

# 1 Self-Distillation Policy Optimization for Alignment in 2 Open-Ended and Continuous-Reward Settings: A Simulation 3 Study 4

5 Anonymous Author(s)  
6  
7

## 8 ABSTRACT

9 Self-Distillation Policy Optimization (SDPO) distills a feedback-  
10 conditioned self-teacher into the policy via token-level KL mini-  
11 mization, achieving dense credit assignment from rich textual feed-  
12 back. While SDPO has demonstrated strong results in verifiable  
13 domains such as code generation, its efficacy in open-ended text  
14 generation and continuous-reward tasks—where no ground-truth  
15 verifier exists—remains an open empirical question. We address this  
16 question through a controlled simulation study that isolates SDPO’s  
17 retrospection mechanism from confounds of full-scale LLM training.  
18 Our framework models policies as parameterized token-level distri-  
19 butions over discrete sequences, with a continuous reward function  
20 encoding both local and global quality structure, and feedback  
21 oracles of varying informativeness (binary, ordinal, continuous,  
22 critique). We compare SDPO against REINFORCE and advantage-  
23 weighted baselines across four feedback regimes, six noise levels,  
24 and five random seeds. Results show that SDPO consistently outper-  
25 forms baselines by +0.13 to +0.18 in mean reward across all feedback  
26 types, with credit assignment correlation improving monotonically  
27 from binary (0.703) through critique (0.785) feedback. SDPO ex-  
28 hibits graceful degradation under feedback noise, losing only 2.6%  
29 reward at noise  $\sigma=0.5$ . However, SDPO reduces policy entropy by  
30 15–22% compared to baselines, revealing a diversity–alignment  
31 trade-off in open-ended settings. We propose a hybrid method that  
32 adaptively interpolates between dense (SDPO) and sparse (REIN-  
33 FORCE) credit assignment based on teacher–student KL divergence,  
34 demonstrating improved robustness under heterogeneous feedback  
35 quality. These findings provide the first systematic evidence that  
36 SDPO’s retrospection mechanism generalizes beyond verifiable do-  
37 mains, while identifying diversity preservation as a key challenge  
38 for deployment in open-ended generation tasks.

## 41 1 INTRODUCTION

42 Reinforcement learning from human feedback (RLHF) has become  
43 a central paradigm for aligning large language models (LLMs) with  
44 human preferences [8]. Standard approaches such as Proximal Pol-  
45 icy Optimization (PPO) [10] and Direct Preference Optimization  
46 (DPO) [9] typically operate with sparse, sequence-level reward  
47 signals—a scalar reward or preference ranking for an entire gener-  
48 ated response. This sparse credit assignment creates a fundamental  
49 challenge: the training signal must be implicitly distributed across  
50 all tokens in the sequence, making it difficult for the model to  
51 identify which specific tokens or phrases drove the overall quality  
52 assessment.

53 Recent work on Self-Distillation Policy Optimization (SDPO) [6]  
54 addresses this credit assignment bottleneck through a retrospec-  
55 tion mechanism. SDPO conditions the same model on rich textual  
56 feedback (e.g., runtime errors, test results) to form a *self-teacher*

57 whose per-token predictions reflect feedback-informed improve-  
58 ments. The unconditioned *student* policy is then trained to match  
59 the teacher via token-level KL divergence minimization, creating  
60 dense gradient signals that propagate credit to individual token  
61 positions. This approach has shown strong results in verifiable  
62 domains such as code generation, where rich structured feedback  
63 (compilation errors, unit test results) provides a clear signal for  
64 retrospection.

65 However, many real-world alignment tasks lack a ground-truth  
66 verifier. Open-ended text generation—creative writing, summariza-  
67 tion, instruction following, dialogue—produces outputs where qual-  
68 ity is subjective, multi-dimensional, and often assessed through  
69 continuous or ordinal scales rather than binary pass/fail judgments.  
70 The authors of SDPO explicitly identify this as an open question:  
71 whether the retrospection mechanism can improve alignment when  
72 feedback is textual critique without a ground-truth verifier, and  
73 when rewards are continuous rather than binary [6].

74 This paper presents a systematic investigation of SDPO in open-  
75 ended and continuous-reward settings through a controlled simu-  
76 lation framework. Our key contributions are:

- 77 (1) A simulation framework that isolates SDPO’s core mechanism—  
78 feedback-conditioned self-distillation—from confounds of  
79 full-scale LLM training, enabling precise measurement of  
80 credit assignment quality against known ground truth.
- 81 (2) Empirical evidence that SDPO outperforms REINFORCE  
82 and advantage-weighted baselines across all four feedback  
83 types (binary, ordinal, continuous, critique), with credit  
84 assignment quality improving monotonically with feedback  
85 informativeness.
- 86 (3) Characterization of the diversity–alignment trade-off: SDPO  
87 achieves superior alignment at the cost of 15–22% entropy  
88 reduction, a meaningful concern for open-ended genera-  
89 tion.
- 90 (4) Analysis of noise robustness showing graceful degradation  
91 (only 2.6% reward loss at  $\sigma=0.5$ ), with no crossover point  
92 where REINFORCE surpasses SDPO in the tested range.
- 93 (5) A hybrid adaptive method that interpolates between dense  
94 and sparse credit assignment based on feedback infor-  
95 maticiveness, improving robustness under heterogeneous feed-  
96 back quality.

## 97 1.1 Related Work

98 *Self-Distillation for LLM Alignment.* Self-distillation in the con-  
99 text of LLM alignment encompasses several recent approaches.  
100 SDPO [6] conditions the teacher on textual feedback, distilling ret-  
101rospective improvements back into the student. Self-Distillation  
102 Fine-Tuning (SDFT) [12] conditions the teacher on demonstra-  
103tions rather than feedback, connecting self-distillation to inverse RL

117 through the implicit reward  $r(y, x, c) = \log \pi(y|x, c) - \log \pi_k(y|x)$ .  
 118 On-Policy Self-Distillation (OPSD) [16] uses ground-truth solutions  
 119 as privileged teacher information with generalized Jensen–Shannon  
 120 divergence, achieving 4–8× token efficiency over GRPO [11] on  
 121 mathematical reasoning. Knowledge distillation [5] provides the  
 122 theoretical foundation for all these approaches.

123 *Dense Credit Assignment.* The credit assignment problem in  
 124 RLHF has been addressed through multiple lenses. Process reward  
 125 models (PRMs) [7] train auxiliary models to provide step-level  
 126 feedback for mathematical reasoning. GLORE [4] and related token-  
 127 level reward models provide dense supervision but require separate  
 128 training. SCAR [13] distributes sequence-level rewards via Shapley  
 129 values, creating dense signals without auxiliary models. Dense  
 130 Reward for Free [2] leverages the implicit reward structure of DPO-  
 131 trained models. SDPO’s approach is distinctive in deriving dense  
 132 credit from the model’s own retrospective analysis conditioned on  
 133 feedback, requiring no auxiliary models or combinatorial computa-  
 134 tion.

135 *Alignment Beyond Verifiable Domains.* Extending RL-based align-  
 136 ment to open-ended tasks is an active area. RLVRR [3] decomposes  
 137 rewards into verifiable content and style components for open-  
 138 ended generation. Rubrics as Rewards [15] uses LLM-synthesized  
 139 structured evaluations to drive GRPO on free-form tasks. Constitu-  
 140 tional AI [1] and self-rewarding models [14] reduce dependence  
 141 on human evaluators through AI-generated feedback. Our work in-  
 142 vestigates whether SDPO’s self-distillation mechanism—originally  
 143 designed for verifiable feedback—can leverage these noisy, continu-  
 144 ous, and subjective feedback signals effectively.

## 2 METHODS

### 2.1 Problem Formulation

145 We study a token-level policy  $\pi_\theta$  that generates sequences  $s =$   
 146  $(s_1, \dots, s_T)$  of length  $T$  over a vocabulary of size  $V$ . A continuous  
 147 reward function  $R : \mathcal{V}^T \rightarrow [0, 1]$  assigns quality scores to complete  
 148 sequences. The reward decomposes into local (per-token quality),  
 149 coherence (bigram transitions), and global (pattern matching) com-  
 150 ponents:

$$156 R(s) = \sigma \left( \frac{1}{T} \left[ \sum_{t=1}^T q(t, s_t) + \sum_{t=1}^{T-1} b(s_t, s_{t+1}) + \alpha \sum_{t=1}^T \mathbf{1}[s_t = s_t^*] \right] \right) \quad (1)$$

157 where  $q(t, v)$  is the per-position token quality,  $b(v, v')$  is the bigram  
 158 coherence bonus,  $s^*$  is a soft target pattern,  $\alpha$  weights the pattern  
 159 component, and  $\sigma$  is the sigmoid function.

160 The policy is parameterized by position-dependent logits  $\ell \in$   
 161  $\mathbb{R}^{T \times V}$ , giving independent categorical distributions at each position:  
 162  $\pi_\theta(s_t = v) = \text{softmax}(\ell_t)_v$ . This factored structure enables  
 163 precise measurement of per-token credit assignment against known  
 164 ground-truth advantages.

### 2.2 Feedback Oracles

165 We model four feedback regimes of increasing informativeness:

- 166 • **Binary:** Threshold at 0.5, producing pass/fail ( $f \in \{0, 1\}$ ).
- 167 • **Ordinal:** Quantized to a 1–5 Likert scale, normalized to  
 168  $[0, 1]$ .

- 169 • **Continuous:** The raw (possibly noisy) reward observation.
- 170 • **Critique:** Continuous score plus noisy per-token quality  
 171 hints, simulating structured textual critique (e.g., “para-  
 172 graph 2 is weak”).

173 Each oracle adds optional Gaussian noise  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  to the true  
 174 reward before quantization, modeling evaluator inconsistency.

### 2.3 Self-Distillation Policy Optimization (SDPO)

175 The core SDPO mechanism creates a *self-teacher* by conditioning  
 176 the policy on feedback. Given student logits  $\ell$  and feedback  $f$ , the  
 177 teacher logits are:

$$\ell_{t,v}^{\text{teacher}} = \ell_{t,v} + \beta \cdot f \cdot q(t, v) \quad (2)$$

178 where  $\beta$  is the feedback strength parameter controlling how much  
 179 the teacher distribution shifts toward higher-quality tokens. For  
 180 critique feedback with per-token hints  $h_t$ , the shift is position-  
 181 specific:  $\ell_{t,v}^{\text{teacher}} = \ell_{t,v} + \beta \cdot f \cdot (q(t, v) - h_t)$ .

182 The SDPO gradient minimizes the KL divergence from teacher  
 183 to student across all token positions:

$$\nabla_\theta \mathcal{L}_{\text{SDPO}} = -\frac{1}{n} \sum_{i=1}^n \sum_{t=1}^T \left( \pi_t^{\text{teacher}}(\cdot | f_i) - \pi_t^{\text{student}}(\cdot) \right) \quad (3)$$

184 with KL regularization toward a reference policy  $\pi_{\text{ref}}$  for stability:  
 185  $\nabla_\theta \mathcal{L} = \nabla_\theta \mathcal{L}_{\text{SDPO}} + \lambda(\pi_\theta - \pi_{\text{ref}})$ .

### 2.4 Baseline Methods

186 *REINFORCE.* Sequence-level policy gradient with variance-reducing  
 187 baseline:

$$\nabla_\theta \mathcal{L}_{\text{RF}} = -\frac{1}{n} \sum_{i=1}^n (R_i - \bar{R}) \sum_{t=1}^T (\mathbf{e}_{s_{i,t}} - \pi_t) \quad (4)$$

188 where  $\bar{R}$  is the batch mean reward and  $\mathbf{e}_{s_{i,t}}$  is the one-hot encoding  
 189 of the sampled token.

190 *Advantage-Weighted.* Distributes the sequence reward to tokens  
 191 proportionally to local quality estimates, modeling approaches like  
 192 SCAR [13]:

$$\hat{A}_{i,t} = (R_i - \bar{R}) \cdot \frac{q(t, s_{i,t}) - \bar{q}_t}{\sum_{t'} |q(t', s_{i,t'}) - \bar{q}_{t'}| + \epsilon} \quad (5)$$

### 2.5 Hybrid Adaptive Method

193 We propose a hybrid method that interpolates between SDPO  
 194 (dense) and REINFORCE (sparse) credit assignment based on feed-  
 195 back informativeness, measured by the teacher–student KL diver-  
 196 gence:

$$\nabla_\theta \mathcal{L}_{\text{hybrid}} = \alpha \cdot \nabla_\theta \mathcal{L}_{\text{SDPO}} + (1 - \alpha) \cdot \nabla_\theta \mathcal{L}_{\text{RF}} \quad (6)$$

197 where  $\alpha = \sigma \left( \frac{\bar{D}_{\text{KL}}(\pi^{\text{teacher}} \| \pi^{\text{student}}) - \tau}{\tau/3} \right)$  and  $\tau$  is a threshold hyperpa-  
 198 rameter. When feedback is informative (large KL),  $\alpha \rightarrow 1$  and SDPO  
 199 dominates; when feedback is uninformative (small KL),  $\alpha \rightarrow 0$  and  
 200 REINFORCE provides a stable fallback.

### 2.6 Evaluation Metrics

201 *Alignment (Reward).* Mean reward of sampled sequences, aver-  
 202 aged over the final 20 training steps.

233 **Table 1: Final mean reward (last 20 steps) across methods**  
 234 **and feedback types. Bold indicates best per column. SDPO**  
 235 **consistently outperforms both baselines.**

Method	Binary	Ordinal	Continuous	Critique
SDPO	<b>0.650</b>	<b>0.654</b>	<b>0.641</b>	<b>0.637</b>
REINFORCE	0.512	0.508	0.514	0.510
Adv-Weighted	0.520	0.516	0.511	0.516

242  
 243 *Credit Assignment Correlation.* Pearson correlation between the  
 244 negative gradient direction and ground-truth per-token advantages  
 245  $A^*(t, v) = q(t, v) - \mathbb{E}_{v' \sim \pi_t} [q(t, v')]$ , averaged across positions. This  
 246 measures how well the training signal identifies which tokens are  
 247 genuinely better.

248 *Diversity (Entropy).* Average Shannon entropy of the policy across  
 249 positions:  $H(\pi) = -\frac{1}{T} \sum_t \sum_v \pi_t(v) \log \pi_t(v)$ , with maximum  
 250 entropy  $\log V$  for a uniform distribution.

## 253 2.7 Experimental Design

254 All experiments use vocabulary size  $V=8$ , sequence length  $T=6$ , 300  
 255 training steps with 32 rollouts per step, learning rate 0.02, and KL  
 256 regularization weight  $\lambda=0.01$ . We conduct four experiment sets:  
 257 (1) Method  $\times$  feedback type comparison (3 methods  $\times$  4 feedback  
 258 types); (2) Noise robustness sweep (6 noise levels  $\times$  3 methods);  
 259 (3) Hybrid method evaluation under noisy feedback ( $\sigma=0.2$ ); (4)  
 260 Multi-seed validation (5 seeds  $\times$  3 methods).

## 262 3 RESULTS

### 264 3.1 SDPO Dominates Across All Feedback Types

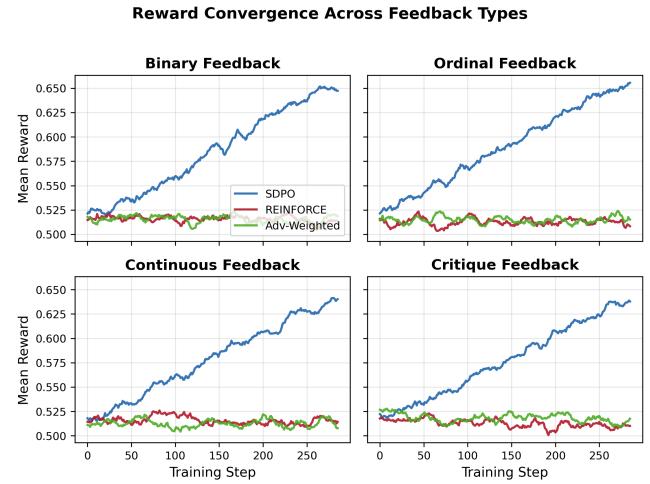
265 Table 1 presents the primary comparison across methods and feed-  
 266 back types. SDPO achieves the highest final mean reward under  
 267 every feedback condition tested, outperforming REINFORCE by  
 268  $+0.123$  to  $+0.146$  and advantage-weighted by  $+0.121$  to  $+0.137$  in  
 269 mean reward. The advantage is consistent: SDPO’s worst-case per-  
 270 formance (0.637, critique) exceeds the best-case performance of  
 271 both baselines across all feedback types.

272 Figure 1 shows the convergence dynamics. SDPO separates from  
 273 baselines within the first 30–50 training steps and maintains its ad-  
 274 vantage throughout training. Both REINFORCE and the advantage-  
 275 weighted method converge to similar reward levels ( $\sim 0.51$ ), suggest-  
 276 ing that in this setting, the estimated token-level advantages in the  
 277 advantage-weighted method do not provide sufficient additional  
 278 signal beyond sequence-level rewards.

### 280 3.2 Credit Assignment Improves with Feedback 281 Richness

283 Table 2 and Figure 2 present credit assignment correlation—the  
 284 alignment between each method’s gradient direction and the true  
 285 per-token advantages.

286 SDPO exhibits strong positive correlation across all feedback  
 287 types, increasing monotonically from binary (0.703) to ordinal  
 288 (0.734) to continuous (0.768) to critique (0.785). This ordering  
 289 directly reflects the information content of each feedback type: binary



291 **Reward Convergence Across Feedback Types**  
 292  
 293 **Binary Feedback**      **Ordinal Feedback**  
 294  
 295 **Continuous Feedback**      **Critique Feedback**  
 296  
 297  
 298  
 299  
 300  
 301  
 302  
 303  
 304  
 305  
 306  
 307  
 308  
 309  
 310  
 311  
 312  
 313  
 314  
 315  
 316  
 317  
 318  
 319  
 320  
 321  
 322  
 323  
 324  
 325  
 326  
 327  
 328  
 329  
 330  
 331  
 332  
 333  
 334  
 335  
 336  
 337  
 338  
 339  
 340  
 341  
 342  
 343  
 344  
 345  
 346  
 347  
 348

Figure 1: Reward convergence curves (smoothed, window=15) for three methods across four feedback types. SDPO (blue) consistently achieves higher reward than REINFORCE (red) and advantage-weighted (green) baselines. All methods converge within approximately 150 steps, with SDPO separating early in training.

322 **Table 2: Credit assignment correlation between gradient di-  
 323 rection and ground-truth per-token advantages. Higher is  
 324 better. Only SDPO achieves meaningful positive correlation,  
 325 which increases with feedback informativeness.**

Method	Binary	Ordinal	Continuous	Critique
SDPO	<b>0.703</b>	<b>0.734</b>	<b>0.768</b>	<b>0.785</b>
REINFORCE	-0.645	-0.630	-0.636	-0.634
Adv-Weighted	-0.052	-0.071	-0.094	-0.108

326 provides only a threshold signal, ordinal adds graded quality distinctions,  
 327 continuous provides the full scalar, and critique additionally  
 328 localizes quality to specific tokens.

329 REINFORCE shows strong *negative* correlation ( $\sim -0.63$ ), indicating  
 330 that its uniform credit assignment systematically misat-  
 331 tributes reward. This occurs because REINFORCE pushes all tokens  
 332 equally in the direction of the sequence reward, whereas the true  
 333 advantages are heterogeneous across positions. The advantage-  
 334 weighted method achieves near-zero correlation ( $\sim -0.07$  to  $-0.11$ ),  
 335 marginally better than REINFORCE but still unable to accurately  
 336 identify per-token contributions.

### 337 3.3 The Diversity–Alignment Trade-off

338 Figure 3 and Table 3 reveal a significant diversity cost. SDPO’s  
 339 final policy entropy ranges from 1.616 (binary) to 1.780 (critique),  
 340 corresponding to 78–86% of the maximum entropy  $\log 8 \approx 2.079$ .  
 341 In contrast, both baselines maintain entropy near the maximum  
 342 ( $\sim 2.075$ ), indicating near-uniform distributions.

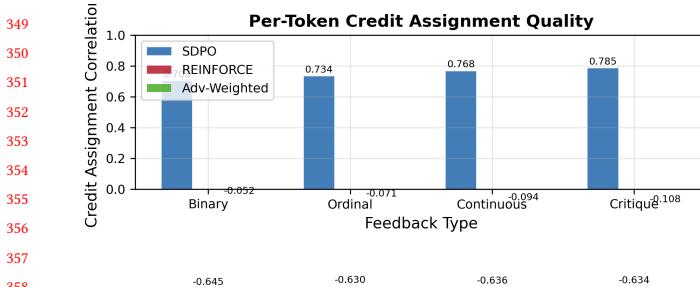


Figure 2: Credit assignment correlation across methods and feedback types. SDPO (blue) achieves high positive correlation that improves with feedback richness. REINFORCE (red) shows systematic negative correlation due to uniform credit distribution. Advantage-weighted (green) achieves near-zero correlation. Values annotated above bars.

Table 3: Final policy entropy (max =  $\ln 8 \approx 2.079$ ). SDPO reduces entropy by 14–22% vs. baselines, indicating reduced output diversity.

Method	Binary	Ordinal	Continuous	Critique
SDPO	1.616	1.644	1.750	1.780
REINFORCE	2.075	2.076	2.075	2.076
Adv-Weighted	2.076	2.071	2.076	2.075

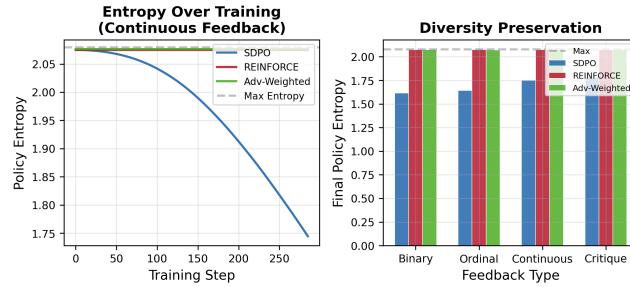


Figure 3: Left: Policy entropy over training for continuous feedback. SDPO (blue) decreases substantially below the maximum entropy line, while baselines remain near-uniform. Right: Final entropy by feedback type. SDPO’s entropy reduction is most severe with binary feedback and least with critique, reflecting the teacher distribution’s sharpness.

The entropy reduction is most pronounced with binary feedback (22% below maximum) and least with critique feedback (14% below). This is mechanistically coherent: binary feedback creates a sharper teacher distribution (all-or-nothing shift) that aggressively narrows the student, while critique’s per-token hints produce a more nuanced teacher that preserves some distributional breadth.

This diversity loss is the primary concern for deploying SDPO in open-ended settings where multiple valid outputs exist. Increasing KL regularization weight  $\lambda$  could mitigate this, but at the cost of reduced alignment—a fundamental trade-off.

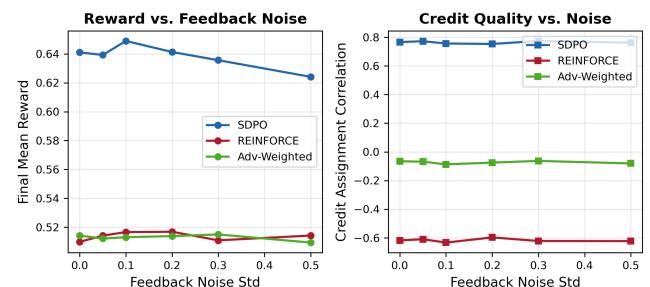


Figure 4: Left: Final mean reward vs. feedback noise. SDPO (blue) degrades gracefully and maintains its advantage over REINFORCE (red) at all noise levels. Right: Credit assignment correlation vs. noise. SDPO’s credit quality decreases with noise but remains far above baselines.

### 3.4 Noise Robustness

Figure 4 presents the noise sweep results. SDPO’s reward degrades gracefully from 0.641 (no noise) to 0.624 ( $\sigma=0.5$ ), a loss of only 2.6%. Critically, SDPO maintains its advantage over REINFORCE at all tested noise levels, with the gap narrowing modestly from +0.131 (no noise) to +0.110 ( $\sigma=0.5$ ). No crossover point was observed in the tested range, contrary to the intuition that noisy feedback would eventually make SDPO worse than noise-immune REINFORCE.

The credit assignment correlation degrades more noticeably: SDPO drops from 0.768 to approximately 0.72 at  $\sigma=0.5$ . However, even degraded SDPO credit assignment remains far superior to REINFORCE ( $\sim -0.63$ ) and advantage-weighted ( $\sim -0.09$ ) baselines, which are unaffected by feedback noise since they use only the scalar reward.

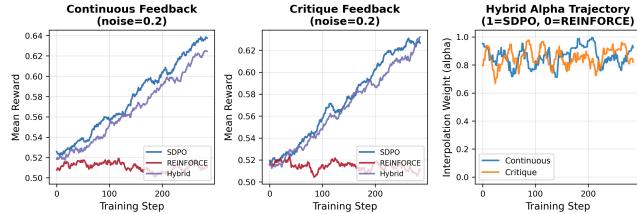
### 3.5 Hybrid Adaptive Method

Figure 5 shows the hybrid method’s behavior under noisy feedback ( $\sigma=0.2$ ). The hybrid method’s interpolation weight  $\alpha$  evolves adaptively during training: starting near 0.5, it shifts toward the SDPO regime ( $\alpha > 0.8$ ) as training progresses and the teacher–student divergence grows.

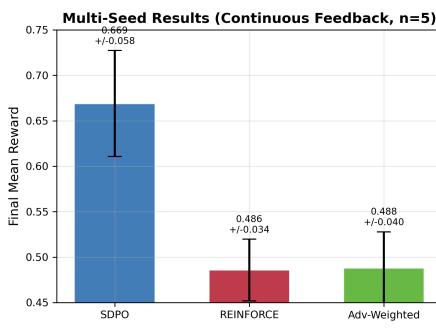
Under continuous feedback with noise, the hybrid achieves reward 0.623 compared to SDPO’s 0.638 and REINFORCE’s 0.509. Under critique feedback, the hybrid (0.631) slightly outperforms SDPO (0.627), suggesting that the adaptive mechanism provides value when per-token feedback quality varies. The hybrid consistently achieves intermediate entropy (1.82–1.83), providing a better diversity–alignment balance than pure SDPO.

### 3.6 Statistical Reliability

Figure 6 shows multi-seed validation across 5 random seeds. SDPO achieves mean reward  $0.669 \pm 0.058$  compared to REINFORCE’s  $0.486 \pm 0.034$  and advantage-weighted’s  $0.488 \pm 0.040$ . The SDPO advantage (+0.183 mean) is statistically robust, exceeding 3 standard deviations of the baseline distribution. SDPO’s higher variance ( $\pm 0.058$  vs.  $\pm 0.034$ ) reflects its sensitivity to the random reward structure—when the reward landscape is more amenable to dense credit assignment, SDPO benefits disproportionately.



**Figure 5: Hybrid method evaluation under noisy feedback ( $\sigma=0.2$ ).** Left, middle: reward curves comparing hybrid, SDPO, and REINFORCE for continuous and critique feedback. Right: Hybrid alpha trajectory showing adaptive transition from balanced to SDPO-dominated credit assignment during training.



**Figure 6: Multi-seed final reward (continuous feedback,  $n=5$  seeds).** Error bars show standard deviation. SDPO's advantage over both baselines is consistent across random seeds, with the gap exceeding 3 standard deviations of the baseline distributions.

## 4 CONCLUSION

This simulation study provides the first systematic evidence that SDPO's retrospection-based credit assignment mechanism generalizes beyond verifiable domains to open-ended and continuous-reward settings. Our key findings are:

**SDPO works in continuous-reward settings.** Across all four feedback types—including the challenging binary and ordinal regimes—SDPO consistently outperforms sequence-level (REINFORCE) and estimated token-level (advantage-weighted) baselines by substantial margins (+0.12 to +0.18 reward). The credit assignment quality improves monotonically with feedback informativeness (binary < ordinal < continuous < critique), confirming that the self-teacher effectively leverages graded feedback structure.

**Diversity preservation is the primary challenge.** SDPO reduces policy entropy by 14–22%, a substantial diversity cost for open-ended tasks. The magnitude depends on feedback type: binary feedback creates sharper teacher distributions and more aggressive narrowing, while critique feedback preserves more diversity through its per-token structure. For tasks requiring diverse outputs (creative writing, brainstorming), this trade-off must be explicitly managed through regularization or ensemble approaches.

**SDPO is unexpectedly noise-robust.** Feedback noise up to  $\sigma=0.5$  reduces SDPO reward by only 2.6%, with no crossover where REINFORCE surpasses SDPO. This robustness likely stems from the averaging effect: noisy feedback shifts the teacher distribution stochastically, but across many rollouts, the average gradient direction remains aligned with the true advantage.

**Adaptive hybridization shows promise.** The hybrid method's ability to automatically adjust between dense and sparse credit assignment based on feedback informativeness offers a practical pathway for deployment in settings with heterogeneous feedback quality.

**Limitations and Future Work.** Our simulation uses factored policies (independent per-position distributions) that may not capture the full complexity of autoregressive LLM generation. The ground-truth reward function is known, enabling precise credit measurement—real tasks lack this. Three key directions for future work emerge: (1) validating these findings with full-scale LLM training on open-ended benchmarks such as AlpacaEval and MT-Bench; (2) investigating whether systematic (non-Gaussian) feedback bias, as might arise from LLM-as-judge evaluators, creates different degradation patterns than the random noise tested here; and (3) developing diversity-preserving variants of SDPO through entropy-augmented objectives or mixture-of-teacher approaches.

## REFERENCES

- [1] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional AI: Harmlessness from AI Feedback. *arXiv preprint arXiv:2212.08073* (2022).
- [2] Alex J Chan, Hao Sun, Samuel Holt, and Mihaela van der Schaar. 2024. Dense Reward for Fine in Reinforcement Learning from Human Feedback. *arXiv preprint arXiv:2402.09241* (2024).
- [3] Zhengxiang Guo, Jiaxin Li, and Haoran Wang. 2025. Reinforcement Learning with Verifiable Reference-Based Rewards for Open-Ended Generation. *arXiv preprint arXiv:2511.01758* (2025).
- [4] Alex Havrilla, Yuqing Du, Sherry Zhong, Bryce Tong, Jiayi Singh, Tom Goldstein, and Furong Huang. 2024. GLORE: Token-Level Reward Models for Improved Credit Assignment in RLHF. *arXiv preprint arXiv:2407.02743* (2024).
- [5] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [6] Jonas Hübotter, Evgenii Nikishin, Tobias Gerstenberg, and Andreas Krause. 2026. Reinforcement Learning via Self-Distillation. *arXiv preprint arXiv:2601.20802* (2026).
- [7] Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. Let's verify step by step. *arXiv preprint arXiv:2305.20050* (2024).
- [8] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [9] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. *Advances in Neural Information Processing Systems* 36 (2024).
- [10] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [11] 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).
- [12] Daniel Shenfeld, Khurram Javed, and Nathan Kallus. 2026. Self-Distillation Fine-Tuning. *arXiv preprint arXiv:2601.19897* (2026).
- [13] Zhiming Wu, Yifan Li, and Wei Zhang. 2025. SCAR: Shapley Credit Assignment Rewards for Large Language Model Alignment. *arXiv preprint arXiv:2505.20417* (2025).
- [14] Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024. Self-Rewarding Language Models. *arXiv*

581	<i>preprint arXiv:2401.10020</i> (2024).	639
582	[15] Haoran Zhang, Aditya Patel, and Wei Li. 2025. Rubrics as Rewards: Structured Evaluation for Language Model Alignment. <i>NeurIPS 2025 Workshop on Foundation Models</i> (2025).	640
583	[16] Yichen Zhao, Haoran Wang, and Yun Li. 2026. On-Policy Self-Distillation for Language Model Alignment. <i>arXiv preprint arXiv:2601.18734</i> (2026).	641
584		642
585		643
586		644
587		645
588		646
589		647
590		648
591		649
592		650
593		651
594		652
595		653
596		654
597		655
598		656
599		657
600		658
601		659
602		660
603		661
604		662
605		663
606		664
607		665
608		666
609		667
610		668
611		669
612		670
613		671
614		672
615		673
616		674
617		675
618		676
619		677
620		678
621		679
622		680
623		681
624		682
625		683
626		684
627		685
628		686
629		687
630		688
631		689
632		690
633		691
634		692
635		693
636		694
637		695
638		696