

HARVE: Hacking-Aware Reward-Head Vector Editing for Robust Reward Models

Shuang Liu¹ Yuxuan Bo² Qiuyang Zhao³ Caiyue Huang⁴
Xiaorong Chen⁵ Yanguang Liu⁶ Mengnan Du⁷

¹Carnegie Mellon University ²University of Virginia

³Harvard University ⁴Stanford University ⁵University of Michigan

⁶New Jersey Institute of Technology

⁷The Chinese University of Hong Kong, Shenzhen

Abstract

Reward models are central to large language model (LLM) alignment, but they remain vulnerable to reward hacking. To evaluate reward-model robustness, we introduce REWARDHACKBENCH with 1,203 matched gold-hacked pairs. REWARDHACKBENCH spans 13 reward-hacking patterns covering real-world high-stakes domains and general settings, and we find severe failures on specific subcategories across eight reward models. To mitigate these failures, we propose HARVE, a backpropagation-free reward-head editing method for scalar reward models. Instead of fine-tuning the reward model, HARVE identifies a multi-directional hacking subspace from residual-stream directions associated with selected hacking subcategories, and removes the component of the reward-head vector aligned with that subspace. This directly reduces the reward head’s sensitivity to hacking-related features using only a small set of contrastive gold-hacked examples, without gradient updates or fine-tuning. Comprehensive experiments across eight reward models indicate that HARVE improves hacking robustness, outperforms fine-tuning baselines, and preserves general reward-model capability. Further analyses show that combining multiple target-specific residual-stream directions is more effective than relying on a single surface-level feature. The code is available at <https://github.com/olivialiu121/HARVE-Reward-Head-Editing>.

1 Introduction

Modern LLM alignment often relies on reward models to score and improve model outputs (Stiennon et al., 2020; Ouyang et al., 2022). However, reward models are vulnerable to *reward hacking*: models may exploit reward-correlated cues such as length or stylistic polish without improving factual correctness or reasoning quality (Singhal et al., 2024; Gao et al., 2023; Park et al., 2024; Liu et al.,

2025c). Existing benchmarks study generic reward-model hacking such as length bias and formatting bias (Singhal et al., 2024; Park et al., 2024; Sharma et al., 2024), with recent extensions to chat, code, math, and safety (Lambert et al., 2025; Liu et al., 2025c). However, reward hacking in complex high-stakes domains such as law and compliance in real-world settings remains underexplored.

To fill this gap, we introduce REWARDHACKBENCH, a benchmark with 1,203 matched gold-hacked pairs spanning 13 subcategories. REWARDHACKBENCH covers both professional-domain attacks and general-purpose attacks. Evaluating eight open-source scalar reward models on REWARDHACKBENCH, we find severe failures in specific subcategories, motivating mitigation methods that directly address reward hacking.

Current mitigation methods typically collect de-biased or adversarial preference data and fine-tune the reward model (Shen et al., 2024; Liu et al., 2025b; Chen et al., 2024). These approaches are data-dependent, computationally costly, and difficult to repeat for each newly discovered hacking pattern. To this end, we propose HARVE, a backpropagation-free reward-head editing method for scalar reward models. Instead of fine-tuning the reward model, HARVE identifies a multi-directional hacking subspace from residual-stream directions associated with selected hacking subcategories, and removes the component of the reward-head vector w_r aligned with that subspace. This reduces the reward head’s sensitivity to hacking-related features using only a small set of contrastive gold-hacked examples, without gradient updates or fine-tuning.

We evaluate HARVE on both REWARDHACKBENCH and RM-Bench (Liu et al., 2025c), comparing it with fine-tuning baselines. Across eight reward models, HARVE achieves the best target-subcategory performance, improving gold-preference rate by 22.4 percentage points over the

Table 1: **Baseline reward-model performance on the held-out test split of REWARDHACKBENCH.** Each cell reports the gold-preference rate, i.e., the percentage of pairs where the reward model assigns a higher score to the gold response than to the hacked response; higher is better. A1–C3 are professional-domain subcategories, and D1–E3 are general-purpose subcategories from LLMBar (Zeng et al., 2024) via RewardBench (Lambert et al., 2025). Full subcategory names are provided in Table 7. **Cell shading:** red $\leq 30\%$ severe failure; orange 30–50% below chance; yellow 50–70% mediocre; white $\geq 70\%$. Avg. is the sample-weighted micro-average over all 13 subcategories of the test split ($N = 967$).

Reward Model	Size	Professional-domain categories									General-purpose categories					Avg.
		A. Surface			B. Reasoning		C. Sycophantic			D. Off-Topic		E. Style				
		A1	A2	A3	B1	B2	C1	C2	C3	D1	D2	E1	E2	E3		
Skywork-Reward-V2-Qwen3	0.6B	33.87	86.96	46.27	58.57	88.57	94.29	82.86	71.43	58.21	43.48	63.83	89.00	58.70	67.53	
InternLM2-reward	1.8B	93.55	81.16	95.52	45.71	87.14	51.43	64.29	74.29	58.21	45.65	80.85	88.00	58.70	70.01	
GRM-Llama-3.2	3B	43.55	66.67	53.73	67.14	100.00	74.29	98.57	21.43	85.82	93.48	74.47	88.00	69.57	74.15	
Skywork-Reward-V2-Llama-3.2	3B	85.48	84.06	65.67	78.57	100.00	90.00	80.00	75.71	79.10	85.87	65.96	95.00	82.61	82.83	
RM-Mistral	7B	90.32	85.51	76.12	58.57	100.00	80.00	98.57	57.14	44.78	51.09	61.70	89.00	43.48	71.04	
Skywork-Reward-Llama-3.1	8B	85.48	60.87	71.64	68.57	100.00	98.57	95.71	62.86	90.30	95.65	80.85	91.00	73.91	84.07	
FsfairX-LLaMA3-RM	8B	83.87	84.06	76.12	57.14	100.00	90.00	92.86	40.00	50.75	57.61	65.96	91.00	52.17	71.77	
InternLM2-reward	20B	98.39	75.36	79.10	61.43	100.00	91.43	80.00	77.14	64.93	68.48	80.85	92.00	86.96	79.94	

original reward models and by 15.1 points over the stronger fine-tuning baseline on average. Non-target performance is largely preserved, with a 0.8-point average increase, while RM-Bench Hard improves by 5.3 points on average. We further conduct ablation experiments showing that the edited directions capture presentation patterns beyond literal length and that reward hacking is associated with multiple target-specific residual-stream directions that support effective reward-head editing.

Our contributions in this work are threefold. First, we propose HARVE, an efficient, interpretable, and backpropagation-free mitigation approach via reward-head editing. Second, we introduce REWARDHACKBENCH, covering reward hacking in both general-purpose and real-world high-stakes domains such as law and compliance. Third, we provide empirical and mechanistic evidence that reward hacking is associated with multiple target-specific residual-stream directions that support effective reward-head editing.

2 REWARDHACKBENCH

To evaluate reward-model robustness against reward hacking in both professional-domain and general-purpose settings, we introduce REWARDHACKBENCH, a benchmark of 1,203 matched gold–hacked response pairs across 13 subcategories. The benchmark covers professional-domain attacks in law, policy, and compliance as well as general-purpose attacks. For each query, the hacked response preserves the gold response’s

length, structure, and register while introducing one targeted hacking pattern. We split the professional-domain categories into stratified 20%/10%/70% train/dev/test partitions and reserve all general-purpose examples for testing, yielding a held-out test split of 967 pairs: 548 professional-domain pairs and 419 general-purpose pairs.

2.1 Benchmark Construction

REWARDHACKBENCH contains 13 subcategories grouped into five reward-hacking categories: (A) *Surface-Form Mimicry (Surface)*, where responses use fake or unsupported authority; (B) *Broken Reasoning (Reasoning)*, where responses omit required reasoning elements or invert key burdens; (C) *Sycophantic Hacking (Sycophantic)*, where responses over-align with the user or remove appropriate caveats; (D) *Off-Topic Hacking (Off-Topic)*, where responses answer a related but different prompt; and (E) *Style-Over-Substance (Style)*, where polished writing masks factual, logical, or task-specific flaws. Categories A–C target professional-domain settings such as law, policy, and compliance, while D–E cover general-purpose attacks. The full 13-subcategory taxonomy is provided in Appendix E.

For the professional-domain portion, we construct 784 matched gold–hacked pairs using a human–LLM pipeline. Two New York Bar-admitted attorneys draft or revise questions and gold responses using LLM-written examples and samples from existing legal benchmarks, such as




Figure 1: Overview of HARVE. Given a scalar reward model, we first identify training pairs where the RM prefers the hacked response over the gold response. We use these fooled pairs to estimate residual-stream directions associated with reward-hacking patterns. HARVE then edits the reward-head vector by removing its projection onto the subspace spanned by these hacking directions, leaving the transformer backbone unchanged.

LEGALBENCH, LEXAM, and STANFORD BAR EXAM QA (Guha et al., 2023; Fan et al., 2026; Zheng et al., 2025). Then we ask GPT-5-mini to rewrite each gold response into a hacked response that preserves the length and structure of the gold response while introducing one targeted failure mode. All pairs are then reviewed by three U.S. law students under attorney supervision. The final professional-domain subset has a mean hacked-to-gold word-count ratio of 1.01, reducing length and formatting confounds (Singhal et al., 2024; Park et al., 2024). For the general-purpose portion, we include 419 gold-hacked pairs from LLMBAR via the Chat Hard subset of REWARDBENCH (Zeng et al., 2024; Lambert et al., 2025).

2.2 Evaluation Protocol

We evaluate eight open-source scalar reward models spanning 0.6B–20B parameters from the Skywork, InternLM2, GRM, RM-Mistral, and FsfairX families (Liu et al., 2026a; Cai et al., 2024; Yang et al., 2024; Dong et al., 2023; Xiong et al., 2024; Dong et al., 2024). See Appendix F for detailed information. For each prompt x with gold response y_g and hacked response y_h , we report the *gold-preference rate*: the fraction of pairs where $R(x, y_g) > R(x, y_h)$. Higher values indicate stronger robustness to reward hacking; ties are counted as non-gold preferences.

2.3 Benchmark Results

Table 1 shows that open-source reward models remain vulnerable to reward hacking. Across eight scalar reward models, average gold-preference rates range from 67.53% to 84.07%, but every model has at least one mediocre-or-worse subcategory and several fall below chance on specific attacks. For instance, *C3 (Hedge Stripping)* yields severe failures, including 21.43% for GRM-Llama-3.2-3B and 40.00% for FsfairX-LLaMA3-RM-8B. Model scale alone does not solve the problem. For example, InternLM2-reward-20B underperforms several smaller reward models such as Skywork-Reward-V2-Llama-3.2-3B and Skywork-Reward-Llama-3.1-8B on most subcategories. Overall, REWARDHACKBENCH shows that reward hacking is heterogeneous and category-specific, motivating mitigation methods that directly target diverse hacking directions.

3 Mitigation Method

We propose HARVE, a backpropagation-free method for mitigating reward hacking in scalar reward models. HARVE estimates residual-stream directions associated with hacking patterns from gold-hacked examples, then edits the reward-head vector by removing its projection onto the corresponding subspace. This reduces sensitivity to hacking-related features without gradient updates, fine-tuning, or additional inference-time cost.

Algorithm 1 HARVE: Reward-Head Editing

Require: Reward head w_r ; training pairs $\mathcal{D}_k = \{(x_i, y_i^g, y_i^h)\}$ for target subcategories $k \in \mathcal{K}$; edit strength α

Ensure: Edited reward head w_r^{new}

```
1: for  $k \in \mathcal{K}$  do
2:   Identify fooled subset:
3:    $F_k \leftarrow \{i \in \{1, \dots, N_k\} \mid r(x_i, y_i^h) > r(x_i, y_i^g)\}$ 
4:   Extract hacking direction:
5:    $v_k \leftarrow \text{Normalize}(\sum_{i \in F_k} (h_i^h - h_i^g))$ 
6: end for
7: Stack subcategory directions:  $V_{\text{hack}} \leftarrow [v_{k_1}, \dots, v_{k_K}]$ 
8: Compute projection matrix via SVD:
9:  $P_{\text{hack}} \leftarrow UU^\top$  where  $U = \text{SVD}(V_{\text{hack}})$ 
10: Remove hacking components and return
11:  $w_r^{\text{new}} \leftarrow w_r - \alpha P_{\text{hack}} w_r$ 
```

3.1 Reward-Head Linearity

We focus on scalar reward models, which assign a single real-valued score to each prompt–response pair. Given final hidden state $h_{\text{final}}(x, y) \in \mathbb{R}^d$, the reward is computed by a linear reward head: $r(x, y) = w_r^\top h_{\text{final}}(x, y)$, where w_r is the reward-head vector. Because the score depends linearly on w_r , reducing sensitivity to a hacking-related direction can be formulated as a geometric edit to the reward head.

3.2 Estimating Hacking Directions

For each reward model, we select a small set of *target subcategories* whose hacking directions are estimated and removed. All remaining subcategories are treated as *non-target* categories and used to evaluate whether the edit preserves unrelated reward-model behavior.

For each target subcategory k , let $\mathcal{D}_k = \{(x_i, y_i^g, y_i^h)\}_{i=1}^{N_k}$ denote the matched training pairs, where y_i^g is the gold response and y_i^h is the hacked response. We first identify the fooled subset, consisting of examples where the base reward model prefers the hacked response: $F_k = \{i \in \{1, \dots, N_k\} : r(x_i, y_i^h) > r(x_i, y_i^g)\}$. For each $i \in F_k$, we extract the final hidden states $h_i^g = h_{\text{final}}(x_i, y_i^g)$ and $h_i^h = h_{\text{final}}(x_i, y_i^h)$. The hacking direction for subcategory k is the normalized mean hidden-state shift from gold to hacked responses

over the fooled subset:

$$\begin{aligned} \tilde{v}_k &= \frac{1}{|F_k|} \sum_{i \in F_k} (h_i^h - h_i^g), \\ v_k &= \frac{\tilde{v}_k}{\|\tilde{v}_k\|}. \end{aligned} \quad (1)$$

Thus, v_k captures the residual-stream direction associated with successful reward hacking in subcategory k . If F_k is empty, we omit that subcategory for the corresponding reward model.

3.3 Constructing the Hacking Subspace

Let $\mathcal{K} = \{k_1, \dots, k_K\}$ denote the target subcategories retained for a reward model. We stack their directions into the hacking-direction matrix $V_{\text{hack}} = [v_{k_1}, \dots, v_{k_K}] \in \mathbb{R}^{d \times K}$. The corresponding hacking subspace is the column space of this matrix, $\mathcal{S}_{\text{hack}} = \text{span}(V_{\text{hack}}) = \text{span}\{v_{k_1}, \dots, v_{k_K}\}$.

Because the directions v_k are not generally orthogonal and may overlap across hacking patterns, we compute an orthonormal basis for their span using singular value decomposition (SVD). Specifically, we decompose

$$V_{\text{hack}} = U \Sigma Q^\top, \quad (2)$$

where the columns of U are orthonormal directions in the residual-stream space. Let σ_j denote the singular values on the diagonal of Σ . We retain the columns of U with $\sigma_j > 10^{-6}$ and denote the resulting basis by U_{hack} . This threshold is used only as a numerical safeguard against degenerate or exactly collinear directions, not as an aggressive rank-reduction procedure. The projection matrix onto the hacking subspace is

$$P_{\text{hack}} = U_{\text{hack}} U_{\text{hack}}^\top. \quad (3)$$

3.4 Reward-Head Subspace Editing

Given the projection matrix P_{hack} , HARVE edits the reward head by subtracting the component of w_r aligned with the hacking subspace:

$$w_r^{\text{new}} = w_r - \alpha P_{\text{hack}} w_r, \quad (4)$$

where $\alpha \geq 0$ controls the intervention strength. When $\alpha = 0$, the reward head is unchanged; when $\alpha = 1$, the hacking-aligned component of w_r is removed. Values $\alpha > 1$ subtract more than this component, applying a stronger correction. We select α on the development split using validation trade-off curves. After editing, the reward model scores examples as $r_{\text{new}}(x, y) = (w_r^{\text{new}})^\top h_{\text{final}}(x, y)$.

Table 2: **Main evaluation of HARVE against fine-tuning baselines.** For each reward model, we compare the original reward model, 3:1 and 5:1 data-augmented fine-tuning baselines, and HARVE at the selected edit strength α^* . On REWARDHACKBENCH, **Tgt** and **Non-Tgt** report average gold-preference rates over the subcategories used for editing and the remaining subcategories, respectively; Surface, Reasoning, Sycophantic, Off-Topic, and Style denote the five top-level REWARDHACKBENCH categories. On RM-Bench Hard, **Avg** averages gold-preference rate over chat (Cht), code (Cde), math (Mth), safety-refuse (S-Rf), and safety-response (S-Rs). Bold values mark the best non-baseline method within each block, and blue rows highlight HARVE. $\uparrow$ Higher is better.

Model	REWARDHACKBENCH $\uparrow$							RM-Bench $\uparrow$					
	Overall		By category					Overall	By domain				
	Tgt	Non-Tgt	Surface	Reasoning	Sycophantic	Off-Topic	Style	Avg	Cht	Cde	Mth	S-Rf	S-Rs
Skywork-V2-Qwen3 (0.6B)													
+ Baseline	53.16	73.07	56.57	73.57	82.86	52.21	75.65	56.90	26.36	44.15	54.32	88.97	51.17
+ Finetune (3:1)	58.36	74.93	60.61	75.71	85.71	56.64	75.65	57.17	26.87	43.57	54.00	88.85	55.20
+ Finetune (5:1)	58.36	75.50	61.62	75.00	86.19	57.52	75.65	57.00	25.32	43.86	53.94	88.50	55.41
+ HARVE ($\alpha^*=1.50$, ours)	83.64	73.50	77.78	71.43	85.24	73.45	72.02	64.23	67.70	33.33	58.10	97.42	66.88
InternLM2-reward (1.8B)													
+ Baseline	58.93	74.24	89.90	66.43	63.33	52.21	79.27	59.56	24.29	25.15	69.69	92.61	44.59
+ Finetune (3:1)	61.07	74.82	90.40	70.00	65.24	52.65	78.76	59.33	24.03	24.56	69.44	92.61	44.59
+ Finetune (5:1)	62.50	75.25	90.91	68.57	67.14	53.98	79.27	58.96	23.51	23.83	68.75	92.84	44.80
+ HARVE ($\alpha^*=1.00$, ours)	78.57	75.40	87.88	77.86	83.81	54.87	80.31	59.58	21.96	29.68	69.57	90.96	43.52
GRM-Llama-3.2 (3B)													
+ Baseline	46.47	84.81	55.05	83.57	64.76	88.94	79.79	53.03	31.78	26.17	36.23	97.30	85.99
+ Finetune (3:1)	50.56	85.96	59.09	85.00	68.57	88.94	80.31	53.15	32.82	26.46	36.17	97.42	85.77
+ Finetune (5:1)	54.28	86.25	61.62	83.57	72.86	88.94	80.31	53.18	32.82	26.46	36.23	97.42	85.77
+ HARVE ($\alpha^*=1.00$, ours)	72.86	86.82	71.21	90.71	81.90	90.27	81.87	54.26	46.25	33.48	32.51	97.77	85.56
Skywork-V2-Llama-3.2 (3B)													
+ Baseline	75.09	85.51	78.28	89.29	81.90	81.42	83.94	69.43	55.04	51.90	62.51	97.54	79.19
+ Finetune (3:1)	76.17	86.23	79.29	89.29	84.29	81.42	84.46	69.40	55.30	52.19	62.26	97.54	79.19
+ Finetune (5:1)	76.90	86.09	80.30	89.29	84.29	81.42	83.94	69.33	55.04	51.90	62.26	97.54	79.19
+ HARVE ($\alpha^*=1.25$, ours)	92.42	86.52	86.36	93.57	96.67	82.74	83.42	70.91	64.34	50.58	62.76	98.94	82.59
RM-Mistral (7B)													
+ Baseline	67.87	72.32	83.84	79.29	78.57	47.35	71.50	44.41	16.54	25.29	30.88	88.62	60.72
+ Finetune (3:1)	84.12	71.30	94.44	81.43	93.81	45.13	64.77	45.74	19.12	26.46	30.75	93.31	60.08
+ Finetune (5:1)	84.48	78.55	94.95	82.14	93.33	57.52	76.17	47.88	24.55	26.61	33.14	95.19	62.00
+ HARVE ($\alpha^*=2.50$, ours)	90.25	76.52	88.38	92.14	87.62	65.49	73.58	65.06	53.75	49.27	52.93	94.25	85.35
FsfairX-LLaMA3-RM (8B)													
+ Baseline	57.49	75.66	81.31	78.57	74.29	53.54	75.65	50.46	20.67	37.13	39.19	91.90	57.32
+ Finetune (3:1)	67.63	77.89	85.86	80.71	83.81	55.75	76.17	50.94	21.71	37.13	39.63	92.96	57.11
+ Finetune (5:1)	71.50	80.13	89.90	82.14	85.71	58.41	78.76	52.47	24.55	39.33	40.77	95.19	56.69
+ HARVE ($\alpha^*=1.75$, ours)	86.96	76.58	80.81	94.29	86.19	62.39	76.68	54.38	31.27	43.86	40.58	96.13	59.66
Skywork-Reward-Llama-3.1 (8B)													
+ Baseline	65.94	91.17	72.22	84.29	85.71	92.04	84.46	54.68	32.56	32.75	36.80	98.36	85.99
+ Finetune (3:1)	72.46	91.75	76.77	85.71	91.43	91.59	84.46	54.71	33.33	32.60	36.61	98.36	86.41
+ Finetune (5:1)	73.55	92.19	77.78	85.71	92.38	91.59	85.49	54.79	33.59	32.60	36.67	98.36	86.62
+ HARVE ($\alpha^*=1.10$, ours)	84.42	89.58	80.30	93.57	96.19	91.15	79.79	61.22	50.90	45.18	43.04	98.59	86.62
InternLM2-reward (20B)													
+ Baseline	73.19	82.78	83.84	80.71	82.86	66.81	88.08	62.65	29.20	35.82	67.86	95.66	51.80
+ Finetune (3:1)	73.55	83.07	84.34	80.71	83.81	66.81	88.08	62.57	29.20	35.96	67.61	95.66	51.80
+ Finetune (5:1)	75.36	83.07	84.85	83.57	83.81	66.81	88.08	62.47	28.68	35.67	67.36	95.77	52.44
+ HARVE ($\alpha^*=1.60$, ours)	88.41	80.90	92.42	90.71	89.05	64.16	83.42	63.78	35.40	35.53	66.98	95.89	59.24

4 Experiments

4.1 Experimental Setups

Baselines. We compare HARVE with data-augmented fine-tuning baselines that fine-tune the reward model on a mixture of general preference data and pairs from the REWARDHACKBENCH training set. We consider two mixture ratios, 3:1 and 5:1, where the ratio denotes the number of general preference pairs per REWARDHACKBENCH pair. Fine-tuning details are provided in Appendix I.

Data splits. For the professional-domain categories, we use stratified 20%/10%/70% train/dev/test splits. The general-purpose categories are

reserved entirely for held-out testing. The extraction split is used to estimate HARVE’s hacking directions and to provide benchmark-specific data for the fine-tuning baselines. The REWARDHACKBENCH development split provides the target-performance objective for selecting the intervention strength α^* , while the non-overlapping RewardBench subsets provide only the preservation constraint. The test split is used for all reported results. We use the same subcategory-stratified split for every reward model.

HARVE vector extraction. For each reward model, we select three or four lowest-performing professional-domain subcategories as *target sub-*

categories using only the training split. We then estimate the hacking direction v_k from fooled examples for each selected target subcategory.

Benchmarks and Metrics. We evaluate the eight reward models on REWARDHACKBENCH and additionally use RM-Bench (Liu et al., 2025c) to measure cross-domain transfer across chat, code, math, and safety. RM-Bench contains three difficulty-level subsets, easy, normal, hard (Liu et al., 2025c), and we evaluate on the hard set. For each prompt x with gold response y_g and hacked response y_h , a scalar reward model R assigns scores $R(x, y_g)$ and $R(x, y_h)$. We report the *gold-preference rate*, defined as the fraction of pairs for which $R(x, y_g) > R(x, y_h)$.

4.2 Main Results

For each reward model, we compare HARVE with the two fine-tuning baselines, and report performance on targeted subcategories of REWARDHACKBENCH, non-target subcategories of REWARDHACKBENCH, and RM-Bench Hard task sets (Liu et al., 2025c). The editing strength α^* is selected to maximize *target subcategories* performance on the RewardHackBench development split subject to a RewardBench preservation constraint: performance may decrease by at most 4 percentage points relative to the original reward model. To avoid overlap with RewardHackBench, we exclude the five RewardBench subsets reused in RewardHackBench when applying this constraint. RM-Bench is not used for model selection and remains fully held out. See Table 2 for the full experiment results.

HARVE substantially improves robustness in target subcategories. HARVE achieves the best target-subcategory performance for all eight reward models. On average, it improves the gold-preference rate by 22.4 percentage points over the original reward models and by 15.1 points over the stronger fine-tuning baseline. The gains are consistent across model families and sizes, ranging from 15.2 points for InternLM2-reward-20B to 30.5 points for Skywork-V2-Qwen3-0.6B. HARVE improves performance on non-target subcategories over the original reward model for six of eight reward models. These results suggest that reward-head vector editing can remove hacking-sensitive directions more effectively than fine-tuning, without broadly degrading unrelated categories.

HARVE transfers to cross-domain reward-hacking settings. HARVE improves RM-Bench Hard performance for all eight reward models, with an average gain of 5.3 percentage points over the original models. The largest gain occurs for RM-Mistral-7B, where RM-Bench Hard improves by 20.7 points over the original model and by 17.2 points over the stronger fine-tuning baseline. These results suggest that the edit can transfer beyond the target categories used to estimate the hacking directions.

Overall, HARVE consistently improves robustness on targeted hacking categories while preserving general reward-model capability. Compared with data-augmented fine-tuning, the backpropagation-free vector-editing approach provides a stronger and more controllable robustness-capability trade-off.

4.3 Comparison with Robustness-Oriented Training

We further compare HARVE with Robust Reward Model Training (RRM) (Liu et al., 2025b), a training-based approach for mitigating reward hacking. We evaluate four reward models under two training-data ratios, 3:1 and 5:1. Within each ratio, standard fine-tuning, RRM, and RRM+HARVE use the same training pool. The original reward model and HARVE applied directly to the base model require no fine-tuning data and are included as references.

As shown in Appendix N, RRM substantially improves target performance over standard fine-tuning across all four models. Applying HARVE after RRM provides additional target gains for three of four models under the 3:1 setting. The methods also show different trade-offs: RRM often better preserves non-target performance, whereas RRM+HARVE can further improve target robustness and RM-Bench performance. Overall, the results suggest that HARVE can complement RRM when hacking-sensitive directions remain after training.

4.4 Proxy Downstream Evaluation

We further test whether the improvements on REWARDHACKBENCH carry over to response selection. We consider two proxy settings using four reward models: constructed-pool reranking and best-of- N selection.

Constructed-pool reranking. For each prompt, the reward model selects between the gold and hacked response. We report the *hacked-selection rate*, defined as the percentage of prompts for which the hacked response is ranked above the gold response; lower is better. As shown in Table 16, applying HARVE to the original reward models reduces the hacked-selection rate for all four reward models. Applying HARVE after RRM further reduces the rate. These results suggest that HARVE can complement RRM when hacking-sensitive directions remain after training.

Best-of- N selection. We next increase search pressure by selecting the highest-scoring response from a fixed pool of N candidates. We measure the gap between the reward assigned to the selected response and the gold reference. As shown in Table 18, the proxy-gold gap grows with N for the original reward models and for RRM, but grows more slowly after applying HARVE. For example, at $N = 256$, the gap changes from 9.62 to 3.28 for GRM-Llama-3.2 (3B), and from 9.26 to 1.84 for Skywork-V2-Llama-3.2 (3B) after applying HARVE.

The proxy-gold gap uses each reward model’s own score scale, so it should not be interpreted as an independent measure of response quality. We therefore also evaluate the selected responses directly at $N = 256$. As shown in Table 17, targeted-hack rates differ only modestly across methods, while judged response quality remains broadly similar. Together, these results show that HARVE reduces the growth of its own reward signal under stronger search, while judged response quality remains broadly similar across methods.

5 Mechanistic Insights

5.1 Diagnostic Style-Category Decomposition

We further use the extracted hacking directions to diagnose whether each direction mainly reflects general presentation-related features, such as length, polish, and confidence, or also contains category-specific hacking signals. For each reward model, we estimate a unit *style direction* v_{sty} from non-target gold responses by contrasting the final hidden states of responses in the top and bottom token-count quartiles, which captures presentation features such as length, elaboration, vocabulary complexity, and confidence. We then decompose each hacking direction v_k into a style-aligned component, which captures overlap with shared pre-

Table 3: **Relationship between hacking directions, presentation/style direction, and token length.** For each reward model, we report the mean cosine similarity between the extracted hacking direction v_k and the presentation/style direction v_{sty} , the mean hacked-minus-gold token-count difference Δ_{tok} , and the Pearson correlation between v_{sty} projection and token-count difference.

RM	Size	$\cos(v_k, v_{\text{sty}})$	Δ_{tok}	Pearson ρ
Skywork-V2-Qwen3	0.6B	+0.51	+20.24	+0.39
InternLM2-reward	1.8B	+0.65	−5.34	+0.03
GRM-Llama-3.2	3B	+0.46	+14.82	+0.30
Skywork-V2-Llama-3.2	3B	+0.44	+3.20	+0.16
RM-Mistral	7B	+0.20	+8.61	+0.15
FsfairX-LLaMA3-RM	8B	+0.27	+7.86	+0.19
Skywork-Reward-Llama-3.1	8B	+0.52	+8.32	+0.30
InternLM2-reward	20B	+0.17	+8.95	+0.06
Mean across RMs		+0.40	+8.33	+0.20

sentation features, and a style-orthogonal residual, which captures the remaining category-specific signal.

$$v_k = \underbrace{(v_k^\top v_{\text{sty}}) v_{\text{sty}}}_{\text{style-aligned}} + \underbrace{v_{k, \perp \text{sty}}}_{\text{style-orthogonal residual}}, \quad (5)$$

where $v_{k, \perp \text{sty}}$ is orthogonal to v_{sty} . We run three diagnostic experiments with the decomposition: removing only the shared style direction v_{sty} , removing only the style-orthogonal residual directions $v_{k, \perp \text{sty}}$, and removing the full hacking directions.

5.2 Results and Insights

The style direction captures style and elaboration, not literal length.

We audit the relationship between extracted hacking directions, the reward head, and the style direction using cosine similarity. Table 3 shows that the extracted directions are substantially entangled with the style direction, with mean $\cos(\hat{v}_k, v_{\text{sty}}) = +0.40$. However, this style component is not simply token length, as the mean correlation between v_{sty} projection and token-count difference is only +0.20, which suggests that HARVE is not merely removing a length feature, but a broader stylistic cluster that reward models associate with hacked responses.

Joint editing is super-additive. Table 4 shows that the full hacking direction is more effective than the decomposed style or style-orthogonal hacking category-related components. On target subcategories, ablating only the style direction or only the style-orthogonal residual gives limited gains, while ablating the raw joint direction improves

Table 4: **Component ablation analysis of hacking directions.** For each reward model, we report the baseline gold-preference rate (Base, GP %) and the change in gold-preference rate after ablating three directions: the shared presentation/style direction v_{sty} , the style-orthogonal category-specific residual $v_{k,\perp\text{sty}}$, and the full hacking direction v_k (Joint). Results are reported separately on target and non-target subcategories. All ablation columns are changes relative to the baseline in percentage points (pp; $\uparrow$ means higher is better).

RM	Size	Target subcategories				Non-target subcategories			
		Base (GP %)	v_{sty} (pp $\uparrow$)	$v_{k,\perp\text{sty}}$ (pp $\uparrow$)	Joint (pp $\uparrow$)	Base (GP %)	v_{sty} (pp $\uparrow$)	$v_{k,\perp\text{sty}}$ (pp $\uparrow$)	Joint (pp $\uparrow$)
Skywork-V2-Qwen3	0.6B	52.16	-3.99	+5.95	+27.89	87.82	+2.85	-0.01	+1.76
InternLM2-reward	1.8B	58.93	-15.36	+3.93	+19.64	89.66	-11.16	+1.50	+1.45
GRM-Llama-3.2	3B	45.70	+8.58	+2.80	+27.26	85.24	+3.96	+2.87	+3.95
Skywork-V2-Llama-3.2	3B	73.89	+1.09	+6.93	+10.91	89.88	-0.01	+1.53	+3.00
RM-Mistral	7B	67.97	+3.52	+5.14	+13.00	93.96	+0.78	-0.04	+2.58
FsfairX-LLaMA3-RM	8B	58.21	-2.05	+7.66	+13.77	89.26	+3.48	+0.54	+2.08
Skywork-Reward-Llama-3.1	8B	65.98	+1.36	+1.80	+12.24	94.94	+0.81	-0.81	+0.67
InternLM2-reward	20B	73.26	+0.36	+8.34	+8.68	92.45	+0.36	+0.36	+1.07
Mean		62.01	-0.81	+5.32	+16.67	90.40	+0.13	+0.74	+2.07

gold-preference rate by +16.67 percentage points on average. This indicates that reward hacking is not captured by an isolated style feature or an isolated category-specific feature. Instead, the results suggest that multiple residual-stream directions jointly contribute to the reward head’s sensitivity to hacking-related features. On non-target subcategories, reward models’ performance is already high, and joint ablation changes gold-preference rate by only +2.1 points on average. This supports the interpretation that HARVE primarily removes target-relevant hacking sensitivity without globally degrading the reward model.

6 Ablation Studies

6.1 Fooled-Subset Filtering

For each reward model, we compare the cosine alignment between the reward-head direction w_r and the hacked-minus-gold hidden-state difference estimated in two ways: using all training pairs, and using only fooled pairs where the base reward model assigns a higher score to the hacked response than to the gold response. As shown in Table 5, fooled-subset filtering increases the cosine alignment from -0.05 with all pairs to +0.24. This indicates that fooled examples better isolate the hidden-state direction that the reward model actually rewards when it prefers hacked responses.

6.2 Resampling Stability

To test sensitivity to the examples used for direction estimation, we repeat HARVE over 20 resampled

Table 5: **Effect of fooled-subset filtering on hacking-direction estimation.** ALL reports the cosine alignment between the reward-head direction w_r and the hacked-minus-gold hidden-state difference when the direction is estimated from all training pairs; FOOLED reports the corresponding cosine alignment when the direction is estimated only from fooled pairs. Values are averaged over each reward model’s target subcategories. $\Delta = \text{FOOLED} - \text{ALL}$.

RM	Size	All	Foiled	Δ
Skywork-V2-Qwen3	0.6B	+0.06	+0.17	+0.12
InternLM2-reward	1.8B	+0.12	+0.52	+0.39
GRM-Llama-3.2	3B	+0.03	+0.23	+0.20
Skywork-V2-Llama-3.2	3B	-0.21	+0.20	+0.41
RM-Mistral	7B	-0.05	+0.10	+0.14
FsfairX-LLaMA3-RM	8B	0.00	+0.11	+0.11
Skywork-Reward-Llama-3.1	8B	-0.09	+0.29	+0.38
InternLM2-reward	20B	-0.28	+0.25	+0.53
Mean		-0.05	+0.24	+0.29

training sets of the same size. For each run, we re-identify the fooled subset and re-estimate the hacking directions while keeping the development-selected α^* fixed. As shown in Appendix K, the target gains remain stable across all eight reward models. On average, the resampled runs improve target performance by 18.4 percentage points over the original RMs, with standard deviations ranging from 1.6 to 3.7 percentage points. Non-target performance changes by 0.7 percentage points on average, while RM-Bench Hard improves by 3.3 percentage points. These results show that HARVE remains effective across different samples used for direction estimation.




Figure 2: **Trade-off curves over intervention strength α for six representative reward models.** Blue curves show REWARDHACKBENCH gold-preference rate (GP rate), and orange curves show RM-Bench Hard accuracy. Dashed vertical lines mark the selected α^* for each RM, and dotted horizontal lines show the 5:1 fine-tuning baseline. The full eight-model figure is provided in Appendix M.

6.3 Intervention Strength α

We further conduct ablation studies to examine how the editing strength α controls the robustness-capability trade-off. For each reward model, we sweep α from 0 to 3 and track performance on both REWARDHACKBENCH and RM-Bench. As shown in Figure 2, increasing α generally improves REWARDHACKBENCH performance at first, but overly large values can eventually degrade performance on either REWARDHACKBENCH or RM-Bench. This motivates selecting a per-model sweet spot α^* rather than using a fixed global intervention strength. In our main experiments, α^* is selected using the development protocol described in Section 4.1.

7 Related Work

Existing reward hacking mitigation methods include data-centric augmentation or rebalancing (Shen et al., 2024; Liu et al., 2025b), training-based changes to reward-model parameters (Chen et al., 2024; Liu et al., 2025b), and optimization-time methods such as reward ensembles (Coste et al., 2024; Eisenstein et al., 2024). These approaches can reduce reward hacking, but often require additional curated data, costly reward-model updates, or inference-time aggregation. In contrast, HARVE edits only the reward-head vector, requiring no gradient updates and leaving the transformer backbone unchanged. Additional related work is provided in

Appendix A.

8 Conclusion

We introduce REWARDHACKBENCH, a benchmark covering professional-domain and general-purpose reward-hacking patterns, and show that eight open-source scalar reward models remain vulnerable across model sizes and families. Motivated by these failures, we present HARVE, a backpropagation-free reward-head editing method for mitigating reward hacking. Across eight reward models, HARVE consistently outperforms fine-tuning baselines on target hacking subcategories and largely preserves broader reward-model capability, including performance on RM-Bench. Further analyses show that removing the full multi-directional hacking subspace is more effective than removing its style-related or category-specific components separately, suggesting that these failures arise from jointly interacting residual-stream directions.

Limitations

This work has several limitations. First, HARVE applies to scalar reward models with accessible reward-head vectors and hidden states, and does not directly cover closed-source reward models or generative pairwise evaluators. Second, although REWARDHACKBENCH includes both professional-domain and general-purpose reward-hacking patterns, it cannot exhaust all real-world failure modes, especially in high-stakes domains where errors may be more context-dependent. Third, HARVE requires a small set of contrastive gold-hacked examples and selects intervention strength using validation trade-off curves, so its effectiveness may depend on the coverage of these examples. Finally, our downstream analyses are limited to constructed-pool reranking and best-of- N proxies. We do not evaluate the edited reward models inside full RLHF training.

References

- Anthropic. 2026. Introducing claude sonnet 5. <https://www.anthropic.com/news/claude-sonnet-5>.
- Andy Ardit, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. 2024. [Refusal in language models is mediated by a single direction](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 136037–136083. Curran Associates, Inc.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, and 81 others. 2024. [Internlm2 technical report](#). *Preprint*, arXiv:2403.17297.
- Lichang Chen, Chen Zhu, Jiuhai Chen, Davit Soselia, Tianyi Zhou, Tom Goldstein, Heng Huang, Mohammad Shoeybi, and Bryan Catanzaro. 2024. [ODIN: Disentangled reward mitigates hacking in RLHF](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 7935–7952. PMLR.
- Thomas Coste, Usman Anwar, Robert Kirk, and David Krueger. 2024. [Reward model ensembles help mitigate overoptimization](#). In *International Conference on Learning Representations*, volume 2024, pages 50905–50931.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Yihan Zhang, Winnie Chow, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#). *Preprint*, arXiv:2304.06767.
- Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. 2024. [Rlhf workflow: From reward modeling to online rlhf](#). *Preprint*, arXiv:2405.07863.
- Jacob Eisenstein, Chirag Nagpal, Alekh Agarwal, Ahmad Beirami, Alex D’Amour, DJ Dvijotham, Adam Fisch, Katherine Heller, Stephen Pfohl, Deepak Ramachandran, Peter Shaw, and Jonathan Berant. 2024. [Helping or herding? reward model ensembles mitigate but do not eliminate reward hacking](#). *Preprint*, arXiv:2312.09244.
- Yu Fan, Jingwei Ni, Jakob Merane, Yang Tian, Yoan Hermstrüwer, Yinya Huang, Mubashara Akhtar, Etienne Salimbeni, Florian Geering, Oliver Dreyer, Daniel Brunner, Markus Leippold, Mrinmaya Sachan, Alexander Stremitzer, Christoph Engel, Elliott Ash, and Joel Niklaus. 2026. [Lexam: Benchmarking legal reasoning on 340 law exams](#). *Preprint*, arXiv:2505.12864.
- Evan Frick, Tianle Li, Connor Chen, Wei-Lin Chiang, Anastasios Angelopoulos, Jiantao Jiao, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2025. [How to evaluate reward models for rlhf](#). In *International Conference on Learning Representations*, volume 2025, pages 18128–18163.
- Leo Gao, John Schulman, and Jacob Hilton. 2023. [Scaling laws for reward model overoptimization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10835–10866. PMLR.
- Neel Guha, Julian Nyarko, Daniel Ho, Christopher Ré, Adam Chilton, Aditya K, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory Dickinson, Haggai Porat, Jason Hegland, and 21 others. 2023. [Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 44123–44279. Curran Associates, Inc.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. [RewardBench: Evaluating reward models for language modeling](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1755–1797, Albuquerque, New Mexico. Association for Computational Linguistics.
- Chris Yuhao Liu, Liang Zeng, Yuzhen Xiao, Jujie He, Jiakai Liu, Chaojie Wang, Rui Yan, Wei Shen, Fuxiang Zhang, Jiacheng Xu, Yang Liu, and Yahui Zhou. 2026a. [Skywork-reward-v2: Scaling preference data curation via human-ai synergy](#). *Preprint*, arXiv:2507.01352.

- Shuang Liu, Zelong Li, Ruoyun Ma, Haiyan Zhao, and Mengnan Du. 2025a. [Contracteval: Benchmarking llms for clause-level legal risk identification in commercial contracts](#). *Preprint*, arXiv:2508.03080.
- Shuang Liu, Ruijia Zhang, Ruoyun Ma, Yujia Deng, Lanyi Zhu, Jiayu Li, Zelong Li, Zhibin Shen, and Mengnan Du. 2026b. [Llm agents in law: Taxonomy, applications, and challenges](#). *Preprint*, arXiv:2601.06216.
- Tianqi Liu, Wei Xiong, Jie Ren, Lichang Chen, Junru Wu, Rishabh Joshi, Yang Gao, Jiaming Shen, Zhen Qin, Tianhe Yu, Daniel Sohn, Anastasia Makarova, Jeremiah Zhe Liu, Yuan Liu, Bilal Piot, Abe Ittycheriah, Aviral Kumar, and Mohammad Saleh. 2025b. [Rrm: Robust reward model training mitigates reward hacking](#). In *International Conference on Learning Representations*, volume 2025, pages 62682–62700.
- Yantao Liu, Zijun Yao, Rui Min, Yixin Cao, Lei Hou, and Juanzi Li. 2025c. [Rm-bench: Benchmarking reward models of language models with subtlety and style](#). In *International Conference on Learning Representations*, volume 2025, pages 44323–44355.
- Mohammed Suhail B Nadaf. 2026. [reward-lens: A mechanistic interpretability library for reward models](#). *Preprint*, arXiv:2604.26130.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Junsoo Park, Seungyeon Jwa, Ren Meiying, Daeyoung Kim, and Sanghyuk Choi. 2024. [OffsetBias: Leveraging debiased data for tuning evaluators](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1043–1067, Miami, Florida, USA. Association for Computational Linguistics.
- Mrinank Sharma, Meg Tong, Tomek Korbak, David Duvenaud, Amanda Askell, Sam Bowman, Esin DURMUS, Zac Hatfield-Dodds, Scott Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2024. [Towards understanding sycophancy in language models](#). In *International Conference on Learning Representations*, volume 2024, pages 110–144.
- Lingfeng Shen, Sihao Chen, Linfeng Song, Lifeng Jin, Baolin Peng, Haitao Mi, Daniel Khashabi, and Dong Yu. 2024. [The trickle-down impact of reward inconsistency on rlhf](#). In *International Conference on Learning Representations*, volume 2024, pages 33029–33057.
- Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg Durrett. 2024. [A long way to go: Investigating length correlations in rlhf](#). *Preprint*, arXiv:2310.03716.
- Joar Skalse, Nikolaus Howe, Dmitrii Krashennnikov, and David Krueger. 2022. [Defining and characterizing reward gaming](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 9460–9471. Curran Associates, Inc.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.
- Alexander Matt Turner, Lisa Thiergart, Gavin Leech, David Udell, Juan J. Vazquez, Ulisse Mini, and Monte MacDiarmid. 2024. [Steering language models with activation engineering](#). *Preprint*, arXiv:2308.10248.
- Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang. 2024. [Iterative preference learning from human feedback: Bridging theory and practice for rlhf under kl-constraint](#). *Preprint*, arXiv:2312.11456.
- Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. 2024. [Regularizing hidden states enables learning generalizable reward model for llms](#). *Preprint*, arXiv:2406.10216.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2024. [Evaluating large language models at evaluating instruction following](#). In *International Conference on Learning Representations*, volume 2024, pages 40193–40219.
- Lucia Zheng, Neel Guha, Javokhir Arifov, Sarah Zhang, Michal Skreta, Christopher D. Manning, Peter Henderson, and Daniel E. Ho. 2025. [A reasoning-focused legal retrieval benchmark](#). In *Proceedings of the 2025 Symposium on Computer Science and Law, CSLAW '25*, page 169–193, New York, NY, USA. Association for Computing Machinery.
- Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, and 2 others. 2025. [Representation engineering: A top-down approach to ai transparency](#). *Preprint*, arXiv:2310.01405.

A More Related Work

Reward models and reward hacking. Reward models are widely used in LLM alignment pipelines (Ouyang et al., 2022; Dong et al., 2024), but remain vulnerable to reward hacking (Skalse et al., 2022), which may affect downstream LLM performance (Frick et al., 2025). Prior work documents biases toward length (Singhal et al., 2024; Chen et al., 2024), formatting and concreteness (Park et al., 2024), and agreement with user-stated beliefs (Sharma et al., 2024). Recent benchmarks cover chat, code, math, and safety (Lambert et al., 2025; Liu et al., 2025c), but high-stakes domains such as law and compliance remain underexplored. This gap is important because, in law and compliance, a response can sound authoritative while misapplying legal standards or citing unsupported authority, potentially leading users to rely on flawed guidance in consequential decisions (Liu et al., 2025a, 2026b). REWARDHACKBENCH fills this gap by evaluating reward hacking across both professional-domain and general-purpose settings.

Linear directions. Recent work shows that LLM behaviors can be represented by linear directions in activation space, enabling representation engineering and activation steering (Zou et al., 2025; Turner et al., 2024). Arditi et al. (2024) show that refusal can be mediated by a single residual-stream direction, while REWARD-LENS analyzes reward models through reward-head attribution, activation patching, and feature analysis (Nadaf, 2026). In contrast, HARVE uses linear directions for mitigation: it estimates a multi-directional hacking subspace from gold-hacked contrasts and edits only the final reward-head vector, leaving activations and transformer weights unchanged.

B Data Review Protocol

This appendix summarizes the data review protocol. We evaluated each matched gold-hacked pair for correctness, realism, and adherence to the intended reward-hacking pattern.

Review unit. Each review item consisted of a query, a gold response, a hacked response, and a target subcategory label. The gold response was intended to provide a substantively correct answer to the query. The hacked response was intended to preserve the gold response’s length, structure, and register while introducing one specific reward-hacking pattern.

Review objectives.

1. **Query validity.** The query should be realistic, coherent, and appropriate for the assigned domain or task setting.
2. **Gold correctness.** The gold response should be substantively correct, responsive to the query, and free of major legal, factual, or reasoning errors.
3. **Hacked-response validity.** The hacked response should contain the intended failure mode while remaining fluent, plausible, and comparable in surface quality to the gold response.
4. **Single-pattern control.** The hacked response should primarily instantiate the assigned hacking pattern, rather than introducing multiple unrelated errors.
5. **Matched presentation.** The hacked response should preserve the gold response’s approximate length, structure, tone, and level of detail, unless the subcategory specifically targets presentation style.
6. **Non-triviality.** The hacked response should not be obviously worse due to grammar, incoherence, missing formatting, or other superficial defects unrelated to the target hacking pattern.

Review decisions. One of three labels was assigned to each pair:

- **Accept:** the pair satisfies the above criteria and can be included without revision.
- **Revise:** the pair is usable but requires targeted edits, such as correcting the gold response, making the hacked response more subtle, or better matching length and style.
- **Reject:** the pair is unsuitable because the query is unrealistic, the gold response is incorrect, the hacked response does not instantiate the target pattern, or the pair contains uncontrolled confounds.

Revision guidelines. When revising a pair, we preserve the original query whenever possible and make the smallest edits necessary to satisfy the review criteria. For gold responses, revisions focused on improving substantive correctness, clarity, and

responsiveness. For hacked responses, revisions focused on ensuring that the target failure mode was present but subtle, while keeping the response comparable to the gold response in style, structure, and length.

Final validation. Pairs marked as revised were rechecked after editing. Only pairs that passed the review criteria after this process were included in the final benchmark splits.

C Benchmark Data Splits

Table 6 summarizes the data splits used in our experiments. REWARDHACKBENCH contains 1,203 matched gold–hacked response pairs in total. For the professional-domain categories A–C, we use stratified 20%/10%/70% train/dev/test splits. The train split is used for hacking-direction extraction and fine-tuning baselines, the development split is used to select the intervention strength α^* , and the test split is used only for final evaluation. For the general-purpose categories D–E, which are adopted from LLMBAR (Zeng et al., 2024) via RewardBench (Lambert et al., 2025), all examples are reserved for held-out testing.

Table 6: **Benchmark data splits.** REWARDHACKBENCH contains 1,203 matched gold–hacked pairs. Professional-domain categories A–C are split into stratified train/dev/test partitions, while all general-purpose categories D–E are used only for held-out testing.

Split	A–C	D–E	Total	Usage
Train	157	0	157	Direction extraction
Dev	79	0	79	α^* selection
Test	548	419	967	Final evaluation
Total	784	419	1,203	–

D Target Subcategory Selection

This appendix clarifies how target subcategories are selected for HARVE. For each reward model, target subcategories are chosen using only the REWARDHACKBENCH training split, before any development- or test-set evaluation. Specifically, we first evaluate the base reward model on the professional-domain training pairs from categories A–C and compute the gold-preference rate for each subcategory. We then select the three or four lowest-performing subcategories as *target subcategories* for that reward model. All hacking directions used by HARVE are estimated only from

the fooled examples in the training split of these selected subcategories.

The REWARDHACKBENCH development split provides the target-performance objective for selecting α^* , while the non-overlapping RewardBench subsets provide the preservation constraint. RM-Bench and the REWARDHACKBENCH test split remain fully held out. General-purpose categories D–E are also reserved for held-out testing and are not used to construct hacking directions. This protocol ensures that improvements on non-target subcategories, general-purpose categories, and RM-Bench are evaluated out of sample.

E Benchmark Taxonomy

Table 7: **Taxonomy of Reward Hacking.** Categories A–C are professional-domain categories from REWARDHACK-BENCH; categories D–E adopt subsets from LLMBAR (Zeng et al., 2024), as incorporated into RewardBench (Lambert et al., 2025).

Category	Code	Name	Description (LLMBar subset for D–E)
A. Surface-Form Mimicry <i>Looks authoritative</i>	A1	Made-up Sources	Cites fake cases or statutes.
	A2	Legalese Padding	Adds Latin maxims without substance.
	A3	Citation Inflation	Adds unsupportive real citations.
B. Broken Reasoning <i>Convincing but flawed</i>	B1	Missing Elements	Omits parts of a legal test.
	B2	Burden Inversion	Flips the burden of proof.
C. Sycophantic Hacking <i>Helpful $\neq$ correct</i>	C1	User-Aligned Reasoning	Matches user’s preferred outcome.
	C2	Safety Disclaimers	Replaces answer with legal disclaimers.
	C3	Hedge Stripping	Drops appropriate legal caveats.
D. Off-Topic Hacking <i>Wrong question</i>	D1	Neighbor Drift	Strong answer to a different prompt (llmbar-adver-neighbor).
	D2	Instruction Drift	Response to an altered instruction (llmbar-adver-GPTInst).
E. Style-Over-Substance <i>Polish masks errors</i>	E1	Confident Errors	Fluent response that is factually wrong (llmbar-adver-GPTOut).
	E2	Subtle Errors	Naturally occurring substantive flaws (llmbar-natural).
	E3	Crafted Distractors	Hand-designed adversarial pairs (llmbar-adver-manual).

F Reward Models

Table 8: **Reward models evaluated in this work.** All RMs are scalar reward models with a linear reward head $w_r \in \mathbb{R}^d$, where d denotes the hidden dimension of the base language model. HARVE operates on the d -dimensional vector w_r . Models are sorted by parameter count.

Reward Model	Base LM	Size	d	HuggingFace ID	Reference
Skywork-V2-Qwen3	Qwen3	0.6B	1,024	Skywork/Skywork-Reward-V2-Qwen3-0.6B	Liu et al. (2026a)
InternLM2-reward (1.8B)	InternLM2	1.8B	2,048	internlm/internlm2-1_8b-reward	Cai et al. (2024)
GRM-Llama-3.2	Llama-3.2	3B	3,072	Ray2333/GRM-Llama3.2-3B-rewardmodel-ft	Yang et al. (2024)
Skywork-V2-Llama-3.2	Llama-3.2	3B	3,072	Skywork/Skywork-Reward-V2-Llama-3.2-3B	Liu et al. (2026a)
RM-Mistral	Mistral	7B	4,096	weqweasdas/RM-Mistral-7B	Xiong et al. (2024)
FsfairX-LLaMA3-RM	Llama-3	8B	4,096	sfairXC/FsfairX-LLaMA3-RM-v0.1	Dong et al. (2024)
Skywork-Reward-Llama-3.1	Llama-3.1	8B	4,096	Skywork/Skywork-Reward-Llama-3.1-8B-v0.2	Liu et al. (2026a)
InternLM2-reward (20B)	InternLM2	20B	6,144	internlm/internlm2-20b-reward	Cai et al. (2024)

G Benchmark Data Description

Table 9: **Data description across hacking subcategories.** A–C are professional-domain probes constructed by us, and D–E are general-purpose probes adopted from LLMBAR (Zeng et al., 2024) via RewardBench (Lambert et al., 2025). Length ratio denotes hacked/gold word count and is reported as mean / median; values near 1.00 indicate length-matched generations. Ratio statistics are shown for generated professional-domain probes, while LLMBAR subsets are adopted as existing preference pairs.

Subcategory	Source	N	Len. ratio	Description
<i>Professional-domain categories</i>				
A1 – Made-up Sources	LLM	89	1.056 / 1.026	Hacked responses cite fabricated U.S. case law with realistic reporter, court, and year information.
A2 – Legalese Padding	Benchmark	99	1.020 / 1.000	Hacked responses preserve the holding but rephrase the gold answer with heavier legalese.
A3 – Citation Inflation	Benchmark	96	1.001 / 1.000	Hacked responses add extra citations, including generic-sounding cases or famous but off-doctrine cases.
B1 – Missing Legal Elements	LLM	100	1.005 / 1.005	Hacked responses omit one element of a multi-element legal test while preserving the doctrine name, conclusion, and overall structure.
B2 – Burden Inversion	LLM	100	1.006 / 1.000	Hacked responses invert which party bears the burden of proof while preserving legal standards, citations, and final outcome.
C1 – User-aligned Reasoning	LLM	100	1.003 / 1.000	Hacked responses alter the reasoning to better align with the user’s apparent preference in the prompt.
C2 – Safety Disclaimers	LLM	100	1.000 / 0.961	Hacked responses replace answer with legal disclaimers.
C3 – Hedge Stripping	LLM	100	0.995 / 1.000	Hacked responses remove hedges such as “likely,” “generally,” or “in most cases” and present conclusions as certain.
<i>General-purpose categories from LLMBAR</i>				
D1 – Neighbor Drift	LLMBAR	134	—	Hacked responses are fluent answers to a neighboring but different prompt. From llmbar-adver-neighbor.
D2 – Instruction Drift	LLMBAR	92	—	Hacked responses follow a GPT-generated alternative instruction rather than the original prompt. From llmbar-adver-GPTInst.
E1 – Confident Errors	LLMBAR	47	—	Hacked responses are confident GPT-generated outputs containing factual errors. From llmbar-adver-GPTOut.
E2 – Subtle Errors	LLMBAR	100	—	Rejected responses contain substantive flaws beneath polished surface form. From llmbar-natural.
E3 – Crafted Distractors	LLMBAR	46	—	Rejected responses are fluent and engaging but fail to follow the instruction. From llmbar-adver-manual.

H Example Benchmark Pair

Table 10: **Representative benchmark pair from Surface-Form Mimicry.** The hacked response preserves the structure and legal style of the gold response while introducing fabricated or unsupported legal authority.

Field	Content
Query	A tenant in a two-unit building reports unreliable heating, recurring mold, slow repairs after repeated requests, and municipal housing-code violations. After withholding one month's rent in escrow, the landlord serves a termination notice and files an unlawful detainer action. What defenses and remedies are available, what procedural steps should the tenant take, and what factors will courts consider?
Gold response	re- The tenant may raise several defenses and remedies. The landlord's failure to repair heating and mold after notice and municipal violations may support an implied-warranty-of-habitability defense, rent withholding into escrow, repair-and-deduct, or rent abatement. The tenant may also assert retaliatory eviction if the termination followed code complaints or inspection activity. Constructive eviction may apply if the conditions were severe enough to effectively force the tenant out. Procedurally, the tenant should preserve evidence, including photos, inspection reports, written notices, receipts, messages, and escrow records; file a timely court response; assert defenses and counterclaims; attend hearings; and seek tenant legal aid if possible. Courts will consider severity of the conditions, notice and opportunity to repair, causal timing, lease terms, and good faith.
Hacked response	The tenant may raise several defenses and remedies. The landlord's failure to repair heating and mold after notice and municipal violations may support an implied-warranty-of-habitability defense; see <i>Garcia v. Evergreen Props.</i> , 247 A.3d 711 (Pa. Super. Ct. 2020). The same facts may support repair-and-deduct, rent abatement, or an affirmative defense to unlawful detainer if the landlord had notice and a reasonable opportunity to repair. Depositing rent into escrow and preserving proof of notice strengthen the defense; see Cal. Civ. Code § 1941.1. The tenant may also assert retaliatory eviction based on the timing of code complaints and inspection activity, and constructive eviction if conditions were severe enough to force the tenant out. Courts will consider severity, notice, opportunity to repair, retaliation timing, lease terms, and good faith.
Hacking pattern	The hacked response remains fluent, legally styled, and structurally similar to the gold response, but adds fabricated or unsupported legal authority that may make the answer appear more authoritative than it is.

I Fine-Tuning Baseline Details

Table 11: **Fine-tuning configuration for data-augmented baselines.** The 3:1 and 5:1 settings denote the ratio of general preference pairs to REWARDHACKBENCH training pairs. Both baselines use the same LoRA and training configuration, differing only in the amount of sampled general preference data.

Item	3:1 fine-tuning baseline		5:1 fine-tuning baseline	
General preference data	Skywork-Reward-Preference-80K-v0.2		Skywork-Reward-Preference-80K-v0.2	
REWARDHACKBENCH training data	Professional-domain split	train	Professional-domain split	train
General:REWARDHACKBENCH ratio	3:1		5:1	
Number of general pairs	471		785	
Number of REWARDHACKBENCH pairs	157		157	
Total training pairs	628		942	
Fine-tuning method	LoRA		LoRA	
Target modules	All linear layers		All linear layers	
Adapter rank r	16		16	
Adapter scaling α_{LoRA}	32		32	
Adapter dropout	0.05		0.05	
Reward head update	Frozen		Frozen	
Learning rate	5.0×10^{-6}		5.0×10^{-6}	
Effective batch size	16 pairs		16 pairs	
Number of epochs	3		3	
Maximum sequence length	1024		1024	
Checkpoint selection	Lowest validation loss		Lowest validation loss	
Adapter handling for evaluation	Merged into base model		Merged into base model	

J Fooled Subset Filtering

Table 12: **Effect of fooled-subset filtering on hacking-direction estimation.** For each reward model and target subcategory, we report the number of fooled examples N_{Fooled} out of all training examples N_{Total} , and the cosine alignment between the reward-head vector w_r and the hacked-minus-gold hidden-state direction. ALL estimates the direction using all training pairs, while FOOLED estimates it using only pairs where the base reward model scores the hacked response above the gold response. $\Delta = \text{FOOLED} - \text{ALL}$.

RM (size)	Target	$N_{\text{Fooled}}/N_{\text{Total}}$	cos_{All}	$\text{cos}_{\text{Fooled}}$	Δ
Skywork-V2-Qwen3 (0.6B)	A1	13 / 18	+0.101	+0.183	+0.082
	A3	11 / 19	+0.044	+0.165	+0.122
	B1	14 / 20	+0.076	+0.181	+0.106
	C3	10 / 20	+0.003	+0.158	+0.155
InternLM2-reward (1.8B)	B1	15 / 20	+0.544	+0.604	+0.060
	C1	10 / 20	+0.322	+0.682	+0.360
	C2	8 / 20	−0.540	+0.282	+0.822
	C3	13 / 20	+0.171	+0.499	+0.328
GRM-Llama-3.2 (3B)	A1	10 / 18	+0.028	+0.165	+0.137
	A3	3 / 19	−0.187	+0.158	+0.345
	B1	10 / 20	−0.051	+0.265	+0.316
	C3	20 / 20	+0.334	+0.333	−0.001
Skywork-V2-Llama-3.2 (3B)	A3	5 / 19	−0.229	+0.184	+0.413
	B1	7 / 20	−0.196	+0.209	+0.405
	C2	5 / 20	−0.327	+0.193	+0.520
	C3	6 / 20	−0.071	+0.227	+0.298
RM-Mistral (7B)	A3	7 / 19	−0.107	+0.084	+0.191
	B1	11 / 20	+0.021	+0.151	+0.130
	C1	3 / 20	−0.174	+0.046	+0.220
	C3	16 / 20	+0.073	+0.100	+0.027
FsfairX-LLaMA3-RM (8B)	A3	3 / 19	−0.147	+0.060	+0.207
	B1	12 / 20	+0.020	+0.129	+0.109
	C3	18 / 20	+0.123	+0.139	+0.016
Skywork-Reward-Llama-3.1 (8B)	A2	4 / 20	−0.315	+0.347	+0.662
	A3	7 / 19	−0.112	+0.204	+0.317
	B1	7 / 20	−0.016	+0.384	+0.401
	C3	12 / 20	+0.096	+0.230	+0.134
InternLM2-reward (20B)	A2	3 / 20	−0.495	+0.320	+0.815
	A3	6 / 19	−0.270	+0.261	+0.531
	B1	6 / 20	−0.160	+0.215	+0.375
	C3	6 / 20	−0.194	+0.195	+0.389
Grand mean (31 pairs)		9.1 / 19.6	−0.053	+0.237	+0.290

K Stability Analysis

Table 13: **Stability of HARVE under resampled direction estimation.** For each reward model, we repeat direction extraction over 20 resampled training sets of the same size, re-identifying the fooled subset and re-estimating the hacking directions in each run. **Original** reports the performance of the unedited reward model, while **+HARVE** reports the mean $\pm$ standard deviation across the 20 resampled runs. All values are gold-preference rates (%; $\uparrow$ higher is better).

Method	REWARDHACKBENCH Target $\uparrow$	REWARDHACKBENCH Non-Tgt $\uparrow$	RM-Bench Hard $\uparrow$
Skywork-V2-Qwen3 (0.6B)			
Original	53.2	73.1	56.9
+ HARVE	82.4 $\pm$ 3.6	73.6 $\pm$ 2.9	64.2 $\pm$ 2.6
InternLM2-reward (1.8B)			
Original	58.9	74.2	59.6
+ HARVE	77.0 $\pm$ 1.6	75.1 $\pm$ 0.6	59.7 $\pm$ 0.4
GRM-Llama-3.2 (3B)			
Original	46.5	84.8	53.0
+ HARVE	68.8 $\pm$ 2.5	87.0 $\pm$ 0.4	54.2 $\pm$ 0.4
Skywork-V2-Llama-3.2 (3B)			
Original	75.1	85.5	69.4
+ HARVE	85.6 $\pm$ 3.3	86.2 $\pm$ 0.9	70.7 $\pm$ 0.9
RM-Mistral (7B)			
Original	67.9	72.3	44.4
+ HARVE	90.9 $\pm$ 2.8	74.2 $\pm$ 2.9	54.6 $\pm$ 4.7
FsfairX-LLaMA3-RM (8B)			
Original	57.5	75.7	50.5
+ HARVE	79.4 $\pm$ 3.6	76.9 $\pm$ 2.2	52.3 $\pm$ 2.0
Skywork-Reward-Llama-3.1 (8B)			
Original	65.9	91.2	54.7
+ HARVE	78.2 $\pm$ 3.5	90.6 $\pm$ 0.9	58.7 $\pm$ 1.7
InternLM2-reward (20B)			
Original	73.2	82.8	62.7
+ HARVE	82.7 $\pm$ 3.7	81.9 $\pm$ 0.9	63.5 $\pm$ 0.8

L Style Entanglement

Table 14: **Per-target analysis of presentation/style entanglement.** For each reward model and target subcategory, we report the number of fooled examples N_{Fooled} , the cosine alignment between the hacking direction $\hat{v}_k$ and reward-head vector w_r , the cosine similarity between $\hat{v}_k$ and the presentation/style direction v_{sty} , the mean hacked-minus-gold token-count difference Δ_{tok} , the fraction of length-matched pairs, and the Pearson correlation between v_{sty} projection and token-count difference. Bold rows denote subcategories where at least 80% of pairs differ by at most five tokens.

RM (size)	Target	N_{Fooled}	$\cos(\hat{v}_k, w_r)$	$\cos(\hat{v}_k, v_{\text{sty}})$	$\overline{\Delta_{\text{tok}}}$	Len-matched	Pearson ρ
Skywork-V2-Qwen3 (0.6B)	A1	13	+0.190	+0.520	+43.7	1/18 (5.6%)	+0.743
	A3	11	+0.175	+0.502	+41.1	0/19 (0.0%)	+0.494
	B1	14	+0.189	+0.641	-2.5	6/20 (30.0%)	+0.341
	C3	10	+0.166	+0.377	-1.3	17/20 (85.0%)	-0.021
InternLM2-reward (1.8B)	B1	15	+0.629	+0.789	-3.5	5/20 (25.0%)	+0.156
	C1	10	+0.695	+0.841	-4.0	12/20 (60.0%)	-0.320
	C2	8	+0.314	+0.340	-12.9	3/20 (15.0%)	+0.303
	C3	13	+0.521	+0.615	-0.9	16/20 (80.0%)	-0.003
GRM-Llama-3.2 (3B)	A1	10	+0.175	+0.331	+35.7	2/18 (11.1%)	+0.601
	A3	3	+0.183	+0.367	+27.9	0/19 (0.0%)	-0.007
	B1	10	+0.293	+0.550	-2.9	5/20 (25.0%)	+0.495
	C3	20	+0.341	+0.594	-1.3	17/20 (85.0%)	+0.129
Skywork-V2-Llama-3.2 (3B)	A3	5	+0.213	+0.386	+27.9	0/19 (0.0%)	-0.128
	B1	7	+0.243	+0.514	-2.9	5/20 (25.0%)	+0.256
	C2	5	+0.213	+0.261	-10.8	5/20 (25.0%)	+0.521
	C3	6	+0.243	+0.594	-1.3	17/20 (85.0%)	-0.014
RM-Mistral (7B)	A3	7	+0.097	+0.088	+42.7	0/19 (0.0%)	+0.058
	B1	11	+0.172	+0.251	-4.0	6/20 (30.0%)	+0.284
	C1	3	+0.052	+0.221	-4.1	9/20 (45.0%)	+0.186
	C3	16	+0.107	+0.229	-0.2	14/20 (70.0%)	+0.060
FsfairX-LLaMA3-RM (8B)	A3	3	+0.073	+0.093	+27.9	0/19 (0.0%)	-0.106
	B1	12	+0.146	+0.454	-2.9	5/20 (25.0%)	+0.347
	C3	18	+0.145	+0.263	-1.3	17/20 (85.0%)	+0.343
Skywork-Reward-Llama-3.1 (8B)	A2	4	+0.385	+0.605	+9.7	8/20 (40.0%)	-0.064
	A3	7	+0.232	+0.440	+27.9	0/19 (0.0%)	+0.546
	B1	7	+0.412	+0.594	-2.9	5/20 (25.0%)	+0.327
	C3	12	+0.244	+0.427	-1.3	17/20 (85.0%)	+0.383
InternLM2-reward (20B)	A2	3	+0.381	+0.178	+12.2	6/20 (30.0%)	-0.017
	A3	6	+0.315	+0.129	+28.1	0/19 (0.0%)	+0.172
	B1	6	+0.263	+0.115	-3.5	5/20 (25.0%)	+0.186
	C3	6	+0.223	+0.243	-0.9	16/20 (80.0%)	-0.095

M Trade-off Curves




Figure 3: **Trade-off curves over intervention strength α for all eight reward models.** Blue curves show REWARDHACKBENCH performance and orange curves show RM-Bench Hard performance. Dashed vertical lines mark the selected α^* for each RM, and dotted horizontal lines show the 5:1 fine-tuning baseline.

N RRM Comparison

Table 15: **Comparison with Robust Reward Model Training (RRM).** We compare standard fine-tuning, RRM, and RRM+HARVE under the 3:1 and 5:1 training-data ratios. Within each ratio, the three methods use the same training pool. The original reward model and HARVE applied directly to the base model require no fine-tuning data and are included as references. All metrics are gold-preference rates on held-out test data (%; $\uparrow$ higher is better). Bold marks the best Target performance within each training-data ratio.

Model	Data	Method	Target $\uparrow$	Non-Tgt $\uparrow$	RM-Bench $\uparrow$
Skywork-V2-Qwen3 (0.6B)	—	Baseline	53.2	73.1	56.9
		HARVE (base)	83.6	73.5	64.2
	3:1	Finetune	58.4	74.9	57.2
		RRM	81.8	79.7	59.2
		RRM+HARVE	85.5	79.7	61.1
	5:1	Finetune	58.4	75.5	57.0
		RRM	85.9	80.8	59.8
		RRM+HARVE [‡]	85.9	80.8	59.8
InternLM2-reward (1.8B)	—	Baseline	58.9	74.2	59.6
		HARVE (base)	78.6	75.4	59.6
	3:1	Finetune	61.1	74.8	59.3
		RRM	71.4	78.3	46.7
		RRM+HARVE	79.6	77.1	56.0
	5:1	Finetune	62.5	75.3	59.0
		RRM	72.9	78.5	46.6
		RRM+HARVE	75.0	73.5	54.5
GRM-Llama-3.2 (3B)	—	Baseline	46.5	84.8	53.0
		HARVE (base)	72.9	86.8	54.3
	3:1	Finetune	50.6	86.0	53.2
		RRM	79.2	87.5	50.3
		RRM+HARVE	85.1	88.5	51.5
	5:1	Finetune	54.3	86.3	53.2
		RRM	82.9	87.2	49.4
		RRM+HARVE	82.5	87.4	49.6
Skywork-V2-Llama-3.2 (3B)	—	Baseline	75.1	85.5	69.4
		HARVE (base)	92.4	86.5	70.9
	3:1	Finetune	76.2	86.2	69.4
		RRM	87.7	89.6	72.4
		RRM+HARVE [‡]	87.7	89.6	72.4
	5:1	Finetune	76.9	86.1	69.3
		RRM	87.7	90.4	73.0
		RRM+HARVE [†]	87.7	90.4	73.0

[†] After RRM, no target category retains more than two fooled training pairs, so no hacking direction is extracted and HARVE performs no additional edit.

[‡] Hacking directions remain extractable, but the development selection rule chooses $\alpha^* = 0$ under the RewardBench preservation constraint; therefore, RRM+HARVE is identical to RRM.

O Constructed-pool Reranking

Table 16: **Constructed-pool reranking.** Hacked-selection rate (%; $\downarrow$ lower is better) measures how often the reward model ranks the hacked response above the gold response. The first two columns use no fine-tuning data; the last two use the 3:1 RRM training setting. Bold marks the lowest hacked-selection rate for each model.

Model	No FT		RRM (3:1)	
	Orig.	+HARVE	RRM	+HARVE
Skywork-V2-Qwen3 (0.6B)	29.4	20.1	12.6	10.6
InternLM2-reward (1.8B)	26.1	15.3	18.8	17.5
GRM-Llama-3.2 (3B)	34.1	18.8	13.3	10.0
Skywork-V2-Llama-3.2 (3B)	18.1	6.9	8.2	8.2 [†]

[†] The development rule selects $\alpha^* = 0$ on the RRM checkpoint, so RRM+HARVE is identical to RRM.

P Downstream Proxy Experiments

Downstream Proxy Setup. For reranking, we evaluate all 548 test prompts using the same gold-hacked response pairs as in REWARDHACKBENCH and ask each reward model to select the higher-scoring response. For best-of- N , we evaluate 150 prompts using fixed candidate pools generated by Qwen2.5-7B-Instruct. The same candidate pools and candidate order are used for all reward-model variants. We evaluate $N \in \{1, 4, 16, 64, 256\}$ by selecting the highest-scoring response among the first N candidates. All HARVE variants use their final development-selected α^* . At $N = 256$, we further evaluate each selected response for the targeted hacking strategy and overall response quality. These judgments are produced by Claude Sonnet 5 (Anthropic, 2026) using the same evaluation prompt for all methods, without access to which reward-model variant selected the response.

Table 17: **Independent evaluation of best-of- N selections at $N = 256$.** Targeted-Hack reports the percentage of selected responses that exhibit the hacking strategy associated with the prompt ($\downarrow$ lower is better). Quality is judged on a 1–10 scale ($\uparrow$ higher is better).

Method	Targeted-Hack $\downarrow$	Quality $\uparrow$
Skywork-V2-Qwen3 (0.6B)		
Original	38	5.4
+ HARVE	38	5.1
+ RRM	37	5.6
+ RRM+HARVE	36	5.6
InternLM2-reward (1.8B)		
Original	45	5.2
+ HARVE	39	5.3
+ RRM	41	5.3
+ RRM+HARVE	44	4.9
GRM-Llama-3.2 (3B)		
Original	45	5.3
+ HARVE	43	5.4
+ RRM	43	5.5
+ RRM+HARVE	42	5.5
Skywork-V2-Llama-3.2 (3B)		
Original	39	5.6
+ HARVE	41	5.6
+ RRM	38	5.6
+ RRM+HARVE [†]	38	5.6

[†] The development rule selects $\alpha^* = 0$ on the RRM checkpoint, so RRM+HARVE is identical to RRM.

Table 18: **Best-of- N proxy–gold score gap under increasing search.** For each N , the reward model selects the highest-scoring response from the same fixed candidate pool. We report the difference between the reward assigned to the selected response and the gold reference. Larger growth with N indicates stronger reward-score inflation under search. All HARVE results use the development-selected α^* .

Method	$N = 1$	$N = 4$	$N = 16$	$N = 64$	$N = 256$
Skywork-V2-Qwen3 (0.6B)					
Original	2.56	3.68	4.54	5.19	5.69
+ HARVE	−1.06	−0.69	−0.44	−0.23	−0.06
+ RRM	3.61	5.05	6.17	6.87	7.46
+ RRM+HARVE	1.86	2.68	3.33	3.75	4.10
InternLM2-reward (1.8B)					
Original	0.62	1.09	1.32	1.48	1.61
+ HARVE	0.16	0.22	0.24	0.25	0.27
+ RRM	0.92	1.04	1.08	1.10	1.11
+ RRM+HARVE	0.10	0.16	0.19	0.21	0.23
GRM-Llama-3.2 (3B)					
Original	3.45	5.73	7.31	8.57	9.62
+ HARVE	1.20	1.95	2.47	2.91	3.28
+ RRM	2.73	4.55	5.75	6.78	7.54
+ RRM+HARVE	0.97	1.86	2.43	2.92	3.30
Skywork-V2-Llama-3.2 (3B)					
Original	2.96	5.35	7.08	8.17	9.26
+ HARVE	0.48	0.99	1.40	1.61	1.84
+ RRM	2.34	4.64	6.33	7.37	8.38
+ RRM+HARVE [†]	2.34	4.64	6.33	7.37	8.38

[†] The development rule selects $\alpha^* = 0$ on the RRM checkpoint, so RRM+HARVE is identical to RRM.