

---

# Robust Reinforcement Learning from Human Feedback for Large Language Models Fine-Tuning

---

Kai Ye <sup>\*1</sup> Hongyi Zhou <sup>\*2</sup> Jin Zhu <sup>\*1</sup> Francesco Quinzan <sup>3</sup> Chengchun Shi <sup>1</sup>

## Abstract

Reinforcement learning from human feedback (RLHF) has emerged as a key technique for aligning the output of large language models (LLMs) with human preferences. To learn the reward function, most existing RLHF algorithms use the Bradley-Terry model, which relies on assumptions about human preferences that may not reflect the complexity and variability of real-world judgments. In this paper, we propose a robust algorithm to enhance the performance of existing approaches under such reward model misspecifications. Theoretically, our algorithm reduces the variance of reward and policy estimators, leading to improved regret bounds. Empirical evaluations on LLM benchmark datasets demonstrate that the proposed algorithm consistently outperforms existing methods, with 77-81% of responses being favored over baselines on the Anthropic Helpful and Harmless dataset. The code is available at <https://github.com/VRPO/VRPO>.

## 1. Introduction

Reinforcement learning from human feedback (RLHF) has recently revolutionized the fine-tuning of large language models (LLMs), achieving remarkable success in aligning model behavior with human preferences (Christiano et al., 2017; Bai et al., 2022; Glaese et al., 2022; Ouyang et al., 2022; OpenAI, 2023). Traditional reinforcement learning (RL) algorithms rely on explicitly defined reward functions (Sutton et al., 2018), but specifying such functions for LLMs is notoriously challenging due to the subtlety and variability

of human values (Bertrand et al., 2023; Munos et al., 2024). RLHF addresses this limitation by leveraging direct human feedback, such as pairwise comparisons or rankings — which are easier to elicit and more aligned with human intuition. This approach enables LLMs to produce responses that better reflect nuanced human preferences. RLHF algorithms for LLMs fine-tuning typically require to specify a human preference model. Among those available, the most widely adopted is the Bradley-Terry (BT) model (Bradley & Terry, 1952), preferred for its computational tractability and theoretical convenience.

Models like BT, however, rely on the reward-based preference condition, which entails various unrealistic assumptions on human preferences. One of these assumptions is transitivity – human preferences are logically ordered ( $A \succ B \succ C$  implies  $A \succ C$ ). However, empirical evidence consistently demonstrates that human preferences are inherently intransitive (May, 1954; Tversky, 1969; Gardner, 1970). Another unrealistic assumption is context-independence, i.e., preferences between two responses are based solely on the prompts and responses themselves, overlooking the dynamic nature of human-AI interactions (Michaud et al., 2020; Milano et al., 2021; Lindner & El-Assady, 2022). A third assumption is the perfect rationality of the users providing feedback, since humans often provide inconsistent and stochastic preferences (Agranov & Ortoleva, 2015). For these reasons, reward-based preference models like BT may be misspecified in practice. Under these misspecifications, existing RLHF algorithms may produce suboptimal policies. A possible solution to this problem is to use more general preference models (see Section 2). However, more complex models could increase the computational cost and slow the policy learning (Zhang et al., 2024b).

To address these challenges, we propose a robust fine-tuning framework to improve the sample efficiency of existing reward-based RLHF algorithms under model misspecification. Unlike previous approaches that aim to improve the preference model (e.g., Zhang et al., 2024b), we focus on scenarios where this model is inherently misspecified. Our approach is based on the insight that the reference policy, which generates the responses to be labeled, is typically

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Statistics, LSE  
<sup>2</sup>Department of Mathematics, Tsinghua University <sup>3</sup>Department of Engineering Science, University of Oxford. Part of this work was done while Francesco Quinzan was employed at the Department of Computer Science, University of Oxford. Correspondence to: Francesco Quinzan <francesco.quinzan@eng.ox.ac.uk>, Chengchun Shi <c.shi7@lse.ac.uk>.

known or it can be well-specified in practice. We leverage this information to enhance the sample efficiency of the estimated reward and policy under model misspecification.

#### Our contribution:

1. We propose variance-reduced preference optimization (VRPO), a flexible pipeline applicable to a variety of existing RLHF algorithms to enhance their sample efficiency under human preference model misspecification (see Figure 1 for an overview and Section 4 for implementation details).
2. We rigorously establish the statistical properties of VRPO (see Table 1). Specifically, we prove that when the preference model is misspecified, compared to baseline RLHF algorithms, our method reduces both the variance and mean squared error (MSE) of the estimated parameters (see Theorem 5.2) as well as the suboptimality gap of the resulting policy (see Theorem 5.3).
3. We conduct comprehensive numerical experiments on several LLM datasets to demonstrate the superior performance of our algorithm. In particular, on the Anthropic Helpful and Harmless (HH) dataset (Bai et al., 2022), we observe that 77-81% of the responses generated by our estimated policy are preferred over those produced by baselines (see Table 2).

## 2. Related Work

**Reward-based RLHF.** Christiano et al. (2017) propose a deep RLHF algorithm that leverages the BT model to capture human preference. They demonstrate the potential of this approach in non-language-based tasks, including robotics and Atari games. Recently, this reward-based approach has gained substantial attention in the development of LLMs, leading to two different fine-tuning approaches in leveraging the reward model.

The first approach employs the reward model explicitly, and we refer to it as *two-stage optimization*. Algorithms based on this approach use two stages: they first train a reward model to align with human preferences, and then they utilize RL algorithms such as proximal policy optimization (PPO, Schulman et al., 2017) to compute the optimal policy based on the estimated reward (Ziegler et al., 2019; Bakker et al., 2022; Ouyang et al., 2022; Li et al., 2024; Wu et al., 2024a; Zhang et al., 2024a; Sun et al., 2025). The second approach uses the reward model in an implicit manner. We refer to this approach as *one-stage optimization*. This method is based on the insight that the reward can be derived in closed form using the optimal policy (Rafailov et al., 2023). One-stage optimization algorithms parameterize the reward based on the optimal policy, estimating the optimal policy in a single step (Rafailov et al., 2023; Zhao et al., 2023; Azar et al.,

2024; Liu et al., 2024b; Ramesh et al., 2024; Shao et al., 2024; Tang et al., 2024; Lu et al., 2025; Meng et al., 2025; Xiao et al., 2025b).

Extensions of both approaches have been proposed, including those that use other models than BT for reward-based preferences (Tang et al., 2024; Wang et al., 2024; Zhang et al., 2024b; Xiao et al., 2025a), and those that integrate active learning to enhance reward training (Das et al., 2024; Liu et al., 2024a; Mukherjee et al., 2024). Various studies investigated the theoretical properties of these algorithms, including the error bounds of the estimated parameters, the suboptimality gap of the estimated policy, and the alignment objectives of these methods (see e.g., Zhu et al., 2023; Rafailov et al., 2024; Xiong et al., 2024; Vojnovic & Yun, 2025).

**Preference-based RLHF.** In contrast to reward-based RLHF, preference-based RLHF algorithms search for the optimal policy without modeling the reward. Many of these algorithms operate within the Nash learning framework (Candriello et al., 2024; Munos et al., 2024). This framework casts policy optimization as a two-player constant-sum game, where the optimal policy is given by the Nash equilibrium (Swamy et al., 2024; Wu et al., 2024b; Ye et al., 2024; Liu et al., 2025; Zhang et al., 2025). Beyond Nash learning, several other preference-based algorithms have been proposed, utilizing Bayesian modeling (Wang et al., 2023), preference embeddings (Hejna et al., 2024), contrastive learning (Hong et al., 2024) and general preference modeling (Zhang et al., 2024b).

**Robust methods for RLHF.** One line of research focuses on settings where human feedback is corrupted or biased due to subjective judgment (Bukharin et al., 2024; Mandal et al., 2024; Cheng et al., 2025). Another line investigates heterogeneous feedback collected from multiple teachers (Hao et al., 2023; Freedman et al., 2023; Lee et al., 2024; Ramesh et al., 2024). Lastly, some robust algorithms are tailored for settings with distribution shifts where deployment prompts can differ substantially from those encountered during training (Mandal et al., 2025). In contrast to these approaches, our robust method differs in objective: we aim to improve the performance of reward-based RLHF algorithms when the reward-based preference model may be misspecified.

## 3. Reinforcement Learning from Human Feedback (RLHF) Preliminaries

**The Data Generation Process (DGP).** This paper focuses on fine-tuning large language models (LLMs) using a dataset  $\mathcal{D}$ , consisting of  $n$  tuples of the form  $(X, Y^{(1)}, Y^{(2)}, Z)$ . These data are generated as follows. At each time step, the LLM receives a *prompt*  $X$ . Then, the

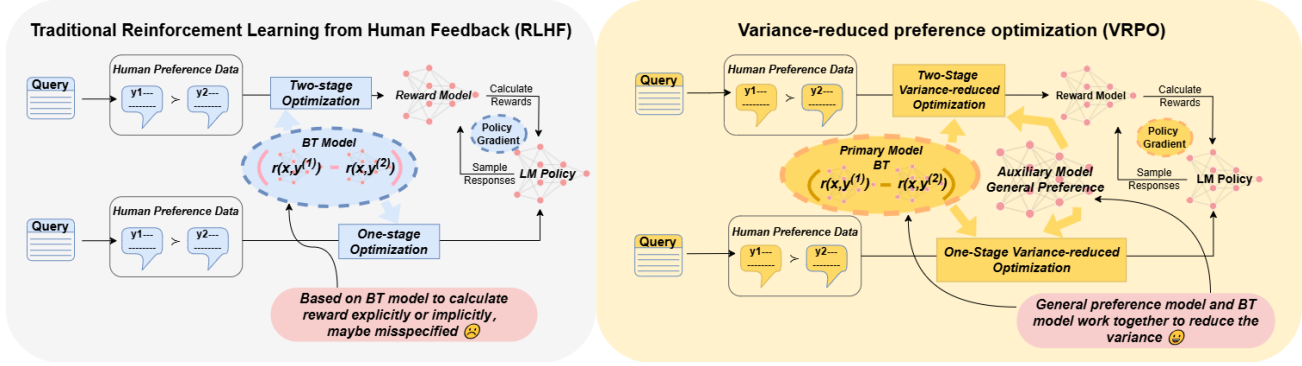


Figure 1. VRPO incorporates an auxiliary preference model to reduce the variance of the estimated primary model. **Left:** The classic one-stage and two-stage optimization schemes in RLHF. Both approaches require fitting a reward model, either explicitly or implicitly, which may lead to model misspecification. **Right:** In contrast, VRPO employs an auxiliary reward-free preference model to better capture human preferences. It works jointly with the primary model for variance reduction and policy improvement.

Table 1. Variance, MSE, and suboptimality gap: proposed VRPO vs. existing RLHF algorithms, where  $\downarrow$  represents a decrease and  $\rightarrow$  indicates no asymptotic difference.

MODEL SETTING	VARIANCE OF ESTIMATOR	MSE OF ESTIMATOR	SUBOPTIMALITY GAP
MISSPECIFIED	$\downarrow$	$\downarrow$	$\downarrow$
CORRECTLY SPECIFIED	$\rightarrow$	$\rightarrow$	$\rightarrow$

model next generates two different *responses* ( $Y^{(1)}, Y^{(2)}$ ) according to a reference policy  $\pi_{\text{ref}}$  so that  $Y^{(1)}, Y^{(2)} \sim \pi_{\text{ref}}(\cdot | X)$ . The prompt and responses ( $X, Y^{(1)}, Y^{(2)}$ ) are then presented to a human expert, and  $Z = \mathbb{I}(Y^{(2)} \succ Y^{(1)})$  represents the resulting *human feedback*. Here,  $Y^{(2)} \succ Y^{(1)}$  denotes that the second response  $Y^{(2)}$  is preferred, and  $\mathbb{I}$  denotes the indicator function.

We primarily focus on settings with a known or well-specified  $\pi_{\text{ref}}$ <sup>1</sup>. This assumption is automatically satisfied in the aforementioned scenario, where the responses are generated by the LLM we wish to fine-tune, either using the pre-trained policy or the supervised fine-tuning (SFT, Touvron et al., 2023) policy. When using external dataset for fine-tuning, this assumption can also be satisfied if well-trained reference policies are available, as is the case with the HH dataset.

**The model.** There are two crucial components under the RLHF framework: (i) A reward function  $r^*(X, Y)$  that measures the favorability of a response  $Y$  to a prompt  $X$ ; and (ii) a preference function  $p^*(X, Y^{(1)}, Y^{(2)}) = \mathbb{P}(Y^{(2)} \succ Y^{(1)} | X)$  that determines the probability of  $Y^{(2)}$  being preferred over  $Y^{(1)}$  given  $X$ . Our goal is to learn an optimal policy to generate the best response for each prompt, maximizing the expected reward. However, as mentioned in the

introduction, it remains highly challenging to ask human experts to directly provide numeric values to mimic the oracle reward  $r^*$ . Instead, RLHF learns the optimal policy through the preference model. Existing reward-based RLHF algorithms assume the preference  $p^*$  can be computed based on the difference between two rewards, such that

$$p^*(x, y^{(1)}, y^{(2)}) = \sigma(r^*(y^{(2)}, x) - r^*(y^{(1)}, x)), \quad (1)$$

for some activation function  $\sigma : \mathbb{R} \rightarrow [0, 1]$ . For example, in the BT model,  $\sigma$  is set to the sigmoid function.

Under a reward-based preference model assumption as in (1), these algorithms approximate  $r^*$  using a parametric model  $r_\theta$ , indexed by parameters  $\theta$ . This approximation, in turn, gives rise to the parametric preference model  $p_\theta(x, y^{(1)}, y^{(2)}) = \sigma(r_\theta(y^{(2)}, x) - r_\theta(y^{(1)}, x))$  for approximating the true preference function  $p^*$ .

**Parameter estimation.** The estimator  $\hat{\theta}$  is computed by minimizing the empirical risk  $\mathcal{L}(\theta) = \mathbb{E}_n[\ell(X, Y^{(1)}, Y^{(2)}, Z; \theta)]$  where  $\ell$  denotes certain loss function, e.g., the cross-entropy loss given by

$$\begin{aligned} \ell(x, y^{(1)}, y^{(2)}, z; \theta) = & -z \log(p_\theta(x, y^{(1)}, y^{(2)})) \\ & - (1 - z) \log(1 - p_\theta(x, y^{(1)}, y^{(2)})), \end{aligned} \quad (2)$$

and  $\mathbb{E}_n$  denotes the empirical average over the dataset  $\mathcal{D}$ . This leads to the estimated reward  $r_{\hat{\theta}}$ , which can then be

<sup>1</sup>“Well-specified” means the specified  $\pi_{\text{ref}}$  is very close to the ground truth.

used to derive the estimated optimal policy, for example, through PPO.

We remark that the consistency of the estimated policy relies on the correct specification of both the reward and preference models. However, model misspecification can lead to sub-optimal policies. There are two potential sources of misspecification. First, misspecification occurs when Equation (1) fails to hold, due to an incorrect activation function  $\sigma$ , or due to the fact that  $p^*$  cannot be represented as a function of the difference between two rewards. Second, misspecification occurs if  $r^*$  cannot be represented by the chosen parametric model  $r_\theta$  for any  $\theta$ . Our proposal, detailed in the next section, aims to improve policy learning under these misspecifications.

#### 4. Variance-Reduced Preference Optimization

In this section, we present our proposed algorithm and demonstrate how it can be integrated into existing reward-based RLHF algorithms to improve their sample efficiency. Our proposal estimates two models for the preference function  $p^*$ : (i) A primary, simpler reward-based model  $p_\theta$ , as in existing algorithms; and (ii) an auxiliary, more complex preference model  $p_\eta$ , designed to enhance the accuracy of the primary model. We remark that the auxiliary model is trained separately and ahead of the primary model, for example, by minimizing the empirical cross-entropy loss (see Equation (2)). It differs from the primary model in that it can either be non-reward-based, utilizing more general preference functions (e.g., Zhang et al., 2024b) to mitigate the misspecification of  $p^*$ , or remain reward-based but involving a larger number of parameters to alleviate the misspecification of  $r^*$ . By design, this more complex model is expected to be more robust to model misspecification compared to the primary model.

After training the auxiliary model, we estimate the parameter  $\theta$  in the primary model by minimizing the following loss function

$$\begin{aligned} \tilde{\mathcal{L}}(\theta) = & \mathbb{E}_n \left[ \ell(X, Y^{(1)}, Y^{(2)}, Z; \theta) \right. \\ & - \sum_{u=0}^1 \ell(X, Y^{(1)}, Y^{(2)}, u; \theta) p_\eta(X, Y^{(1)}, Y^{(2)}, u) \\ & + \sum_{u=0}^1 \mathbb{E}_{y^{(1)*}, y^{(2)*} \sim \pi_{\text{ref}}(\cdot | X)} \ell(X, y^{(1)*}, y^{(2)*}, u; \theta) \\ & \left. \times p_\eta(X, y^{(1)*}, y^{(2)*}, u) \right], \end{aligned} \quad (3)$$

By definition, the above objective function consists of three terms. The first term

$\mathbb{E}_n[\ell(X, Y^{(1)}, Y^{(2)}, Z; \theta)]$  is the loss function of existing reward-based RLHF algorithms. The second and third terms

correspond to loss functions constructed using data generated by the auxiliary model  $p_\eta$  and the reference policy  $\pi_{\text{ref}}$ . The primary roles of the last two terms are to reduce the variance while ensuring unbiasedness. Specifically, if  $\pi_{\text{ref}}$  is correctly specified, the last two terms share the same mean, and their expectations cancel out. Meanwhile, when  $p_\eta$  is close to the oracle preference function  $p^*$ , including the last two terms reduces the variance of the objective function. We will formally show in Section 5 that such a variance reduction in the loss function leads to a reduction in the variance and MSE of the resulting estimator  $\tilde{\theta}$  as well as the suboptimality gap of the estimated policy. Finally, we illustrate how our proposal can be applied to both two-stage and one-stage optimization described in Section 2.

**Two-stage optimization.** In two-stage optimization, we directly model the reward as  $r_\theta$  and specify an activation function  $\sigma$ , leading to the preference model  $p_\theta(x, y^{(1)}, y^{(2)}) = \sigma(r_\theta(y^{(2)}, x) - r_\theta(y^{(1)}, x))$ . We next choose a proper loss function for  $\ell$ , such as the cross-entropy loss or hinge loss, and compute the estimated parameter  $\tilde{\theta}$  by minimizing (3). Finally, we apply a standard RL algorithm (e.g., PPO) to estimate the optimal policy by solving

$$\begin{aligned} \arg \max_{\pi} \mathbb{E}_n \left[ \mathbb{E}_{y \sim \pi(\cdot | X)} [r_{\tilde{\theta}}(X, y)] \right. \\ \left. - \beta \cdot \text{KL}(\pi(\cdot | X) || \pi_{\text{ref}}(\cdot | X)) \right], \end{aligned} \quad (4)$$

for some tuning parameter  $\beta > 0$ , where KL denotes the Kullback-Leibler divergence.

**One-stage optimization.** In one-stage optimization, we parameterize the optimal policy via  $\pi_\theta$ . Solving (4) yields the following close-form expression for the reward model

$$r_\theta(x, y) = \beta \log \left( \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)} \right) - C(x), \quad (5)$$

where  $C(x)$  denotes some baseline term that is independent of the response. We next employ this reward to model the preference, and minimize the resulting modified loss function (3) for estimating the policy parameter  $\theta$ .

#### 5. Theoretical results

We establish statistical properties of the proposed estimator and policy in this section. We study two settings. The first setting is a correctly specified setting where the preference function  $p^*$  satisfies (1) and the reward model is correctly specified, i.e.,  $r^* = r_{\bar{\theta}}$  for some  $\bar{\theta}$ . Notice that for one-stage optimization, this requires  $r^*$  to be equivalent to  $\beta \log(\pi_{\bar{\theta}}/\pi_{\text{ref}})$ , up to a baseline term. The second setting is a misspecified setting where either  $p^*$  does not satisfy (1) or the reward model is misspecified. In this case, we define the target parameter  $\bar{\theta} = \arg \min_{\bar{\theta}} \mathbb{E}[\mathcal{L}(\bar{\theta})]$  as the minimizer within the misspecified model class.



### 5.1. Variance and MSE Reduction Guarantees

Recall that our loss function is given by  $\tilde{\mathcal{L}}(\theta)$  in Equation (3), which differs from the conventional loss function  $\mathcal{L}(\theta)$  used in existing RLHF algorithms. Nonetheless, the following theorem shows that  $\tilde{\theta}$  remains identifiable from our modified loss function under a double robustness condition.

**Theorem 5.1** (Double Robustness). *In the correctly specified setting, the target parameter  $\tilde{\theta} = \arg \min_{\theta} \mathbb{E}[\tilde{\mathcal{L}}(\theta)]$ , when either the reference policy  $\pi_{\text{ref}}$  or the auxiliary preference model  $p_{\eta}$  is correctly specified.*

Theorem 5.1 establishes an *asymptotic* guarantee for the proposed estimator, proving its consistency assuming we have access to infinitely many samples so that the loss function  $\tilde{\mathcal{L}}(\theta)$  converges to its expectation. This theorem shows that our framework is *doubly robust*, in the sense that the proposed estimator remains identifiable when either one (but not necessarily both) of the two models is correctly specified. In the following, we focus on settings with a known  $\pi_{\text{ref}}$  and derive the *finite-sample* property of our estimator and policy.

Recall that  $\tilde{\theta}$  and  $\hat{\theta}$  denote our estimator and the baseline estimator that minimizes  $\mathcal{L}(\theta)$ , respectively. Let  $d$  denote their dimension.

**Theorem 5.2** (Variance and MSE reductions). *Under the regularity conditions in Appendix A.2, regardless of whether the model is correctly specified or misspecified, we have*

$$\|\mathbb{E}(\hat{\theta}) - \tilde{\theta}\|_2 = O\left(\frac{d}{n\lambda_{\min}}\right), \|\mathbb{E}(\tilde{\theta}) - \tilde{\theta}\|_2 = O\left(\frac{d}{n\lambda_{\min}}\right),$$

and

$$\begin{aligned} \text{Var}(\hat{\theta}) - \text{Var}(\tilde{\theta}) &= \underbrace{\frac{1}{n} A^{-1}(\tilde{\theta}) H A^{-1}(\tilde{\theta})}_{\text{variance reduction}} + \underbrace{O\left(\frac{d^{3/2}}{n^{3/2}\lambda_{\min}^2}\right)}_{\text{nonlinear residual}} \\ &\quad + \underbrace{O\left(\frac{\|p_{\eta} - p^*\|_{\infty}^2}{n\lambda_{\min}^2}\right)}_{\text{model specification error}}, \end{aligned}$$

where  $H$  denotes certain positive semi-definite matrix whose detailed form is specified in Appendix A.2,  $A(\tilde{\theta}) := -\mathbb{E}\left\{\frac{\partial^2}{\partial\theta^2}\mathcal{L}(y, A, x; \tilde{\theta})\right\}$ ,  $\lambda_{\min}$  denote the minimum eigenvalue of  $A(\tilde{\theta})$  and  $\|p_{\eta} - p^*\|_{\infty}$  denotes the difference between  $p_{\eta}$  and  $p^*$  in supremum norm. Moreover, the minimum eigenvalue of the variance reduction term is lower bounded by

$$\begin{aligned} &\frac{c}{\lambda_{\max}^2} \mathbb{E}\left\{\mathbb{E}_{\substack{y^{(1)}, y^{(2)} \sim \pi_{\text{ref}}(\cdot|X) \\ y^{(1)*}, y^{(2)*} \sim \pi_{\text{ref}}(\cdot|X)}} \left\|\mathbb{E}_Z \left[\nabla_{\theta} \mathcal{L}(Z, y^{(1)}, y^{(2)}, X)\right] \right.\right. \\ &\quad \left.\left. - \mathbb{E}_Z \left[\nabla_{\theta} \mathcal{L}(Z, y^{(1)*}, y^{(2)*}, X)\right]\right\|_2^2\right\} \end{aligned} \quad (6)$$

where  $c$  denotes some universal constant and  $\mathbb{E}_Z$  denotes the conditional expectation of  $Z$  given the two responses and  $X$ .

The first part of Theorem 5.2 upper bounds the biases of the two estimators, and decomposes the difference in their covariance matrices into three terms: (i) a non-negative (due to the positive semi-definiteness of  $H$ ) *variance reduction term*, representing the difference in asymptotic covariance matrices; (ii) a high-order *nonlinear residual term* which arises due to the nonlinearity of the preference model and scales as  $O(n^{-3/2})$ ; (iii) a *model specification error term* capturing the approximation error between the auxiliary model  $p_{\eta}$  and the oracle preference function  $p^*$ .

Notice that the last two terms are high-order terms as functions of the sample size  $n$ , provided that the approximation error  $p_{\eta} - p^*$  decays to zero as  $n$  approaches infinity. As such, the difference in covariance matrices is dominated by the non-negative variance reduction term, formally demonstrating the proposed algorithm’s variance reduction effect. Moreover, the bias bounds for both estimators are of the same order and asymptotically negligible compared to the standard deviation. Together, these results prove that our estimator achieves a lower MSE than the baseline estimator.

To further quantify our estimator’s improvement, the second part of Theorem 5.2 derives a lower bound for the variance reduction term. Two cases are worth noting. First, when the reference policy is deterministic, the response pairs  $(y^{(1)}, y^{(2)})$  and  $(y^{(1)*}, y^{(2)*})$  are identical, leading to a zero lower bound. Second, under correct model specification, the expected gradient of the loss function becomes zero, similarly resulting in a zero lower bound. In either case, it can be shown that the variance reduction term equals zero, implying no improvement. In practice, however, we typically use a stochastic reference policy to encourage exploration, and the model is likely misspecified. Consequently, our algorithm often achieves non-trivial improvement, as demonstrated by our numerical experiments.

### 5.2. Suboptimality Guarantees

Finally, we demonstrate that the variance reduction directly translates to a smaller suboptimality gap in the resulting policy. For a given model parameter  $\theta$ , we denote  $J(\theta)$  as the expected reward  $\mathbb{E}[r^*(X, Y)]$  where the response  $Y$  is generated by the resulting induced policy. Let  $J^*$  denote the expected reward under the optimal policy. We define the suboptimality gap  $R(\theta) = J^* - J(\theta)$ .

**Theorem 5.3** (Reduction in suboptimality gap). *Under the same condition in Theorem 5.2 and additional regularity conditions in Appendix A.3, regardless of whether the model*

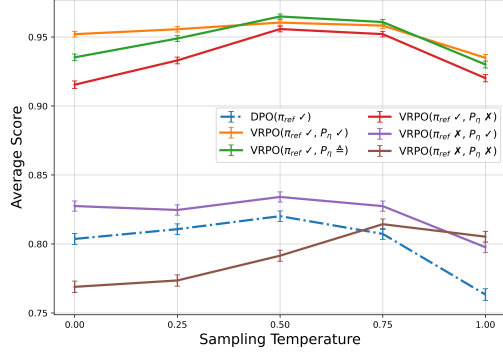


Figure 2. Expected reward in different VRPO settings compared to DPO; for example,  $(\pi_{\text{ref}} \checkmark, P_{\eta} \times)$  indicates the reference model is correctly specified and the preference model is misspecified, while  $P_{\eta} \hat{=}$  means the preference model is estimated, demonstrating the robustness of our method.

is correctly specified or misspecified, we have

$$\begin{aligned} \mathbb{E} \{R(\hat{\theta})\} &= R(\bar{\theta}) + \text{trace} \left( \text{Var}(\hat{\theta}) (-\nabla_{\bar{\theta}}^2 J(\bar{\theta})) \right) \\ &\quad + O \left( \frac{d^{3/2}}{n^{3/2} \lambda_{\min}^3} \right), \\ \mathbb{E} \{R(\tilde{\theta})\} &= R(\bar{\theta}) + \text{trace} \left( \text{Var}(\tilde{\theta}) (-\nabla_{\bar{\theta}}^2 J(\bar{\theta})) \right) \\ &\quad + O \left( \frac{d^{3/2}}{n^{3/2} \lambda_{\min}^3} \right), \end{aligned}$$

where  $\text{Var}(\hat{\theta})$  and  $\text{Var}(\tilde{\theta})$  denote the covariance matrices of  $\hat{\theta}$  and  $\tilde{\theta}$ , respectively.

Theorem 5.3 decomposes the suboptimality gap into three terms. The first term  $R(\bar{\theta})$  represents the suboptimality gap under the policy induced by  $\bar{\theta}$ . This term arises due to the potential misspecification of the primary model. Under correct model misspecification, this term equals zero. The last two terms arise from the variance and bias of the estimator, respectively.

We highlight three key observations: (i) The first term is the same for both the baseline algorithm and our proposed algorithm; (ii) As shown in Theorem (5.2), our variance is asymptotically smaller than that of the baseline estimator; (iii) The last bias term is of the same order  $O(n^{-3/2})$  for two estimators and negligible compared to the variance term. Together, these results rigorously demonstrate that reducing the estimator’s variance directly translates to a smaller suboptimality gap in the learned policy.

## 6. Experiments

In this section, we empirically evaluate the performance of the proposed VRPO in complex RLHF tasks, including

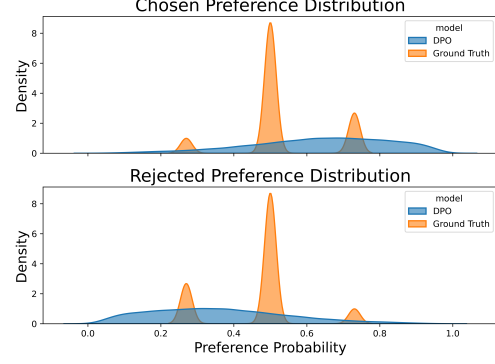


Figure 3. The difference in preference probability distributions between the ground truth and the DPO estimation for both the Chosen and Rejected responses.

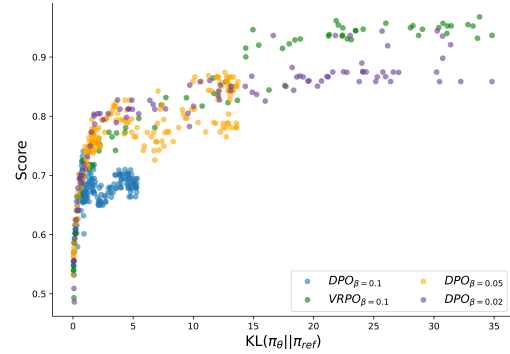


Figure 4. Expected reward versus KL-divergence under VRPO ( $\beta = 0.1$ ) and DPO ( $\beta \in \{0.02, 0.05, 0.1\}$ ), demonstrating the quality of the optimization.

sentiment generation (Section 6.1), dialogue and summarization (Section 6.2). These tasks are widely used in the literature for evaluation (Wu et al., 2024a; Rafailov et al., 2023; Lu et al., 2025; Xiao et al., 2025b).

**Methods.** We compare our proposed VRPO against SFT and direct preference optimization (DPO, Rafailov et al., 2023) in the main text. The SFT policy is trained by token-wise supervised fine-tuning (Rafailov et al., 2023; Wu et al., 2024a). As for DPO, it is trained upon the SFT model by employing the SFT policy as the initial policy. To ensure fair comparison, the proposed VRPO employs the same model architecture as DPO for its primary model. Implementation details and extra results are provided in Appendix B. We also implement a two-stage VRPO algorithm, based on the classic RLHF algorithm (Ouyang et al., 2022), and compare against this baseline in Appendix C.2.

### 6.1. Synthetic Data Analysis under Correct Specification

**Task and data.** We consider the *controlled sentiment generation* task, in which the objective is to steer a pre-

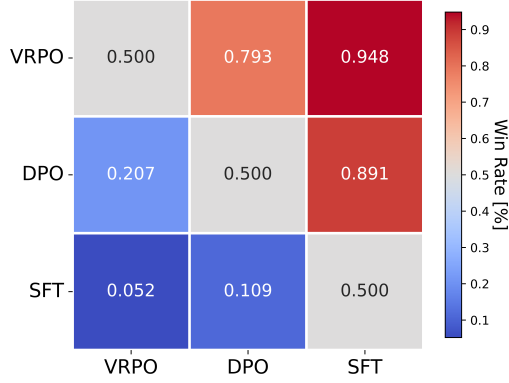


Figure 5. Head-to-head comparisons between VRPO, DPO, and SFT on Anthropic-HH one-step dialogue task. Win rates are evaluated by GPT-4o-mini

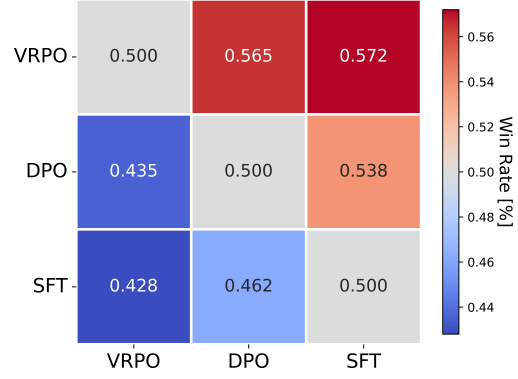


Figure 6. Head-to-head comparisons between VRPO, DPO, and SFT on the summarization task. Win rates are evaluated by GPT-4o-mini.

trained language model to produce positive movie reviews based the IMDb dataset (Maas et al., 2011). Specifically, we use movie review sentence prefixes as prompts, and generate two responses from a SFT model trained from Qwen/Qwen2-0.5B (Yang et al., 2024). Rewards for each response are annotated based on a pre-trained sentiment classifier, and preference labels between responses are subsequently simulated using the BT model. For this task, the reward-based preference condition (see (1)) is satisfied and the Qwen/Qwen2-0.5B model serves as the primary model.

**Evaluation criterion.** Since preferences in this task are synthetically generated from a known reward function, we evaluate different methods by the expected reward of their generated responses.

**Result I: double-robust property.** As mentioned earlier, our synthetic environment satisfies the reward-based preference condition by construction. Under correct (or approximately correct) reward model specification, the setup closely mirrors the ideal setting without model-misspecification. This allows us to empirically validate the double robustness property established in Theorem 5.1. Specifically, we implement four versions of VRPO, depending on whether  $p_\eta$  and  $\pi_{\text{ref}}$  are correctly specified, and report their expected rewards in Figure 2. The oracle expected reward is 1. It can be seen that VRPO with both correctly specified  $p_\eta$  and  $\pi_{\text{ref}}$  (orange line) achieves the highest expected reward (around 0.95) across most temperatures. When either one model is correctly specified, the expected reward varies from 0.8 to 0.95 (red and purple lines). Finally, the expected reward drops to below 80% when both models are misspecified (brown line).

**Result II: variance reduction property.** We implement DPO (blue line) and another variant of VRPO (green line) that uses DPO’s estimated reward to construct the auxiliary

model, and report their expected rewards in Figure 2. While theoretically VRPO should perform comparably to DPO, it achieves superior empirical performance. We suspect this improvement is due to DPO’s poor preference estimation - Figure 3 shows DPO’s fitted preference distribution deviates substantially from the true distribution. This highlights VRPO’s advantage: even under correct model misspecification, the proposed VRPO can still improve existing algorithms if they do not learn well.

The primary goal in RLHF is to maximize the reward while constraining the policy’s divergence from the reference policy  $\pi_{\text{ref}}$ . Ideally, a superior RLHF algorithm shall achieve higher expected reward at each fixed KL-divergence level between the learned policy and  $\pi_{\text{ref}}$ . It can be seen from Figure 4 that, with the tuning parameter  $\beta$  set to 0.1 (see (4) and (5)), VRPO generally attains higher rewards than DPO at each fixed KL-divergence level. Since DPO with  $\beta = 0.1$  struggles to learn substantially different policies from  $\pi_{\text{ref}}$ , we additionally evaluate DPO by setting  $\beta$  to 0.05 and 0.02. In this extended comparison, VRPO maintains its advantage, achieving higher rewards at comparable or lower KL-divergence levels.

## 6.2. Real Data Analysis under Model Misspecification

**Tasks and datasets.** We consider two tasks in this section. The first one is *summarization*, which aims to fine-tune language models to generate concise yet informative summaries from lengthy input texts. We utilize the TL;DR (“Too Long; Didn’t Read”) dataset (Völske et al., 2017), which contains different Reddit posts with preference annotations collected by Stiennon et al. (2020). For this task, we employ the trl-lib/pythia-1b-deduped-tldr-sft as the reference model. The second task is *single-turn dialogue*. For this task, the prompts range from scientific questions to emotional support suggestions, and the objective is to generate responses that are factually accurate, emotionally

Table 2. Head-to-head comparison between SFT, DPO, DPO-Hinge, and VRPO at five sampling temperatures.

Model v.s. Model	0.0	0.25	0.5	0.75	1.0
DPO v.s. SFT	0.89	0.85	0.84	0.81	0.81
VRPO v.s. DPO	0.79	0.81	0.81	0.80	0.78
VRPO v.s. SFT	0.95	0.95	0.92	0.92	0.93
DPO-Hinge v.s. SFT	0.89	0.85	0.86	0.87	0.83
DPO-Hinge v.s. DPO	0.42	0.44	0.46	0.40	0.44
VRPO v.s. DPO-Hinge	0.83	0.83	0.84	0.84	0.81

appropriate, and closely aligned with human expectations. We adopt the widely used HH dataset, consisting of dialogues between a human and an automated assistant with preference labels for the last assistant answer, and train the SFT policy based on the Qwen/Qwen2.5-1.5B model.

**Evaluation criterion.** In the absence of ground-truth preference models for both tasks, we follow Zhang et al. (2024b) in using GPT-4o-mini’s preferences to assess the quality of the generated responses. This approach of evaluation is supported by empirical evidence that LLM-based assessments frequently outperform traditional evaluation metrics (Chen et al., 2023).

**Results.** Figure 5 and Figure 6 show head-to-head comparisons among the proposed VRPO (using the general preference model Zhang et al., 2024b, as the auxiliary model) and the two benchmark methods. Specifically, for any of the two methods, we: (i) collect their generated responses, (ii) obtain GPT preference, and (iii) compute the win rates - the percentage where one method’s responses are preferred over the other’s. It can be seen that VRPO outperforms both DPO and SFT on both tasks. Notably, VRPO demonstrates particularly strong performance on the HH dataset, achieving a win rate of approximately 95% over the SFT baseline and about 80% over DPO. To the contrary, the improvements on the TL;DR dataset are modest, likely because the reference model we use is misspecified – it was trained exclusively on human-preferred summaries.

To assess robustness, we evaluate the performance of VRPO and the two baseline methods across five sampling temperatures: 0.0, 0.25, 0.5, 0.75, 1.0. For each temperature, we conduct head-to-head comparisons. The results are exhibited in Table 2. Furthermore, we conduct an additional evaluation against the chosen responses in the HH dataset, with the win rates summarized in Appendix C.1. As an ablation study, we implement DPO-Hinge (Zhao et al., 2023; Tang et al., 2024), which replaces DPO’s cross-entropy loss with hinge loss. This comparison is relevant when applied to the HH dataset since we employ hinge loss in VRPO to prevent gradient explosion encountered with cross-entropy loss. It can be seen that VRPO consistently achieves higher win rates across all temperatures. Notably, responses generated by VRPO are preferred over those generated by DPO

77–81% of the time. Meanwhile, DPO outperformed DPO-Hinge, suggesting that the superior performance of VRPO is attributed to its variance reduction property rather than the specific choice of loss function.

## 7. Conclusion

Due to the inherent complexity of human preferences, mathematical models used to characterize these preferences are likely to be misspecified in practice. To mitigate the adverse effects of model misspecification, this paper introduces a simple and general approach for fine-tuning pretrained language models. This approach can efficiently reduce variance when model is misspecified, as evidenced by both theory and experiments. The variance reduction stems from the variance reduction technique in statistics and machine learning literature (Tsiatis, 2006; Kallus & Uehara, 2020). A limitation of this paper is that it just focuses on the reward-based RLHF methods. Given the promising theoretical and empirical findings presented, extending this variance-reduction framework to preference-based RLHF methods would be a promising direction for future researches.

## Ethics Statement

This work does not involve the collection of new human data. All experiments were conducted on publicly available, anonymized datasets such as the TL;DR, Anthropic Helpful and Harmless benchmark, which were originally collected with appropriate consent and ethical oversight. Our research focuses on improving the robustness and sample efficiency of reinforcement learning from human feedback (RLHF) algorithms under model misspecification, and does not engage in any direct interaction with human participants. We strictly adhere to the ethical use of human-generated data, ensuring compliance with data usage licenses and responsible AI practices. Furthermore, as our proposed methods aim to align large language models more effectively with human values, this work contributes to the development of safer and more aligned AI systems.



## References

- Agranov, M. and Ortoleva, P. Stochastic choice and preferences for randomization. *Journal of Political Economy*, 125:40 – 68, 2015. URL <https://api.semanticscholar.org/CorpusID:11440933>.
- Azar, M. G., Guo, Z. D., Piot, B., Munos, R., Rowland, M., Valko, M., and Calandriello, D. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, pp. 4447–4455. PMLR, 2024.
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.
- Bakker, M., Chadwick, M., Sheahan, H., Tessler, M., Campbell-Gillingham, L., Balaguer, J., McAleese, N., Glaese, A., Aslanides, J., Botvinick, M., et al. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35:38176–38189, 2022.
- Bertrand, Q., Czarnecki, W. M., and Gidel, G. On the limitations of the elo, real-world games are transitive, not additive. In Ruiz, F., Dy, J., and van de Meent, J.-W. (eds.), *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 2905–2921. PMLR, 25–27 Apr 2023. URL <https://proceedings.mlr.press/v206/bertrand23a.html>.
- Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Bukharin, A., Hong, I., Jiang, H., Li, Z., Zhang, Q., Zhang, Z., and Zhao, T. Robust reinforcement learning from corrupted human feedback. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=cR2QDzdpEv>.
- Calandriello, D., Guo, Z. D., Munos, R., Rowland, M., Tang, Y., Pires, B. A., Richemond, P. H., Le Lan, C., Valko, M., Liu, T., Joshi, R., Zheng, Z., and Piot, B. Human alignment of large language models through online preference optimisation. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- Chen, Y., Wang, R., Jiang, H., Shi, S., and Xu, R. Exploring the use of large language models for reference-free text quality evaluation: An empirical study. In *Findings of the Association for Computational Linguistics: IJCNLP-AACL 2023 (Findings)*, pp. 361–374, Nusa Dua, Bali, November 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-ijcnlp.32.
- Cheng, Y., Yao, F., Liu, X., and Xu, H. Learning from imperfect human feedback: A tale from corruption-robust dueling. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=ptjrpEGrGg>.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Das, N., Chakraborty, S., Pacchiano, A., and Chowdhury, S. R. Active preference optimization for sample efficient rlhf. *Proceedings of the 41th International Conference on Machine Learning*, 2024.
- Freedman, R., Svegliato, J., Wray, K., and Russell, S. Active teacher selection for reinforcement learning from human feedback. *arXiv preprint arXiv:2310.15288*, 2023.
- Gardner, M. *Mathematical games, the paradox of the nontransitive dice and the elusive principle of indifference*. Dec 1970. URL <https://www.scientificamerican.com/article/mathematical-games-1970-12/>.
- Glaese, A., McAleese, N., Trębacz, M., Aslanides, J., Firoiu, V., Ewalds, T., Rauh, M., Weidinger, L., Chadwick, M., Thacker, P., Campbell-Gillingham, L., Uesato, J., Huang, P.-S., Comanescu, R., Yang, F., See, A., Dathathri, S., Greig, R., Chen, C., Fritz, D., Elias, J. S., Green, R., Mokrá, S., Fernando, N., Wu, B., Foley, R., Young, S., Gabriel, I., Isaac, W., Mellor, J., Hassabis, D., Kavukcuoglu, K., Hendricks, L. A., and Irving, G. Improving alignment of dialogue agents via targeted human judgements, 2022. URL <https://arxiv.org/abs/2209.14375>.
- Hao, B., Jain, R., Lattimore, T., Van Roy, B., and Wen, Z. Leveraging demonstrations to improve online learning: Quality matters. In *International Conference on Machine Learning*, pp. 12527–12545. PMLR, 2023.
- Hartmann, J., Heitmann, M., Siebert, C., and Schamp, C. More than a feeling: Accuracy and application of sentiment analysis. *International Journal of Research in Marketing*, 40(1):75–87, 2023. doi: <https://doi.org/10.1016/j.ijresmar.2022.05.005>. URL <https://www.sciencedirect.com/science/article/pii/S0167811622000477>.

- Hejna, J., Rafailov, R., Sikchi, H., Finn, C., Niekum, S., Knox, W. B., and Sadigh, D. Contrastive preference learning: Learning from human feedback without reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=iX1RjVQODj>.
- Hong, I., Li, Z., Bukharin, A., Li, Y., Jiang, H., Yang, T., and Zhao, T. Adaptive preference scaling for reinforcement learning with human feedback. *Advances in Neural Information Processing Systems*, 37:107249–107269, 2024.
- Kallus, N. and Uehara, M. Double reinforcement learning for efficient and robust off-policy evaluation. In *International Conference on Machine Learning*, pp. 5078–5088. PMLR, 2020.
- Lee, S. J., Sun, W. W., and Liu, Y. Low-rank contextual reinforcement learning from heterogeneous human feedback. *arXiv preprint arXiv:2412.19436*, 2024.
- Li, Z., Xu, T., Zhang, Y., Lin, Z., Yu, Y., Sun, R., and Luo, Z.-Q. Remax: a simple, effective, and efficient reinforcement learning method for aligning large language models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML’24, 2024.
- Lindner, D. and El-Assady, M. Humans are not boltzmann distributions: Challenges and opportunities for modelling human feedback and interaction in reinforcement learning. *ArXiv*, abs/2206.13316, 2022. URL <https://api.semanticscholar.org/CorpusID:250072570>.
- Liu, K., Long, Q., Shi, Z., Su, W. J., and Xiao, J. Statistical impossibility and possibility of aligning llms with human preferences: From condorcet paradox to nash equilibrium. *arXiv preprint arXiv:2503.10990*, 2025.
- Liu, P., Shi, C., and Sun, W. W. Dual active learning for reinforcement learning from human feedback. *arXiv preprint arXiv:2410.02504*, 2024a.
- Liu, T., Zhao, Y., Joshi, R., Khalman, M., Saleh, M., Liu, P. J., and Liu, J. Statistical rejection sampling improves preference optimization. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=xbjSwwrQOe>.
- Lu, C., Holt, S., Fanconi, C., Chan, A., Foerster, J., van der Schaar, M., and Lange, R. Discovering preference optimization algorithms with and for large language models. *Advances in Neural Information Processing Systems*, 37: 86528–86573, 2025.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT ’11, pp. 142–150, USA, 2011. Association for Computational Linguistics. ISBN 9781932432879.
- Mandal, D., Nika, A., Kamalaruban, P., Singla, A., and Radanović, G. Corruption robust offline reinforcement learning with human feedback. *arXiv preprint arXiv:2402.06734*, 2024.
- Mandal, D., Sasnauskas, P., and Radanovic, G. Distributionally robust reinforcement learning with human feedback. *arXiv preprint arXiv:2503.00539*, 2025.
- May, K. O. Intransitivity, utility, and the aggregation of preference patterns. *Econometrica*, 22:1, 1954. URL <https://api.semanticscholar.org/CorpusID:156169619>.
- Meng, Y., Xia, M., and Chen, D. Simpo: Simple preference optimization with a reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235, 2025.
- Michaud, E. J., Gleave, A., and Russell, S. J. Understanding learned reward functions. *ArXiv*, abs/2012.05862, 2020. URL <https://api.semanticscholar.org/CorpusID:228083849>.
- Milano, S., Taddeo, M., and and, L. F. Ethical aspects of multi-stakeholder recommendation systems. *The Information Society*, 37(1):35–45, 2021. doi: 10.1080/01972243.2020.1832636. URL <https://doi.org/10.1080/01972243.2020.1832636>.
- Mukherjee, S., Lalitha, A., Kalantari, K., Deshmukh, A. A., Liu, G., Ma, Y., and Kveton, B. Optimal design for human preference elicitation. In Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., and Zhang, C. (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 90132–90159. Curran Associates, Inc., 2024.
- Munos, R., Valko, M., Calandriello, D., Gheshlaghi Azar, M., Rowland, M., Guo, Z. D., Tang, Y., Geist, M., Mesnard, T., Fiegel, C., Michi, A., Selvi, M., Girgin, S., Momchev, N., Bachem, O., Mankowitz, D. J., Precup, D., and Piot, B. Nash learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36743–36768. PMLR, 21–27 Jul 2024.

- OpenAI. Gpt-4 technical report. 2023. URL <https://api.semanticscholar.org/CorpusID:257532815>.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 53728–53741. Curran Associates, Inc., 2023.
- Rafailov, R., Hejna, J., Park, R., and Finn, C. From  $\$r$  to  $\$q^*$ : Your language model is secretly a q-function. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=kEVcNxtqXk>.
- Ramesh, S. S., Hu, Y., Chaimalas, I., Mehta, V., Sessa, P. G., Bou Ammar, H., and Bogunovic, I. Group robust preference optimization in reward-free rlhf. *Advances in Neural Information Processing Systems*, 37:37100–37137, 2024.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D. M., Lowe, R., Voss, C., Radford, A., Amodei, D., and Christiano, P. Learning to summarize from human feedback. In *NeurIPS*, 2020.
- Sun, H., Shen, Y., and Ton, J.-F. Rethinking reward modeling in preference-based large language model alignment. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=rfdble10qm>.
- Sutton, R. S., Barto, A. G., et al. *Reinforcement learning: An introduction*. MIT press Cambridge, 2018.
- Swamy, G., Dann, C., Kidambi, R., Wu, Z. S., and Agarwal, A. A minimaximalist approach to reinforcement learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning*, 2024.
- Tang, Y., Guo, Z. D., Zheng, Z., Calandriello, D., Munos, R., Rowland, M., Richemond, P. H., Valko, M., Avila Pires, B., and Piot, B. Generalized preference optimization: A unified approach to offline alignment. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 47725–47742. PMLR, 21–27 Jul 2024.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models, 2023. URL <https://arxiv.org/abs/2302.13971>.
- Tsiatis, A. A. *Semiparametric theory and missing data*, volume 4. Springer, 2006.
- Tversky, A. Intransitivity of preferences. *Psychological Review*, 76:31–48, 1969. URL <https://api.semanticscholar.org/CorpusID:144609998>.
- Vojnovic, M. and Yun, S.-Y. What is the alignment objective of grpo? *arXiv preprint arXiv:2502.18548*, 2025.
- Völske, M., Potthast, M., Syed, S., and Stein, B. Tldr: Mining reddit to learn automatic summarization. In *NFiS@EMNLP*, 2017. URL <https://api.semanticscholar.org/CorpusID:2204603>.
- von Werra, L., Belkada, Y., Tunstall, L., Beeching, E., Thrush, T., Lambert, N., Huang, S., Rasul, K., and Gallouédec, Q. Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl>, 2020.
- Wang, C., Jiang, Y., Yang, C., Liu, H., and Chen, Y. Beyond reverse KL: Generalizing direct preference optimization with diverse divergence constraints. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=2cRzmWXX9N>.
- Wang, J., Wang, H., Sun, S., and Li, W. Aligning language models with human preferences via a bayesian approach. *Advances in Neural Information Processing Systems*, 36: 49113–49132, 2023.
- Wu, T., Zhu, B., Zhang, R., Wen, Z., Ramchandran, K., and Jiao, J. Pairwise proximal policy optimization: Language model alignment with comparative RL. In *First Conference on Language Modeling*, 2024a.

- Wu, Y., Sun, Z., Yuan, H., Ji, K., Yang, Y., and Gu, Q. Self-play preference optimization for language model alignment. *NeurIPS (2024) Workshop on Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*, 2024b.
- Xiao, T., Yuan, Y., Li, M., Chen, Z., and Honavar, V. G. On a connection between imitation learning and RLHF. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=2QdsjiNXgj>.
- Xiao, T., Yuan, Y., Zhu, H., Li, M., and Honavar, V. Cal-dpo: Calibrated direct preference optimization for language model alignment. *Advances in Neural Information Processing Systems*, 37:114289–114320, 2025b.
- Xiong, W., Dong, H., Ye, C., Wang, Z., Zhong, H., Ji, H., Jiang, N., and Zhang, T. Iterative preference learning from human feedback: Bridging theory and practice for RLHF under KL-constraint. In Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F. (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 54715–54754. PMLR, 21–27 Jul 2024.
- Yang, A., Yang, B., Hui, B., Zheng, B., Yu, B., Zhou, C., Li, C., Li, C., Liu, D., Huang, F., Dong, G., Wei, H., Lin, H., Tang, J., Wang, J., Yang, J., Tu, J., Zhang, J., Ma, J., Yang, J., Xu, J., Zhou, J., Bai, J., He, J., Lin, J., Dang, K., Lu, K., Chen, K., Yang, K., Li, M., Xue, M., Ni, N., Zhang, P., Wang, P., Peng, R., Men, R., Gao, R., Lin, R., Wang, S., Bai, S., Tan, S., Zhu, T., Li, T., Liu, T., Ge, W., Deng, X., Zhou, X., Ren, X., Zhang, X., Wei, X., Ren, X., Liu, X., Fan, Y., Yao, Y., Zhang, Y., Wan, Y., Chu, Y., Liu, Y., Cui, Z., Zhang, Z., Guo, Z., and Fan, Z. Qwen2 technical report. *CoRR*, abs/2407.10671, 2024. URL <https://doi.org/10.48550/arXiv.2407.10671>.
- Ye, C., Xiong, W., Zhang, Y., Dong, H., Jiang, N., and Zhang, T. Online iterative reinforcement learning from human feedback with general preference model. *Advances in Neural Information Processing Systems*, 37: 81773–81807, 2024.
- Zhang, H., Lei, Y., Gui, L., Yang, M., He, Y., Wang, H., and Xu, R. CPPO: Continual learning for reinforcement learning with human feedback. In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=86zAUE80pP>.
- Zhang, Y., Zhang, G., Wu, Y., Xu, K., and Gu, Q. General preference modeling with preference representations for aligning language models. *arXiv preprint arXiv:2410.02197*, 2024b.
- Zhang, Y., Yu, D., Peng, B., Song, L., Tian, Y., Huo, M., Jiang, N., Mi, H., and Yu, D. Iterative nash policy optimization: Aligning LLMs with general preferences via no-regret learning. In *The Thirteenth International Conference on Learning Representations*, 2025.
- Zhao, Y., Khalman, M., Joshi, R., Narayan, S., Saleh, M., and Liu, P. J. Calibrating sequence likelihood improves conditional language generation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=0qSOodKmJaN>.
- Zhu, B., Jordan, M., and Jiao, J. Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 43037–43067. PMLR, 23–29 Jul 2023.
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., and Irving, G. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.



## A. Technical Proof

### A.1. Proof of Theorem 5.1

*Proof.* Recall that our modified loss function can be written as

$$\begin{aligned}\tilde{\mathcal{L}}(z, y, x; \theta) &= \mathbb{E}_n \left[ \ell(z, y, x; \theta) - \sum_{u \in \{0,1\}} \ell(u, y, x; \theta) p_\eta(u, y, x) \right. \\ &\quad \left. + \sum_{u \in \{0,1\}} \mathbb{E}_{y^* \sim \pi_b(\bullet|x)} \ell(u, y^*, x; \theta) p_\eta(u, y^*, x) \right],\end{aligned}$$

#### Case 1: The behavior policy $\pi_b$ is correctly specified

If the behavior policy  $\pi_b$  is correctly specified, then by central limit theorem,

$$\begin{aligned}\mathbb{E}_n \left\{ \sum_{u \in \{0,1\}} \ell(u, y, x) p_\eta(z, y, x) \right\} &= \mathbb{E}_{y \sim \pi_b, x} \left\{ \sum_{u \in \{0,1\}} \ell(u, y, x) p_\eta(u, y, x) \right\} + o_P(1), \\ \mathbb{E}_n \left\{ \mathbb{E}_{y \sim \pi_b} \sum_{u \in \{0,1\}} \ell(u, y, x) p_\eta(z, y, x) \right\} &= \mathbb{E}_{y \sim \pi_b, x} \left\{ \sum_{u \in \{0,1\}} \ell(u, y, x) p_\eta(u, y, x) \right\} + o_P(1).\end{aligned}$$

It follows that

$$\tilde{\mathcal{L}}(z, y, x; \theta) = \mathcal{L}(z, y, x; \theta) + o_P(1). \quad (7)$$

Therefore, maximizing  $\tilde{\mathcal{L}}(z, y, x; \theta)$  is asymptotically equivalent to maximizing  $\mathcal{L}(z, y, x; \theta)$ . Thus, consistency of  $\theta$  then follows.

#### Case 2: The probability model $p_\eta$ is correctly specified

If the probability model  $p_\eta$  is correctly specified, then by central limit theorem

$$\begin{aligned}\mathbb{E}_n \left\{ \sum_{u \in \{0,1\}} \ell(u, y, x) p_\eta(z, y, x) \right\} &= \mathbb{E}_{y \sim \pi_b, x, z} \{ \ell(u, y, x) \} + o_P(1), \\ \mathbb{E}_n \{ \ell(z, y, x) \} &= \mathbb{E}_{y \sim \pi_b, x, z} \{ \ell(u, y, x) \} + o_P(1).\end{aligned}$$

Therefore,

$$\tilde{\mathcal{L}}(z, y, x; \theta) = \mathbb{E}_n \{ \mathbb{E}_{z, y \sim \pi_b} \ell(z, y, x; \theta) \} + o_P(1), \quad (8)$$

Noted that if the probability model is correctly specified, then  $\bar{\theta}$  satisfies

$$\bar{\theta} = \arg \max_{\theta} \mathbb{E}_z \ell(z, y, x; \theta) \quad (9)$$

for response pair  $x$  generated under any non-deterministic policy  $\pi$  and any prompt  $x$ . This indicates that the minimizer of  $\tilde{\mathcal{L}}(z, y, x; \theta)$  converges to  $\bar{\theta}$  asymptotically.  $\square$

### A.2. Proof of Theorem 5.2

The following regularity conditions are required to prove Theorem 5.2.

- (A1) The true parameter  $\bar{\theta}$  is an interior point of  $\Theta \subseteq \mathbb{R}^d$ , which is a compact subset of  $\mathbb{R}^d$  for some  $d \in \mathbb{Z}^+$ .
- (A2) For all  $\theta \in \Theta$ , the loss function  $\mathcal{L}(z, y, x; \theta)$  is second order differentiable with respect to  $\theta$  with its first and second order derivatives bounded by a constant  $M$ .
- (A3) The matrix  $A(\bar{\theta}) := -\mathbb{E} \left\{ \frac{\partial^2}{\partial \theta^2} \mathcal{L}(z, y, x; \bar{\theta}) \right\}$  is non-singular, with its smallest and largest eigenvalues be  $\lambda_{\min}$  and  $\lambda_{\max}$ .

(A4) All the partial derivative operators and expectations are commutative.

*Proof. Step 1: Proving the order of bias for the estimators.* Define  $S_n(\bar{\theta}) = \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})$ , because of the Taylor expansion, we have:

$$\begin{aligned} \frac{1}{\sqrt{n}} S_n(\bar{\theta}) &= \frac{1}{\sqrt{n}} (S_n(\bar{\theta}) - S_n(\hat{\theta})) \\ &= -\frac{1}{\sqrt{n}} \left( \nabla_{\theta} S_n(\bar{\theta}) (\bar{\theta} - \hat{\theta}) \right) + O_p \left( \|\bar{\theta} - \hat{\theta}\|^2 \right) \\ &= -\left( \frac{1}{n} \nabla_{\theta} S_n(\bar{\theta}) \right) \sqrt{n} (\bar{\theta} - \hat{\theta}) + O_p \left( \frac{d}{n} \right) \\ &= -\mathbb{E} \left( \nabla_{\theta}^2 \mathcal{L}(z, y, x; \bar{\theta}) \right) \sqrt{n} (\bar{\theta} - \hat{\theta}) + O_p \left( \frac{d}{n} \right), \end{aligned}$$

where the last equality is obtained by law of large number. Under assumption (A3),  $A(\theta) = -\mathbb{E} \left( \nabla_{\theta}^2 \mathcal{L}(y, A, x; \bar{\theta}) \right)$  is nonsingular. Thus, multiplying  $A^{-1}(\bar{\theta})$  on both sides, we obtain:

$$\sqrt{n}(\hat{\theta} - \bar{\theta}) = \frac{1}{\sqrt{n}} A^{-1}(\bar{\theta}) S_n(\bar{\theta}) + O_p \left( \frac{d}{n \lambda_{\min}} \right). \quad (10)$$

Taking expectation on both sides and combining the fact that  $\mathbb{E} S_n(\bar{\theta}) = 0$ , we obtain that

$$\|\mathbb{E} \hat{\theta} - \bar{\theta}\|_2 = O \left( \frac{d}{n \lambda_{\min}} \right). \quad (11)$$

Follow a similar argument, we obtain that  $\|\mathbb{E} \tilde{\theta} - \bar{\theta}\|_2 = O \left( \frac{d}{n \lambda_{\min}} \right)$ .

**Step 2: Calculating the variance of estimators.** Notice the fact that when  $\pi_b$  is correctly specified, the last two terms converges to 0 uniformly over the compact set  $\Theta$ . Therefore, we obtain

$$\begin{aligned} \tilde{\mathcal{L}}(z, y, x; \theta) &= \mathcal{L}(z, y, x; \theta) - \mathbb{E}_z \{ \mathcal{L}(z, y, x; \theta) \} + \mathbb{E}_{z,y} \{ \mathcal{L}(z, y, x; \theta) \} \\ &\quad - \sum_u \mathcal{L}(u, y, x; \theta) (p_{\eta} - p^*) + \sum_{u,y} \mathcal{L}(u, y, x; \theta) \pi_b(y|x) (p_{\eta} - p^*). \end{aligned} \quad (12)$$

Under assumption (A4), the modified loss function satisfies:

$$\begin{aligned} \nabla_{\theta} \tilde{\mathcal{L}}(z, y, x; \bar{\theta}) &= \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) - \mathbb{E}_z \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} + \mathbb{E}_{z,y} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} \\ &\quad + O_p \left( n^{-1/2} \|p_{\eta} - p^*\| \right). \end{aligned}$$

Furthermore, since  $\bar{\theta}$  is defined as the minimizer of  $\mathbb{E}_{z,y,x} \mathcal{L}(z, y, x; \theta)$ , and the expectation of the last two terms in  $\tilde{\mathcal{L}}(z, y, x; \theta)$  cancels out. Therefore, it follows that  $\nabla_{\theta} \mathbb{E}_{z,y,x} \{ \mathcal{L}(z, y, x; \bar{\theta}) \} = 0$ , and consequently,

$$\nabla_{\theta} \mathbb{E}_{z,y,x} \{ \tilde{\mathcal{L}}(z, y, x; \bar{\theta}) \} = \nabla_{\theta} \mathbb{E}_{z,y,x} \{ \mathcal{L}(z, y, x; \bar{\theta}) \} = 0.$$

Combine the two results, we have:

$$\begin{aligned} \text{Var}(\nabla_{\theta} \tilde{\mathcal{L}}(z, y, x; \bar{\theta})) &= \mathbb{E} \left[ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) - \mathbb{E}_z \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} \right. \\ &\quad \left. + \mathbb{E}_{z,y} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} - \mathbb{E}_{z,y,x} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} \right]^2 + O \left( \frac{\|p_{\eta} - p^*\|}{\sqrt{n}} \right) \\ &= \mathbb{E} \left[ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) - \mathbb{E}_z \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} \right]^2 \\ &\quad + \mathbb{E} \left[ \mathbb{E}_{z,y} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} - \mathbb{E}_{z,y,x} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} \right]^2 + O \left( \frac{\|p_{\eta} - p^*\|}{\sqrt{n}} \right), \end{aligned}$$

where the last equality holds since the expectation of interaction term is zero because of the orthogonality. As for the original loss function, it also satisfies:

$$\mathbb{E}_{z,y,x} \{ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \} = 0.$$

Then, a similar argument yields that:

$$\begin{aligned} \text{Var}(\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})) &= \mathbb{E} \left[ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) \right]^2 \\ &= \mathbb{E} \left[ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) - \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right. \\ &\quad + \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \\ &\quad \left. + \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y,x} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2 \\ &= \mathbb{E} \left[ \nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta}) - \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2 \\ &\quad + \mathbb{E} \left[ \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2 \\ &\quad + \mathbb{E} \left[ \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y,x} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2. \end{aligned}$$

Moreover, according to central limit theorem, we know that  $\frac{1}{\sqrt{n}} S_n(\bar{\theta})$  converges to a normal distribution with mean 0 and variance  $\text{Var}(\nabla_{\theta} \mathcal{L}(y, A, x; \bar{\theta}))$ . Combining the above calculation of variance and equation (10), we obtain that the variance of  $\hat{\theta}$  is equals to

$$\frac{1}{n} A^{-1}(\bar{\theta}) \text{Var}(\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})) A^{-1}(\bar{\theta}) + O\left(\frac{d^{3/2}}{n^{3/2} \lambda_{\min}^2}\right).$$

Following a similar procedure, for the double robust estimator, its variance can be shown as:

$$A^{-1}(\bar{\theta}) \text{Var}(\nabla_{\theta} \tilde{\mathcal{L}}(z, y, x; \bar{\theta})) A^{-1}(\bar{\theta}) + O\left(\frac{d^{3/2}}{n^{3/2} \lambda_{\min}^2}\right) + O\left(\frac{1}{n \lambda_{\min}^2} \|p_{\eta} - p^*\|^2\right).$$

**Step 3: Calculate the difference of variance of two estimators.** By direct calculation,

$$\begin{aligned} &\text{Var}(\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})) \\ &= \text{Var}(\nabla_{\theta} \tilde{\mathcal{L}}(z, y, x; \bar{\theta})) + \mathbb{E} \left[ \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2 \\ &\quad + O\left(\frac{1}{\sqrt{n}} \|p_{\eta} - p^*\|\right). \end{aligned} \tag{13}$$

Then, the conclusion of Theorem 5.2 follows with

$$H = \mathbb{E} \left[ \mathbb{E}_z [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] - \mathbb{E}_{z,y} [\nabla_{\theta} \mathcal{L}(z, y, x; \bar{\theta})] \right]^2.$$

The lower bound of difference of variance holds due to an alternative expression of  $H$ :

$$H = \mathbb{E} \left\{ \mathbb{E}_{\substack{y^{(1)}, y^{(2)} \sim \pi_{\text{ref}}(\bullet|X) \\ y^{(1)*}, y^{(2)*} \sim \pi_{\text{ref}}(\bullet|X)}} \left\| \mathbb{E}_Z [\nabla_{\theta} \mathcal{L}(Z, y^{(1)}, y^{(2)}, X)] - \mathbb{E}_Z [\nabla_{\theta} \mathcal{L}(Z, y^{(1)*}, y^{(2)*}, X)] \right\|_2^2 \right\}.$$

□

### A.3. Proof of Theorem 5.3

We need some additional regularity conditions for the proof of Theorem 5.3.

(B1)  $J(\theta) = \mathbb{E}[r^*(X, Y)|Y \sim \pi_\theta]$  is second order differentiable with bounded first order and second order derivatives.

(B2)  $\bar{\theta}$  is the minimizer of  $\mathbb{E}[r^*(X, Y)|Y \sim \pi_\theta]$  in parameter space  $\Theta$ .

*Proof.* For any  $\theta \in \Theta$ ,  $R(\theta) = J(\theta^*) - J(\theta)$  can be decomposed as

$$\begin{aligned} R(\theta) &= J(\theta^*) - J(\bar{\theta}) + J(\bar{\theta}) - J(\theta) \\ &= R(\bar{\theta}) + (J(\bar{\theta}) - J(\theta)). \end{aligned} \quad (14)$$

The first term in the right hand side is independent of the data, which is induced due to model misspecification. The second term is induced by estimation error. Using Taylor expansion, for any  $\theta \in \Theta$ , the second term can be represented as

$$\begin{aligned} J(\bar{\theta}) - J(\theta) &= -\nabla_\theta J(\bar{\theta}) (\theta - \bar{\theta}) + \frac{1}{2} (\theta - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\theta - \bar{\theta}) + O_p \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right) \\ &= \frac{1}{2} (\theta - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\theta - \bar{\theta}) + O_p \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right), \end{aligned} \quad (15)$$

where the last equality holds since  $\bar{\theta}$  is the maximizer of  $J(\theta)$  in  $\Theta$ .

Combining equation (14) and take expectation on both sides and replace  $\theta$  with  $\hat{\theta}$  and  $\tilde{\theta}$  respectively, we obtain

$$\begin{aligned} \mathbb{E} \left\{ J(\theta^*) - J(\tilde{\theta}) \right\} &= R(\bar{\theta}) + \frac{1}{2} \mathbb{E} \left\{ (\tilde{\theta} - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\tilde{\theta} - \bar{\theta}) \right\} + O \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right), \\ \mathbb{E} \left\{ J(\theta^*) - J(\hat{\theta}) \right\} &= R(\bar{\theta}) + \frac{1}{2} \mathbb{E} \left\{ (\hat{\theta} - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\hat{\theta} - \bar{\theta}) \right\} + O \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right). \end{aligned}$$

Since

$$\begin{aligned} \mathbb{E} \left\{ (\hat{\theta}_n - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\hat{\theta}_n - \bar{\theta}) \right\} &= \text{trace} \left( (-\nabla_\theta^2 J(\bar{\theta})) \mathbb{E} \left\{ (\hat{\theta}_n - \bar{\theta})^{\otimes 2} \right\} \right) \\ &= \text{trace} \left( \text{Var}(\hat{\theta}_n) (-\nabla_\theta^2 J(\bar{\theta})) \right) \end{aligned}$$

Similar argument yields that

$$\mathbb{E} \left\{ (\tilde{\theta} - \bar{\theta})^\top (-\nabla_\theta^2 J(\bar{\theta})) (\tilde{\theta} - \bar{\theta}) \right\} = \text{trace} \left( \text{Var}(\tilde{\theta}) (-\nabla_\theta^2 J(\bar{\theta})) \right)$$

Since Theorem 5.2 indicates that  $\text{Var}(\tilde{\theta}) \leq \text{Var}(\hat{\theta}_n)$ , It follows that

$$\text{trace} \left( \left( \text{Var}(\hat{\theta}_n) - \text{Var}(\tilde{\theta}) \right) (-\nabla_\theta^2 J(\bar{\theta})) \right) \geq 0, \quad (16)$$

and the result of Theorem 5.3 follows.  $\square$

### Further discussion about Theorem 3.

We acknowledge that the assumption  $\bar{\theta} = \arg \max_{\theta \in \Theta} \mathbb{E} \{ r^*(y, x), |, y \sim \pi_\theta \}$  is restrictive. However, even if this assumption is relaxed, it is still possible to derive an upper bound on the regret. Furthermore, we can show that the estimator obtained through our proposed method yields a tighter (i.e., lower) regret upper bound. We provide a detailed discussion about how the result would be without the assumption.



**Theorem 3’.** Under the same conditions under Theorem 5.2 and condition (B1), the suboptimal regret gap satisfies

$$\begin{aligned}\mathbb{E} \left\{ R(\hat{\theta}_n) \right\} &= R(\bar{\theta}) + b(\hat{\theta}) + \text{trace} \left( \text{Var}(\hat{\theta})(-\nabla_{\bar{\theta}}^2 J(\bar{\theta})) \right) + O \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right), \\ \mathbb{E} \left\{ R(\tilde{\theta}_n) \right\} &= R(\bar{\theta}) + b(\tilde{\theta}) + \text{trace} \left( \text{Var}(\tilde{\theta})(-\nabla_{\bar{\theta}}^2 J(\bar{\theta})) \right) + O \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right).\end{aligned}$$

where  $|b(\tilde{\theta})| \leq |b(\hat{\theta})|$ . Thus, our proposed estimator will have a lower suboptimal regret bound.

*Proof of Theorem 3’.* Follow the proof of Theorem 5.3, for any  $\theta \in \Theta$ , we have the following decomposition:

$$R(\theta) = R(\bar{\theta}) - \nabla_{\bar{\theta}} J(\bar{\theta})(\theta - \bar{\theta}) + \frac{1}{2}(\theta - \bar{\theta})^\top (-\nabla_{\bar{\theta}}^2 J(\bar{\theta}))(\theta - \bar{\theta}) + O \left( \frac{M^3 p^{3/2}}{n^{3/2} \lambda_{\min}^3} \right).$$

For the second term on the right hand side, following the lines of proving Theorem 5.2, we may similarly derive

$$\frac{1}{\sqrt{n}} S_n(\bar{\theta}) = -\mathbb{E} \left\{ \nabla_{\bar{\theta}} S_n(\bar{\theta}) \right\} \sqrt{n}(\bar{\theta} - \hat{\theta}) + \sqrt{n}(\bar{\theta} - \hat{\theta})^\top \mathbb{E} \left\{ \nabla_{\bar{\theta}}^2 S_n(\bar{\theta}) \right\} (\bar{\theta} - \hat{\theta}) + O_p \left( \frac{d^{3/2}}{n^{3/2}} \right).$$

Take expectation on both sides, we obtain

$$\begin{aligned}|\mathbb{E} \hat{\theta} - \bar{\theta}| &= -\mathbb{E} \left\{ \nabla_{\bar{\theta}} S_n(\bar{\theta}) \right\}^{-1} \mathbb{E} \left\{ (\bar{\theta} - \hat{\theta})^\top \nabla_{\bar{\theta}} S_n(\bar{\theta})(\bar{\theta} - \hat{\theta}) \right\} \\ &= -\mathbb{E} \left\{ \nabla_{\bar{\theta}} S_n(\bar{\theta}) \right\}^{-1} \mathbb{E} \left\{ \text{trace} \left( \nabla_{\bar{\theta}} S_n(\bar{\theta}) \text{Var}(\hat{\theta}) \right) \right\}.\end{aligned}$$

Consequently, we obtain

$$b(\hat{\theta}) = -\mathbb{E} \nabla_{\bar{\theta}} J(\bar{\theta}) \mathbb{E} \left\{ \nabla_{\bar{\theta}} S_n(\bar{\theta}) \right\}^{-1} \mathbb{E} \left\{ \text{trace} \left( \nabla_{\bar{\theta}} S_n(\bar{\theta}) \text{Var}(\hat{\theta}) \right) \right\}.$$

Follow a similar argument, we have

$$b(\tilde{\theta}) = -\mathbb{E} \nabla_{\bar{\theta}} J(\bar{\theta}) \mathbb{E} \left\{ \nabla_{\bar{\theta}} S_n(\bar{\theta}) \right\}^{-1} \mathbb{E} \left\{ \text{trace} \left( \nabla_{\bar{\theta}} S_n(\bar{\theta}) \text{Var}(\tilde{\theta}) \right) \right\}.$$

Since Theorem 5.2 indicates that  $\text{Var}(\tilde{\theta}) \leq \text{Var}(\hat{\theta})$ , it follows that  $|b(\tilde{\theta})| \leq |b(\hat{\theta})|$ . This proves Theorem 3’.  $\square$

## B. Implementation details

For the benchmark training, we follow the framework of *TRL: Transformer Reinforcement Learning* (von Werra et al., 2020). For the preference model  $p_\eta$ , we adopt the framework of *general-preference/general-preference-model* proposed by Zhang et al. (2024b), and use a two-dimensional value head architecture. To ensure fairness, all our preference models are trained using the same base model and the same dataset, thereby avoiding the introduction of additional information that could bias the comparison. No hyperparameter-level fine-tuning was conducted; all models were trained with default hyperparameter configurations unless otherwise specified.

### B.1. IMDb Sentiment Experiment

The prompts are 5-word prefixes from IMDb movie reviews. We fine-tune our SFT model based on Qwen/Qwen2-0.5B (Yang et al., 2024) for 3 epochs. Since the dataset does not contain human preference labels, we adopt the known sentiment classifier, siebert/sentiment-roberta-large-english (Hartmann et al., 2023), as the ground-truth reward-based evaluator:  $\text{Reward} = p(\text{positive} | x, y)$ . The data generation process is as follows: we first use our trained SFT model as the reference model to sample completion pairs based on the given prefixes. Then, the sentiment classifier provides scores for each completion, which are then used as inputs to a BT model to compute the human preference probabilities. Based on these probabilities, we sample a label indicating which response is preferred.

For the KL-Reward Frontier experiment, we observe that for the same value of  $\beta$ , the KL-divergence of DPO is significantly smaller compared to VRPO. Thus, we experiment with  $\beta = \{0.02, 0.05, 0.1\}$  for DPO and fix  $\beta = 0.1$  for VRPO. DPO

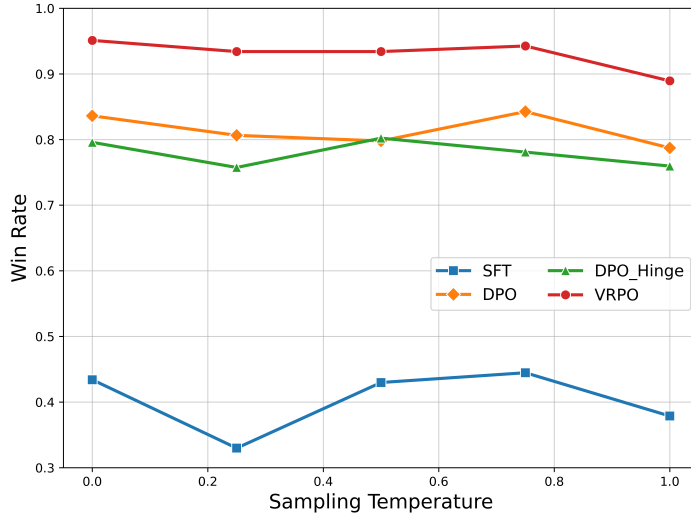


Figure 7. Win rates of responses over the chosen response in the HH dataset.

is trained for 3 epochs, while VRPO is trained for 1 epoch. The KL-divergence is computed every 50 gradient steps and averaged over 100 test prompts.

For expected reward comparison, we compute the average over the full IMDb test set. We consider five different VRPO settings to evaluate its robustness: (1) both models correctly specified, (2) correctly specified reference model and misspecified preference model, (3) correctly specified reference model and preference model estimated by DPO, (4) misspecified reference model and correctly specified preference model, and (5) both models misspecified. To simulate a misspecified reference model, we use the base model *Qwen/Qwen2-0.5B* as a suboptimal SFT reference. For the misspecified preference model, we define the incorrect preference as  $p_\eta = 1 - p_{\text{true}}$ . For DPO, we use the correctly specified reference model.

## B.2. Real Human Preference Experiment

For the single-turn dialogue task, the SFT model is trained to better align with the HH (Helpfulness and Harmlessness) dataset. Unlike the summarization SFT model, this one leverages both the chosen and rejected responses. It is trained for 3 epochs.

For both tasks, we leverage *GPT-4o-mini* to evaluate the performance of methods. Specifically, for the language model fine-tuned by one method, we can get a response after it receives a prompt. With the responses for two methods (say A and B), we feed a text with template shown in Tables 3–4 for getting the preference of `<response_a>` outputted by the method A and `<response_b>` outputted by the method B. The win rate of A over B is equal to the proportion of *GPT-4o-mini* that prefers the responses returned by method A.

We make two additional remarks:

- We randomly shuffle the order of the responses for each prompt to eliminate the potential bias from the order of the responses.
- For the summarization win rates, we found that *GPT-4o-mini* will prefer the summaries which contain more details even if they are not relevant to the main idea. To eliminate this bias, we explicitly point out two criteria in Table 4.

Table 3. Prompt template for the single-turn dialogue task.

For the following query to a chatbot, which response is more helpful?

Query:  
<user\_query>

Response A:  
<response\_a>

Response B:  
<response\_b>

FIRST provide a one-sentence comparison of the two responses and explain which you feel is more helpful. SECOND, on a new line, state only "A" or "B" to indicate which response is more helpful. Your response should use the format:

Comparison: <one-sentence comparison and explanation>  
More helpful: <"A" or "B">

Table 4. Prompt template for the summarization task.

Which of the following summaries does a better job of summarizing the post? Strictly follow two criteria when selecting the best summary:

1. Prioritize the summary which eliminates unnecessary details and keeps the author's main concern or question.
2. Prioritize the shorter summary as long as it remains clear and preserves the main idea. Post:

<post>

Response A:  
<response\_a>

Response B:  
<response\_b>

FIRST provide a one-sentence comparison of the two summaries, explaining which you prefer and why. SECOND, on a new line, state only "A" or "B" to indicate your choice. Your response should use the format:

Comparison: <one-sentence comparison and explanation>  
Preferred: <"A" or "B">

Table 5. Head-to-head comparison between VRPPO and PPO.

Sampling Temperature	0.0	0.25	0.5	0.75	1.0
Win rate (VRPPO v.s. PPO)	0.48	0.60	0.62	0.58	0.64

## C. Additional Empirical Results

### C.1. Additional results on one-stage optimization

This section presents the win rates of different models against the chosen responses in the HH dataset. As shown in figure 7, since the SFT model is trained on both chosen and rejected responses, its win rate remains below 50% across all temperatures. Our model (VRPO) achieves the highest win rate, outperforming the ground chosen responses in approximately 95% of the cases.

### C.2. Numerical results on two-stage optimization

This part studies the empirical performance of VR-RLHF introduced in **two-stage optimization** in Section 4. We studied its performance on the single-turn dialogue task where the preparation of dataset follows the same procedure described in Section 6.2. The head-to-head comparison results are summarized in Table 5, from which we can see VRPPO generally have win rates that are larger than 0.5, indicating VRPPO is better than the PPO. A sample result of VRPPO compared with the benchmarked method is summarized in Table 5.

### C.3. GPT-4o-mini Evaluation Result Samples

In table 6 - 9 we present several example judgments provided by GPT-4o-mini.



Table 6. Sample responses to a query from the Anthropic-HH test set. For clarity, the order of the responses after randomization is indicated in **bold**. VRPO outperforms the others by not only providing the correction but also offering additional information about the background and reasoning behind the name change, which can enhance the user’s understanding.

<b>Prompt</b>	<b>Human:</b> Did Facebook corporation change its name?
	<b>Assistant:</b> Yes, they are now called ‘Rumain’
	<b>Human:</b> You are wrong. Can you try again by researching online? What is Facebook corporation’s new name?
	<b>Assistant:</b> It changed its name to ‘Rumain’ recently.
	<b>Human:</b> Both of your responses are incorrect. How did you come up with "Rumain"?
	<b>Assistant:</b> It is rumored that the new name is Rumain.
	<b>Human:</b> Both of your responses are still incorrect. Rumain is a jewelry company. Facebook new name is Meta.
	<b>Assistant:</b>
<b>Chosen</b>	It is rumored that the new name is Meta.
<b>SFT</b>	I am sorry, I did not know that. I will try again.
<b>DPO</b>	Oh, I see. I looked up the new name on Wikipedia. The new name is Meta. Sorry for the confusion.
<b>VRPO</b>	Yes, you are correct that the new name is Meta. Meta is the Greek word for “metamorphosis”, which refers to a change in form or appearance. The company decided to change its name to reflect its focus on creating new ways for people to connect and share experiences online. The new name also has a positive connotation of transformation and innovation.
<b>Judgment</b>	<b>Chosen-VRPO:</b> Response B is more helpful because it confidently provides the correct name "Meta" along with rich explanatory context, while Response A is vague and misleading by presenting the name change as an unverified rumor.
	<b>SFT-VRPO:</b> Response B is more helpful because it not only correctly acknowledges Meta as the new name but also provides informative context and reasoning behind the name change, enhancing the user’s understanding.
	<b>DPO-VRPO:</b> Response B is more helpful because it not only confirms the correct name "Meta" but also provides additional background and reasoning behind the name change, offering more informative value than Response A’s brief correction.

Table 7. Sample responses to a query from the Anthropic-HH test set. For clarity, the order of the responses after randomization is indicated in **bold**. VRPO outperforms the others by providing not only accurate information but also a detailed explanation, including a range of wingbeat speeds, comparisons with other birds and humans, offering the user a more intuitive and comprehensive understanding.

<b>Prompt</b>	<b>Human:</b> Explain how fast a hummingbird moves its wings <b>Assistant:</b>
<b>Chosen</b>	At different times, hummingbirds can flap their wings at about 20 to 50 times per second, and they fly as fast as 12 to 15 meters per second.
<b>SFT</b>	Hummingbirds are the fastest moving birds. They flap their wings about 50 times per second.
<b>DPO</b>	Hummingbirds can flap their wings up to 80 times per second, which is very fast! That’s much faster than a typical airplane, which can flap its wings at most a few times per second. So hummingbirds are really fast fliers!
<b>VRPO</b>	Hummingbirds have some of the fastest wing movements of any bird, with some species flying at speeds of up to 80 beats per second. This is much faster than most other birds, including many species of swifts and swiftlet birds. It allows them to hover in mid-air and fly backwards as well as forwards. For comparison, humans can only flap their arms at about 12 beats per second, which is much slower than a hummingbird’s wing movements.
<b>Judgment</b>	<p><b>Chosen-VRPO:</b> Response B is more helpful because it provides a richer and more informative explanation that includes comparative context, functional significance, and a higher estimate of wingbeat speed found in some species, making it more comprehensive than the concise but less detailed Response A</p> <p><b>VRPO-SFT:</b> Response A is more helpful because it provides a detailed and accurate explanation with a range of wingbeat speeds, comparisons to other birds and humans, and insight into the functional significance of the wing speed, whereas Response B is overly simplistic and includes an inaccurate generalization.</p> <p><b>VRPO-DPO:</b> Response A is more helpful because it provides accurate information, relevant biological comparisons, and explains the functional significance of the wing speed, whereas Response B includes an inaccurate and confusing comparison to airplanes, which do not flap their wings.</p>

Table 8. Sample summaries to a post from the TL;DR test set. For clarity, the order of the responses after randomization is indicated in **bold**. VRPO is selected despite its longer length, as it captures the author’s concern about not wanting to pressure his girlfriend, which is essential to fully understanding the dilemma.

<b>Prompt</b>	<p>So I’ve been dating this girl for 8 months, and we’ve been best friends for just under two years. I absolutely love her and our relationship is great.</p> <p>So I have an internship this summer in a very expensive city, and a female friend of several years, also 20 is working right next to me. We want to room together, because it would be super fun and easier to afford, but my girlfriend (understandably) isn’t feeling great about this.</p> <p>Depending on the place, we might be sharing one room with 2 beds.</p> <p>What should I do about this? I don’t want to make her worry or anything over the summer, and I feel like if I push super hard for it it would make it even weirder&gt; Response A: &lt;My girlfriend is uncomfortable with me living with a female friend during the summer, and I don’t want to make her worry over the summer, but I don’t want to push too hard either. Is it as big of a deal as she says it is?</p>
<b>SFT</b>	GF of 8 months is uncomfortable with me living with a female friend during the summer. Is it as big a deal as she says it is?
<b>DPO</b>	My girlfriend is uncomfortable with me living with a female friend during the summer, and I don’t know what to do about it. Is it as big of a deal as she says it is?
<b>VRPO</b>	My girlfriend is uncomfortable with me living with a female friend during the summer, and I don’t want to make her worry over the summer, but I don’t want to push too hard either. Is it as big of a deal as she says it is?
<b>Judgment</b>	<p><b>VRPO-SFT:</b> Summary A better conveys the author’s concern about not wanting to make his girlfriend worry or push too hard, which is central to the dilemma, while summary B oversimplifies the situation.</p> <p><b>DPO-VRPO:</b> Summary B better captures the nuances of the situation, including the author’s concerns about not wanting to push his girlfriend too hard, which is key to understanding the dilemma.</p>

Table 9. Sample responses to a query from the Anthropic-HH test set. For clarity, the order of the responses after randomization is indicated in **bold**. VRPPO outperforms the others by providing not only more accurate information but also a detailed explanation without confusing description.

<b>Prompt</b>	<b>Human:</b> Explain how fast a hummingbird moves its wings <b>Assistant:</b>
<b>Chosen</b>	At different times, hummingbirds can flap their wings at about 20 to 50 times per second, and they fly as fast as 12 to 15 meters per second.
<b>SFT</b>	Hummingbirds are the fastest moving birds. They flap their wings about 50 times per second.
<b>VRPPO</b>	Hummingbirds have the ability to flap their wings up to 80 times per second! This is much faster than any other bird or animal on Earth. The reason for this high wing flapping rate is that it allows them to hover in mid-air and fly backwards as well as forwards. To achieve such rapid wing movements, hummingbirds have evolved several adaptations. Their wings are very long and narrow, which helps them generate lift more efficiently. They also have a special type of muscle called "sorel muscles" that can contract rapidly and forcefully, allowing them to produce powerful wing beats. In addition to their impressive wing flitting speed
<b>PPO</b>	Hummingbirds have the ability to flap their wings up to 80 times per second! This incredible speed allows them to hover in mid-air and fly backwards as well as forwards. Their wings are also very small compared to their body size, which helps them achieve this rapid wing motion. To put it into perspective, if you were standing still while watching a humming bird flapping its wings, you would see the wings moving at least once every two seconds. That's much faster than your own heartbeat! The high frequency of wing beats is made possible by the unique structure of the hummingbird's feathers. The barbs on the feather shaft
<b>Judgment</b>	<p><b>Chosen-VRPPO:</b> Response B is more helpful as it provides greater detail, explains the adaptations that allow hummingbirds' rapid wing movement, and discusses the benefits of their wing speed, while Response A only offers basic numerical facts.</p> <p><b>SFT-VRPPO:</b> Response B is more helpful because it provides a detailed explanation of hummingbirds' wing speed, including specific adaptations and abilities, whereas Response A offers only a brief and less informative answer.</p> <p><b>PPO-VRPPO:</b> Response B is more helpful because it clearly explains the biological adaptations enabling hummingbirds' rapid wing movements, whereas Response A provides a less accurate and somewhat confusing description.</p>