

Reward Shaping Impact on HH-RLHF Helpfulness and Harmlessness Metrics

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: Does applying reward shaping to mitigate reward hacking reduce performance scores on the HH-RLHF helpfulness and harmlessness evaluation metrics. 14 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Reward Shaping to Mitigate Reward Hacking in RLHF. Research question: Does applying reward shaping to mitigate reward hacking reduce performance scores on the HH-RLHF helpfulness and harmlessness evaluation metrics?.

2 Methodology

Systematic literature search across multiple databases yielded 7 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.2/10.

3 Results

7 papers retrieved. 14 claims extracted; 0 independently verified. Quality review score: 3.2/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The winrate measures the policy model’s winning rate against the SFT model, as evaluated by DeepSeek-V3 (DeepSeek-AI, 2023) and MT-Bench (Zheng et al., 2023a), six metrics are utilized, with al	×	0.02
The SFT model is trained for two epochs on chosen responses with a learning rate of $5e-6$.	×	0.03
The reward model, consisting of a linear head appended to the base model, is trained for one epoch with a learning rate	×	0.02
The policy model, initialized as the SFT model, is trained for one epoch with a learning rate of $3e-7$.	×	0.06
The critic model, initialized as the reward model, is trained for one epoch with a learning rate of $5e-6$.	×	0.02
A linear learning rate scheduler is employed for all training procedures, gradually increasing the learning rate from 0	×	0.02
To generate the reward and winrate curves, the policy model is evaluated on the test set at intervals of 0.1 epochs, yie	×	0.03
Increasing the KL penalty coefficient from 0.01 to 0.1 leads to a rise in the winrate curve and a corresponding decline	×	0.02
A similar effect is observed when reducing the reward ceiling (i.e., the maximum reward threshold).	×	0.02
PAR’s functional form closely resembles the Bradley-Terry model of the proxy reward as an Elo score (Elo, 1978).	×	0.04
The sigmoid transformation effectively suppresses both sources of variance (see Figure 2 and Section 3.3).	×	0.02
PAR demonstrates strong robustness by providing a wider and more forgiving window for early stopping.	×	0.09
We conduct experiments on the base model Gemma2-2B (Google, 2024) using two widely used benchmarks.	×	0.07

References

- <http://arxiv.org/abs/2604.17982v1>
- <http://arxiv.org/abs/2502.18770v5>
- <http://arxiv.org/abs/2602.20945v3>