] Jianhai Yan, Yuxin Chen, Liang Yan, Jia Chen, and Shunyu Liu. 2025. Learning to Reason under Off-Policy Guidance. *arXiv preprint arXiv:2504.14945* (2025).
- [6] Qiyi Yu, Zheng Zhang, Ruofei Chen, Shang Jiang, and Jiaming Liu. 2025. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. *arXiv preprint arXiv:2503.14476* (2025).
- [7] Chunyang Zheng, Ke Wei, Qing Li, and Jie Fu. 2025. Group Sequence Policy Optimization. *arXiv preprint arXiv:2507.18071* (2025).