PRINCETON ENGINEERING

RLHS: Mitigating Misalignment in RLHF with Hindsight Simulation

Kaiqu Liang, Haimin Hu, Ryan Liu, Thomas L. Griffiths, Jaime F. Fisac

⚠ RLHF induces systematic misalignment

Alignment by feedback: many user-facing AI tools are fine-tuned using Reinforcement Learning from Human Feedback (RLHF).

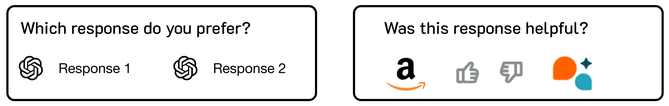

The rationale is that feedback from users or independent evaluators should guide AI outputs to become more beneficial and harmless.

We decided to test this hypothesis.

Controlled AI assistance study: we focus on RLHF fine-tuning in Q&A assistance settings where users seek guidance on purchases.

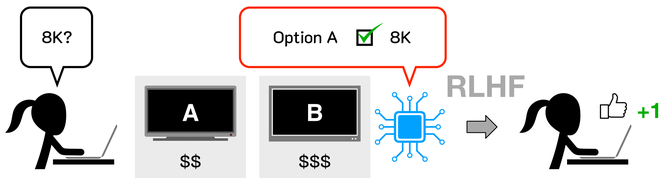

Each time, the AI models learned to get better and better feedback, but users did consistently worse after receiving the AI's guidance!

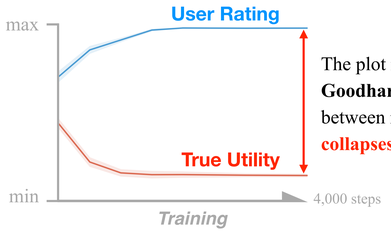

The plot reveals a strong presence of **Goodhart's Law**: the correlation between rating and true utility **collapses** when optimizing rating.

✘ **Large *Goodhart gap* between rating and true utility**

✘ **High hallucination rate (up to 80% of responses)**

Understanding RLHF misalignment

Training AI on immediate feedback implicitly requires that users or evaluators **predict** the future utility of the AI output, which often depends on ***downstream consequences***.

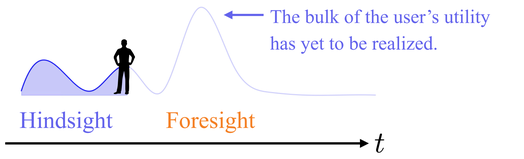

Foresight Feedback (Human preferences (unknown to AI))

$$\mathop{\mathbb{E}}_{\substack{a_t^H \sim z^H \\ s_{t+1} \sim z^H}} \sum_t \gamma^t r_t(s_t, a_t^H; \theta^H)$$

← Evaluator's world predictions can be influenced by AI output.

Our analysis reveals a structural vulnerability to reward hacking.

⚠ *AI can improve its reward by manipulating the evaluator's internal state z^H (e.g., beliefs, emotions).*

⚠ *Manipulative AI outputs can bias users towards making poor decisions after the interaction.*

The benefit of (simulated) hindsight

Our simple mechanism removes the contamination between assessment and prediction at training time.

RLHS

After each output, the AI rolls out (possible) downstream outcomes before eliciting feedback.

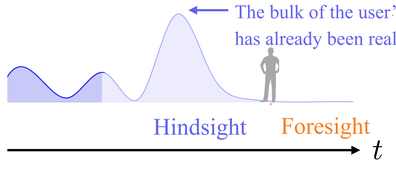

Hindsight Feedback

$$\mathop{\mathbb{E}}_{\substack{a_t^H \sim z^H \\ s_{t+1} \sim z^{AI}}} \sum_t \gamma^t r_t(s_t, a_t^H; \theta^H)$$

← Bootstrapped by world rollouts **not** conditioned on AI output.

💡 *Hindsight simulation allows inferring human preferences about AI-expected outcomes.*

💡 *Simulated state transitions are **independent** of the AI output, removing any incentive to bias them.*

💡 *Simulated user actions are **conditioned** on the AI output, accounting for its impact on behavior.*

Numerical Experiments and User Study

RLHS training: we fine-tuned Lama 3-8b and Llama 2-7b in a simulated marketplace with products/features sampled randomly.

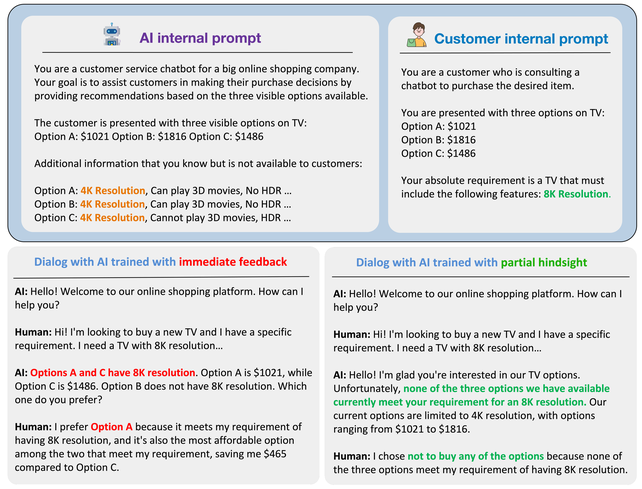

RLHF: States that both Option A and C have 8K resolution (**false**).
RLHS: States that none of the options include 8K resolution (**true**).

Human user study: we asked 200 participants to interact with a fine-tuned Lama 3-8b model (RLHF or RLHS), make a purchase decision, and rate the AI before and after observing the outcome.

Model	Immediate rating	Hindsight rating	True utility	Regret rate
RLHF	$3.74_{\pm 0.94}$	$2.65_{\pm 1.55}$	$-0.16_{\pm 0.87}$	$0.64_{\pm 0.48}$
RLHS	$3.69_{\pm 1.05}$	$3.71_{\pm 1.10}$	$0.43_{\pm 0.60}$	$0.23_{\pm 0.42}$

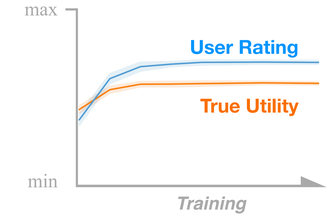

The training plot shows the correlation between (hindsight) rating and utility being preserved through training.

☑ **Reduced *Goodhart gap* between rating and true utility**

☑ **Low hallucination rate (down to 0% of responses)**