PAIRWISE PROXIMAL POLICY OPTIMIZATION: HARNESSING RELATIVE FEEDBACK FOR LLM ALIGNMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

Large Language Models (LLMs) can acquire extensive world knowledge through pre-training on large corpora. However, due to exposure to low-quality data, LLMs may exhibit harmful behavior without aligning with human values. The dominant approach for steering LLMs towards beneficial behavior involves Reinforcement Learning with Human Feedback (RLHF), with Proximal Policy Optimization (PPO) serving as the default RL optimizer. Despite its effectiveness, PPO has limitations when optimizing rewards trained from comparison-based loss. Primarily, PPO is not invariant to equivalent reward functions containing identical preference information due to the need to calibrate the reward scale. Additionally, PPO's necessity for token-wise updates introduces complexity in both function approximation and algorithm design compared to trajectory-wise optimization. This paper proposes a new framework, reinforcement learning with relative feedback, and a novel trajectory-wise policy gradient algorithm, Pairwise Proximal Policy Optimization (P3O) that operates directly on comparative rewards. We show theoretically that P3O is invariant to equivalent rewards and avoids the complexity of PPO. Empirical evaluations demonstrate that P3O outperforms PPO in the KL-Reward trade-off and can align with human preferences as well as or better than prior methods. In summary, this work introduces a simpler yet effective approach for aligning LLMs to human preferences through relative feedback.

1 INTRODUCTION

Large Language Models (LLMs) have made remarkable progress, profoundly influencing the AI community (Chowdhery et al., 2022; Brown et al., 2020; Touvron et al., 2023; Bubeck et al., 2023). However, due to the reliance on massive corpora of internet data, which encompasses a high portion of low-quality data, LLMs are likely to express unintended behavior. These include fabricating facts, generating biased or toxic text, and even harmful content to humans (Perez et al., 2022; Ganguli et al., 2022). Consequently, it is crucial to align LLMs with human values, *e.g.*, helpful, honest, harmless (Bai et al., 2022a).

A dominant approach in the realm of AI Alignment for LLMs named Reinforcement Learning with Human Feedback (RLHF) involves a three-stage procedure: supervised fine-tuning, reward learning, and reinforcement learning (RL) fine-tuning (Ziegler et al., 2019; Ouyang et al., 2022). In the critical third stage, Proximal Policy Optimization (PPO) is widely adopted as the default RL optimizer (Schulman et al., 2017). Despite the acclaimed efficiency of PPO, recent studies have highlighted several intriguing questions and potential issues that require additional attention:

Instability of PPO. Despite its acclaimed effectiveness, recent studies have identified instability associated with PPO. Factors such as reward normalization, reward scaling, reward clipping, KL control, advantage normalization and critic initialization (Zheng et al., 2023; Engstrom et al., 2020) can contribute to this instability. Moreover, we identify another source of instability: there is an inconsistency when it comes to optimizing the reward trained with the Bradley-Terry Loss (BTL) comparative reward model. In essence, BTL is invariant to constant shift while PPO is not. This means that even if two reward models carry identical information about human preferences, optimizing them using PPO can produce different outcomes. (Section 4.4)

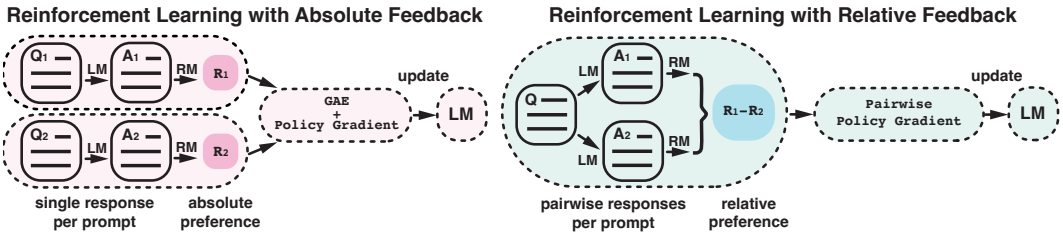

Figure 1: The figure on the left illustrates the prevalent method for fine-tuning LMs using RL, which relies on **Absolute Feedback**. In this paradigm, algorithms like PPO has to learn a V function, which capture not only the valuable relative preference information, but also less part, which is the scale of the reward for a given prompt. Contrastingly, the figure on the right presents paradigm for optimizing reward model trained via comparative loss, *e.g.*, Bradley-Terry Loss (Bradley & Terry, 1952). Our algorithm generate a pair of responses per prompt, leveraging only the **Relative Feedback** - derived from the difference in reward - for policy gradient updates. This method obviates the need for additional V function approximations and intricate components like GAE (Schulman et al., 2015b).

Trajectory-wise vs Token-wise. A prevailing tension in fine-tuning LLMs using RL revolves around whether tuning should be trajectory-wise or token-wise. This dichotomy arises primarily from the nature of the reward model training phase wherein the reward model is trained using comparison dataset with human labelers ranking multiple completions of a single prompt. The obvious reason for taking this indirect approach based on trajectory-wise comparison is that it is often challenging for labelers to make a direct score assignment or provide comparison beyond a trajectory-wise perspective. An exception might be specific tasks with less ambiguity, such as math problems (Lightman et al., 2023), where partial credit for partially correct answers is feasible. This difficulty arises from the fact that the reward model returns only a single scalar reward after receiving the entire response (when the model hits the `<eos>` token). Given this complexity, two dominant viewpoints emerge: Should language generation be approached as a Contextual Bandit (CB) or a Markov Decision Process (MDP)? Although the prevalent approach is to use PPO to perform token-wise updates, it introduces additional complexity both in function approximation and in algorithmic design, *i.e.*, PPO requires to additionally learn a value function V and use Generalized Advantage Estimation (GAE) (Schulman et al., 2015b) to redistribute the sparse reward to all tokens.

In this paper, we provide new insights to address the above two issues. We summarize our primary contributions as:

(1) We define an equivalent relationship for reward functions trained from human preferences. We identify that BTL is invariant under this equivalent relationship, while PPO is not.

(2) We observed that Direct Preference Optimization (DPO) (Rafailov et al., 2023) can be used jointly with a trained reward model, leading to online-DPO. Further, we demonstrate that online-DPO remains invariant under the reward equivalent relationship. Empirical evidence suggest that online-DPO can achieve higher reward while being less KL efficient than PPO.

(3) We introduce a novel policy gradient algorithm, Pairwise Proximal Policy Optimization (P3O) under the framework of Reinforcement Learning with Relative Feedback (Figure 1). Our algorithm aligns perfectly with the comparative nature of the reward model, avoiding complications like estimating the V function, GAE and various normalization techniques (Zheng et al., 2023). Empirical evaluations show that P3O consistently outperforms PPO and DPO in terms of KL-Reward trade-off and GPT-4 Evaluation.

2 Related Work

Significant efforts have been made towards aligning LLMs with human values. These alignment strategies broadly fall into two categories: offline training and online training.

Offline training typically involve a static dataset, and doesn't require additional evaluations or generations. For instance, Thoppilan et al. (2022); Gunasekar et al. (2023) use instruction fine-tuning to update the model on a high quality dataset tailored to a specific downstream task of interest. Snell et al. (2022) proposed to employ offline Q Learning to learn an add-on term for decoding. While

Rafailov et al. (2023) introduced DPO, an offline approach that can directly align LM with human preference data, drawing from the closed-form solution of the Contextual Bandit with KL control problem. There are also methods like PRO (Song et al., 2023) and RRHF (Yuan et al., 2023) that fine-tune the model based on ranking of the rewards.

Our work is categorized under online training, which consist of a loop of generating new responses from the updated policy, evaluating them with the reward model and update the policy. The current dominant approach RLHF relies on online RL methods such as PPO (Schulman et al., 2017), A2C (Mnih et al., 2016) or their variants (Ramamurthy et al., 2022; Zhu et al., 2023b). There are also few methods that deviate from this standard. For instance, Gulcehre et al. (2023) introduce ReST, which use offline RL instead of online RL in the policy improvement phase; Dong et al. (2023) proposed RAFT, which iteratively fine-tune the policy on the responses generated by the Best-of-N policy. Another paradigm parallel to RLHF is Reinforcement Learning with AI Feedback (RLAIF) (Bai et al., 2022b; Lee et al., 2023), which aim for AI self-improvement. RLAIF substitute the role of human with AI in the feedback and yield comparable results with models of smaller scale.

Outside of the context of language, preference-driven policy learning has been explored in both bandit and RL. Contextual dueling bandit (Dudík et al., 2015; Yue et al., 2012) use preferences or rankings of actions to adjust the policy, rather than rewards. Similarly, PbRL (Jain et al., 2013; Busa-Fekete et al., 2014; Christiano et al., 2017; Sadigh et al., 2017; Kupcsik et al., 2018) learn from binary preferences generated by some unknown scoring function. Our work share resemblance to the pairwise policy gradient outlined in Xu et al. (2020), though we integrate the critical clipping technique and the KL control that better fit into the language generation setting.

3 PRELIMINARIES

In this section, we review the RLHF pipeline and discuss how language generation fits into CB and RL setting(Sutton & Barto, 2018; Yang et al., 2021; Lattimore & Szepesvári, 2020; Wu et al., 2022).

- **SFT Phase (Supervised Fine-Tuning):** This stage start with a pre-trained LM, and then fine-tuned with supervised learning (typically maximum likelihood loss) on a high quality dataset for the downstream task of interest. Outcome of this stage is denoted as π^{SFT}.
- **Reward Learning Phase.** In the second phase the SFT model is prompted with prompts $\boldsymbol{x}$ to produce pairs of answers $\boldsymbol{y}_1, \boldsymbol{y}_2 \sim \pi^{\text{SFT}}(\boldsymbol{y}|\boldsymbol{x})$. The responses pairs are then presented to human labelers who express preferences for one answer, denoted as $\boldsymbol{y}_w \succ \boldsymbol{y}_l|\boldsymbol{x}$, where $\boldsymbol{y}_w$ is the one favored by the labeler and $\boldsymbol{y}_l$ is the one less favored. Under these preferences is the inaccessible latent reward model $r^*(\boldsymbol{y}, \boldsymbol{x})$. According to the Bradley-Terry (Bradley & Terry, 1952) model, the human preference distribution p^* can be expressed as:

  $$p^*(\boldsymbol{y}_1 \succ \boldsymbol{y}_2|\boldsymbol{x}) = \frac{\exp(r^*(\boldsymbol{y}_1|\boldsymbol{x}))}{\exp(r^*(\boldsymbol{y}_1|\boldsymbol{x})) + \exp(r^*(\boldsymbol{y}_2|\boldsymbol{x}))} = \frac{1}{1 + \exp\left(r^*(\boldsymbol{y}_1|\boldsymbol{x}) - r^*(\boldsymbol{y}_2|\boldsymbol{x})\right)}$$

  Assuming the access to a dataset $\{(\boldsymbol{x}^{(i)}, \boldsymbol{y}_w^{(i)}, \boldsymbol{y}_l^{(i)})\}_{i=1}^N$ sampled from p^*. We parameterize the reward as r_ϕ and estimate it via maximum log-likelihood:

  $$\mathcal{L}_R = \sum_{i=1}^{N} \frac{1}{N} \log \sigma \left(r_\phi(\boldsymbol{y}_w^{(i)}|\boldsymbol{x}^{(i)}) - r_\phi(\boldsymbol{y}_l^{(i)}|\boldsymbol{x}^{(i)}) \right) \tag{1}$$

  where σ is the sigmoid function. r_ϕ is initialized with π^{SFT} augmented by additional linear layers on top. Constraints like $\mathbb{E}\left[r(\boldsymbol{y}|\boldsymbol{x})\right] = 0$ might be incorporated to lower the variance.
- **RL Fine-Tuning Phase.** Prior work formulate the optimization problem as:

  $$\max_{\pi_\theta} \mathbb{E}_{\boldsymbol{x}\sim\mathcal{D}, \boldsymbol{y}\sim\pi_\theta(\cdot|\boldsymbol{x})} \left[r_\phi(\boldsymbol{y}|\boldsymbol{x}) - \beta D_{\text{KL}}(\pi_\theta(\cdot|\boldsymbol{x}) \| \pi^{\text{SFT}}(\cdot|\boldsymbol{x})) \right] \tag{2}$$

  The $\beta D_{\text{KL}}(\pi_\theta(\cdot|\boldsymbol{x}) \| \pi^{\text{SFT}}(\cdot|\boldsymbol{x}))$ term is used to regulate the deviation from the SFT model, it is important to prevent the model from completely forget the world knowledge acquired in the pre-training stage. The standard approach is to directly employ PPO (Schulman et al., 2017; Ouyang et al., 2022) to optimize the modified reward $r_\phi(\boldsymbol{y}|\boldsymbol{x}) - \beta \left(\log \pi_\theta(\boldsymbol{y}|\boldsymbol{x}) - \log \pi^{\text{SFT}}(\boldsymbol{y}|\boldsymbol{x}) \right)$.

3.1 MDP VS CB

Language Generation modeled as an MDP: A language model takes a sequence of tokens $(t_1, ..., t_h)$ as input and produce the distribution for the next token t_{h+1}. The action space of this MDP consists of all potential tokens, with the state being a concatenation of all the history tokens as well as the prompt. The transition in the model is deterministic, the next state equal to the concatenation of current state and the recently produced token. Formally,

$$P\left(s_{h+1} = (t_1, ..., t_{h+1}) | s_h = (t_1, ..., t_h), a_{h+1} = t_{h+1}\right) = 1$$

An episode end either when the total tokens generated attain a predefined limit or when the language model produce the special token <eos>. Upon the end of an episode, the entire sequence is scored by a reward model, which produces a scalar reward r. PPO follows this formulation, where the single scalar reward will be assigned to the last token and combine with the token-wise KL penalty to form the ultimate reward. To tackle the challenge of sparse feedback, PPO employs the Generalized Advantage Estimation (Schulman et al., 2015b) to reallocate the scalar reward across all tokens.

Language Generation modeled as a CB: A language model takes a prompt $\boldsymbol{x} = (t_1, ..., t_h)$ and generates a response $\boldsymbol{y} = (t_{h+1}, ..., t_{h'})$. The response either terminates with with a <eos> token or continues until a predefined token limit is met. Within this setting, every possible response is treated as an action, and there's no transition at all. The entire sequence is scored by a reward model, producing a scalar reward. This formulation is adopted by (Rafailov et al., 2023; Zhu et al., 2023a) as well as our paper.

4 ALGORITHM

4.1 PROXIMAL PAIRWISE POLICY OPTIMIZATION (P3O)

To derive P3O, we start from Vanilla Policy Gradient (VPG, Pseudocode 1) (Sutton et al., 1999; Schulman et al., 2017). For simplicity, we omit the prompt $\boldsymbol{x}$ in the formula and focus on the derivation for bandit setting and extend to the more general contextual bandit in theorem 1.

In the bandit setting, assume we are updating a parameterized policy π_θ with actions denoted as $\boldsymbol{y}$. The VPG aims for estimating the following formula with samples:

$$\nabla \mathcal{L}^{\text{VPG}} = \mathop{\mathbb{E}}_{\boldsymbol{y} \sim \pi_\theta} r(\boldsymbol{y}) \nabla \log \pi_\theta(\boldsymbol{y}) = \sum_{\boldsymbol{y}} r(\boldsymbol{y}) \nabla \pi_\theta(\boldsymbol{y}) \tag{3}$$

Previous works on PG propose to subtract a baseline b to reduce variance, resulting in

$$\nabla \mathcal{L}^{\text{VPG}} = \sum_{\boldsymbol{y}} (r(\boldsymbol{y}) - b) \nabla \pi_\theta(\boldsymbol{y})$$

We plug in the popular choice of the baseline, $b = \sum_{\boldsymbol{y}} r(\boldsymbol{y}) \pi_\theta(\boldsymbol{y})$:

$$\begin{aligned}
\nabla \mathcal{L}^{\text{VPG}} &= \sum_{\boldsymbol{y}_1} r(\boldsymbol{y}_1) \nabla \pi_\theta(\boldsymbol{y}_1) - \sum_{\boldsymbol{y}_1, \boldsymbol{y}_2} r(\boldsymbol{y}_2) \pi_\theta(\boldsymbol{y}_2) \nabla \pi_\theta(\boldsymbol{y}_1) \\
&= \sum_{\boldsymbol{y}_1, \boldsymbol{y}_2} r(\boldsymbol{y}_1) \pi_\theta(\boldsymbol{y}_2) \nabla \pi_\theta(\boldsymbol{y}_1) - \sum_{\boldsymbol{y}_1, \boldsymbol{y}_2} r(\boldsymbol{y}_2) \pi_\theta(\boldsymbol{y}_2) \nabla \pi_\theta(\boldsymbol{y}_1) \\
&= \sum_{\boldsymbol{y}_1, \boldsymbol{y}_2} \left(r(\boldsymbol{y}_1) - r(\boldsymbol{y}_2)\right) \pi_\theta(\boldsymbol{y}_2) \nabla \pi_\theta(\boldsymbol{y}_1)
\end{aligned} \tag{4}$$

Equation (4) gives rise to an expression that directly rely on the relative difference $r(\boldsymbol{y}_1) - r(\boldsymbol{y}_2)$. We further introduce $\pi_{\theta_{\text{old}}}$ and rewrite Eq (4) with importance sampling:

$$\nabla \mathcal{L}^{\text{VPG}} = \mathop{\mathbb{E}}_{\boldsymbol{y}_1, \boldsymbol{y}_2 \sim \pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1) - r(\boldsymbol{y}_2)\right) \frac{\pi_\theta(\boldsymbol{y}_2)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2)} \frac{\nabla \pi_\theta(\boldsymbol{y}_1)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1)}$$

Swap actions $\boldsymbol{y}_1, \boldsymbol{y}_2$ and average together we get a symmetry gradient expression we termed Pairwise Policy Gradient (PPG):

$$\nabla \mathcal{L}^{\text{PPG}} = \mathop{\mathbb{E}}_{\boldsymbol{y}_1, \boldsymbol{y}_2 \sim \pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1) - r(\boldsymbol{y}_2)\right) \left(\frac{\pi_\theta(\boldsymbol{y}_2)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2)} \frac{\nabla \pi_\theta(\boldsymbol{y}_1)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1)} - \frac{\pi_\theta(\boldsymbol{y}_1)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1)} \frac{\nabla \pi_\theta(\boldsymbol{y}_2)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2)} \right) / 2 \tag{5}$$

$$= \mathop{\mathbb{E}}_{\boldsymbol{y}_1, \boldsymbol{y}_2 \sim \pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1) - r(\boldsymbol{y}_2)\right) \frac{\pi_\theta(\boldsymbol{y}_1)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1)} \frac{\pi_\theta(\boldsymbol{y}_2)}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2)} \nabla \left(\log \frac{\pi_\theta(\boldsymbol{y}_1)}{\pi_\theta(\boldsymbol{y}_2)} \right) / 2 \tag{6}$$

Its immediate generalization to contextual bandit is the following theorem:

Theorem 1 (Pairwise Policy Gradient with Importance Sampling)**.** *For any prompt $\boldsymbol{x}$, the policy gradient can be expressed as $\nabla\mathcal{L}^{VPG} = \mathop{\mathbb{E}}_{\boldsymbol{x}\sim\mathcal{D}} \nabla\mathcal{L}^{VPG}(\boldsymbol{x})$, where $\nabla\mathcal{L}^{VPG}(\boldsymbol{x})$ can be expressed as:*

$$\begin{aligned}
&\mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{old}}} (r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})) \left(\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_2|\boldsymbol{x})} \frac{\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_1|\boldsymbol{x})} - \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_1|\boldsymbol{x})} \frac{\nabla\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_2|\boldsymbol{x})} \right) /2 \\
=&\mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{old}}} (r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_1|\boldsymbol{x})} \frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{old}}(\boldsymbol{y}_2|\boldsymbol{x})} \nabla \left(\log \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})} \right) /2
\end{aligned}$$

4.2 Combine with Clipping

A significant modification of PPO over VPG is to clip the loss function, which disincentivizes large updates to the policy. Specifically, PPO dictates that the ratio $\pi_\theta/\pi_{\theta_{\text{old}}}$ should remain close to 1, guided by the sign of the Advantage function Adv. Understanding this in an intuitive sense, if $\mathsf{Adv}(\boldsymbol{y}|\boldsymbol{x}) > 0$, it implies that taking the action $\boldsymbol{y}$ is beneficial on average. Hence, we aim to increase the probability $\pi_\theta(\boldsymbol{y}|\boldsymbol{x})$. However, if the policy ratio $\pi_\theta/\pi_{\theta_{\text{old}}}$ exceeds $1+\epsilon$, we consider the change sufficient and halt the gradient; otherwise, the gradient is computed for further learning. Conversely, if $\mathsf{Adv}(\boldsymbol{y}|\boldsymbol{x}) < 0$, we strive to optimize the ratio towards $1-\epsilon$ instead of $1+\epsilon$. This intuition guides us to derive two variants of the algorithm, differentiated by whether clipping is applied separately or jointly for actions $\boldsymbol{y}_1$ and $\boldsymbol{y}_2$.

Clipping Separately (Version 1): For $\{i,j\} = \{1,2\}$,

$$\begin{aligned}
\mathcal{L}_i^{\text{P3O}}(\boldsymbol{x}) &= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \texttt{sg}\left((r(\boldsymbol{y}_i|\boldsymbol{x}) - r(\boldsymbol{y}_j|\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}_j|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_j|\boldsymbol{x})} \right) \frac{\pi_\theta(\boldsymbol{y}_i|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_i|\boldsymbol{x})} \\
\mathcal{L}_{i,\texttt{clip}}^{\text{P3O}}(\boldsymbol{x}) &= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \texttt{sg}\left((r(\boldsymbol{y}_i|\boldsymbol{x}) - r(\boldsymbol{y}_j|\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}_j|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_j|\boldsymbol{x})} \right) \texttt{clip}(\frac{\pi_\theta(\boldsymbol{y}_i|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_i|\boldsymbol{x})}, 1-\epsilon, 1+\epsilon) \\
\mathcal{L}_{\text{sep}}^{\text{P3O}} &= \mathop{\mathbb{E}}_{\boldsymbol{x}\sim\mathcal{D}} \left[\min(\mathcal{L}_1^{\text{P3O}}(\boldsymbol{x}), \mathcal{L}_{1,\texttt{clip}}^{\text{P3O}}(\boldsymbol{x})) + \min(\mathcal{L}_2^{\text{P3O}}(\boldsymbol{x}), \mathcal{L}_{2,\texttt{clip}}^{\text{P3O}}(\boldsymbol{x})) \right] /2
\end{aligned}$$

Clipping Jointly (Version 2):

$$\begin{aligned}
\mathcal{L}^{\text{P3O}}(\boldsymbol{x}) &= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \texttt{sg}\left((r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})} \frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} \right) \log \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})} \\
\mathcal{L}_{\texttt{clip}}^{\text{P3O}}(\boldsymbol{x}) &= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \texttt{sg}\left((r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})} \frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} \right) \\
&\qquad \times \texttt{clip}\left(\log \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}, \log \frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} - \epsilon, \log \frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} + \epsilon \right) \\
\mathcal{L}_{\text{joi}}^{\text{P3O}} &= \mathop{\mathbb{E}}_{\boldsymbol{x}\sim\mathcal{D}} \min(\mathcal{L}^{\text{P3O}}(\boldsymbol{x}), \mathcal{L}_{\texttt{clip}}^{\text{P3O}}(\boldsymbol{x}))
\end{aligned}$$

The `sg` and `clip` refers to the stop-gradient and clip operator respectively. We present the entire pipeline of P3O in Pseudocode 2.

4.3 Relationship with PPO and DPO

In this section, we briefly discuss the relationship of our proposed algorithm and two existing algorithms, PPO and DPO. We further elaborate their theoretical properties in Section 4.4.

Comparison with PPO: The objective of PPO, before clipping is applied and reduced to the contextual bandit setting, can be expressed as follows:

$$\mathcal{L}_{\text{no clip}}^{\text{PPO}} = -\mathop{\mathbb{E}}_{\boldsymbol{y}\sim\pi_{\theta_{\text{old}}}(\cdot|\boldsymbol{x})} (r(\boldsymbol{y}|\boldsymbol{x}) - V_\phi(\boldsymbol{x})) \frac{\pi_\theta(\boldsymbol{y}|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}|\boldsymbol{x})}$$

Where $V_\phi(\boldsymbol{x})$ is a proxy to the value function $V^{\pi_{\theta_{\text{old}}}} = \mathbb{E}_{\boldsymbol{y}\sim\pi_{\theta_{\text{old}}}} r(\boldsymbol{y}|\boldsymbol{x})$. In contrast to PPO, P3O expands $V_\phi(\boldsymbol{x})$ using another sample $\boldsymbol{y}'$ to construct an unbiased estimator, bypassing the need to learn an additional value function V, which leads to more stable result.

Comparison with DPO: Like P3O, DPO also adopts the contextual bandit (CB) formulation. The gradient of DPO's objective function takes the following form:

$$\nabla \mathcal{L}^{\text{DPO}}(\boldsymbol{x}, \boldsymbol{y}_w, \boldsymbol{y}_l) = -\beta\sigma\left(\beta \log \frac{\pi_\theta(\boldsymbol{y}_l|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_l|\boldsymbol{x})} - \beta \log \frac{\pi_\theta(\boldsymbol{y}_w|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_w|\boldsymbol{x})}\right) \nabla \left(\log \frac{\pi_\theta(\boldsymbol{y}_w|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_l|\boldsymbol{x})}\right) /2$$

The gradient direction of DPO resembles that of our formulation in Theorem 1. However, DPO uses a different weight in front of the gradient. We hypothesize that the reason for DPO usually achieving a higher reward but falling short on KL-control (Figure 2 and 3) compared to PPO is because DPO aligns the policy towards the goal policy while doesn't directly consider the reward. Unlike PPO and P3O, which applies policy gradient based on the idea of strict policy improvement for every gradient update (Schulman et al., 2015a), DPO aligns the policy via an alternate "distance", where the intermediary steps are not guaranteed to maximize the KL-Reward trade-off. We note that P3O combines the benefits of PPO and DPO, offering guaranteed policy improvement akin to policy gradient.

4.4 Reward Equivalence & Instability of PPO

In this section, we formally define the concept of reward equivalence (Definition 1). We show that BTL is invariant under this equivalent relationship in lemma 1. We then discuss why it leads to a desirable property named invariance (Definition 2) that we want RL algorithms to satisfy. In the end, we present our main theorem (Theorem 2) which shows that PPO does not satisfy this property, contributing to its instability.

Definition 1 (Reward Equivalence)**.** *Two reward functions $r(\boldsymbol{y}|\boldsymbol{x})$ and $r'(\boldsymbol{y}|\boldsymbol{x})$ are termed equivalent, denoted as $r \sim r'$, if and only if there exist a function $\delta(\boldsymbol{x})$ depend solely on the prompt $\boldsymbol{x}$, such that for every prompt and response pair $(\boldsymbol{x}, \boldsymbol{y})$,*

$$r(\boldsymbol{y}|\boldsymbol{x}) - r'(\boldsymbol{y}|\boldsymbol{x}) = \delta(\boldsymbol{x})$$

The equivalent class associated with reward r is represented as $[r]$.

Note that comparative losses such as Bradley-Terry loss and Plackett-Luce loss, is unaffected by a shift in the prompt's reward as in definition 1. This observation leads to the following Lemma:

Lemma 1 (Invariance of BTL)**.** *For two reward functions that satisfy $r \sim r'$, they both yield identical loss for any response pairs (or K responses) under the Bradley-Terry Loss (or Plackett-Luce Loss).*

Lemma 1 underscores that the only information we can learn from the preference data is the reward difference of two responses to the same prompt. This implies that direct comparison of responses stemming from different prompts should be avoided. This is because we can craft an arbitrary function denoted as δ and replace $\hat{r}$ with the identical $\hat{r} + \delta$, while flipping the sign of $\hat{r}(\boldsymbol{y}|\boldsymbol{x}) - \hat{r}'(\boldsymbol{y}'|\boldsymbol{x}')$. As a result, an ideal algorithm should focus only on the relevant information within the reward function, filtering out the noise represented by δ. This leads the following definition:

Definition 2 (Invariance)**.** *An algorithm is said to be **invariant** with respect to the equivalent relation "$\sim$", if for any two equivalent reward functions $r \sim r'$ and a fixed set of prompt and response pairs, the algorithm perform identical updates to the policy.*

To illustrate definition 2, assume that we have two equivalent reward functions $\hat{r}$ and $\hat{r}' = \hat{r} + \delta$. Notably, even when initialized with the same random seed, PPO can result in distinct updates for an identical batch. This behavior can be attributed to PPO's reliance on learning a V function to estimate advantage. In the simplest scenario, where the advantage is estimated via one-step TD ($\mathsf{Adv}(\boldsymbol{y}|\boldsymbol{x}) = r(\boldsymbol{y}|\boldsymbol{x}) - V(\boldsymbol{x})$, corresponding to $\lambda_{\text{GAE}} = 0$) and $\boldsymbol{y}$ is a single token, we should expect the advantage function to stay unchanged. However, following the derivation

$$\begin{aligned} & \mathsf{Adv}_{\hat{r}}(\boldsymbol{y}|\boldsymbol{x}) = \mathsf{Adv}_{\hat{r}'}(\boldsymbol{y}|\boldsymbol{x}) \\ \iff & \hat{r}(\boldsymbol{y}|\boldsymbol{x}) - V_{\hat{r}}(\boldsymbol{x}) = \hat{r}'(\boldsymbol{y}|\boldsymbol{x}) - V_{\hat{r}'}(\boldsymbol{x}) \\ \iff & V_{\hat{r}'}(\boldsymbol{x}) - V_{\hat{r}}(\boldsymbol{x}) = \delta(\boldsymbol{x}) \end{aligned}$$

We can see that even though $\hat{r}$ and $\hat{r}'$ are equivalent, they yield different updates for V function. This give rise to our main theorem (full proof in Appendix B.2):

Theorem 2 (Non-invariance of PPO)**.** *P3O and DPO are invariant with respect to "$\sim$". In contrast, PPO is not, given the same initialization of V.*

5 EXPERIMENTS

In this section, we empirically study how well can P3O align with human preference. We conduct experiments on two widely-adopted RLHF tasks, summarization and question-answering, and we find that P3O achieves better performance in terms of both KL-Reward trade-off and quality of generation, against several strong baselines. We will first briefly introduce the tasks, compared methods, and evaluations in our experiments, and then elaborate on these findings in detail.

Tasks. We explore two different open-ended text generation tasks, *i.e.* **summarization** and **question-answering**. For both tasks, algorithms are given a reward model pre-trained from a dataset of preference $\mathcal{D} = \{\boldsymbol{x}^{(i)}, \boldsymbol{y}_w^{(i)}, \boldsymbol{y}_l^{(i)}\}$, and the goal is to obtain a policy $\pi(\boldsymbol{y}|\boldsymbol{x})$ that can generate high-quality response $\boldsymbol{y}$ given prompt $\boldsymbol{x}$. In summarization, we use the **TL;DR** "too long; didn't read" dataset (Völske et al., 2017), where $\boldsymbol{x}$ is a forum post from Reddit, and $\boldsymbol{y}$ is a corresponding summary. We use a 6B SFT model `CarperAI/openai_summarize_tldr_sft` as the initial policy and `EleutherAI/gpt-j-6b` as the reward model. In question-answering, $\boldsymbol{x}$ is a human query, which may come from diverse topics, and the policy should learn to produce an engaging and helpful response $\boldsymbol{y}$. Following prior work, we use the Anthropic Helpful and Harmless (**HH**) dataset (Bai et al., 2022a). We fine-tune two policies of sizes {1B,6B}, `Dahoas/pythia-1B-static-sft` and `Dahoas/pythia-6B-static-sft`. Both models have gone through supervised fine-tuning with labeled prompt-response pairs, similar to the protocol in Ouyang et al. (2022) and Ramamurthy et al. (2022). For the reward model, we use the 6B model `Dahoas/gptj-rm-static` trained from the same dataset based on `EleutherAI/gpt-j-6b` as a proxy of human preference.

Methods. We compare two versions of P3O, **P3O-V1** and **P3O-V2**, which represent clipping separately and jointly respectively, with several effective and representative approaches for LLM alignment. We start with the **SFT** policy trained by token-wise supervised fine-tuning. It hasn't gone through further alignment; Every other method uses the SFT model as initialization. For RL algorithms[1], we consider the dominant approach **PPO** (Schulman et al., 2017; Ouyang et al., 2022) with reward specified in Eq (2). We follow the implementation of `trlx` (Castricato et al., 2023). Besides, we also consider the newly proposed **DPO** (Rafailov et al., 2023), a method that directly optimizes the policy towards the closed-form solution of the KL-constrained reward maximization. Although DPO is proposed as an offline alignment method, we notice that we can make it online with the help of a proxy reward function. (More details can be found in Appendix A.2)

Evaluations. Deviating too much from the reference policy (*e.g.* SFT model) would lead the online policy to cut corners of the reward model and produce incoherent continuations, as pointed out by previous works (Ziegler et al., 2019). Gao et al. (2023) studied the scaling law of reward over-optimization in a synthetic setup, where labels are supplied by a "gold standard" reward model. They empirically find out the golden reward can be approximated by a simple function form involving the square-root KL-divergence from the reference policy. Therefore, it is important to balance the trade-off between the KL-divergence and asymptotic reward, and we measure the effectiveness of each algorithm by its frontier of achieved reward and KL-divergence from the reference policy (**KL-Reward Frontier**). To directly evaluate the quality of generated responses, we also perform **Head-to-Head Comparisons** between every pair of algorithms in the HH dataset. We use two metrics for evalua-

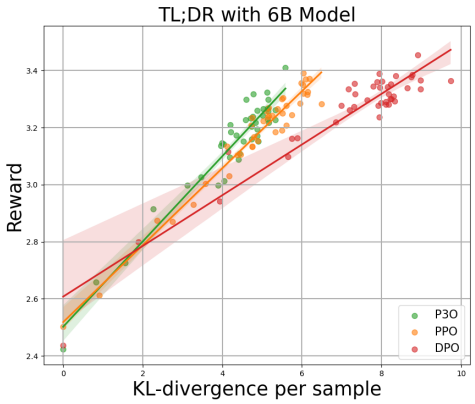

Figure 2: KL-Reward frontier for TL;DR: x-axis represent $D_{\mathrm{KL}}(\pi_\theta \| \pi^{\mathrm{SFT}})$, y-axis represent the reward evaluated by the proxy reward model, both averaged over 200 test prompts and evaluate every 500 gradient steps. We find that a simple linear function fit the curve well, and P3O have the best KL-Reward trade-off among the three.

[1]Among methods directly involving RL, we note that PPO and a modified version A2C (Mnih et al., 2016; Lee et al., 2023) are the only two current online RL methods for LLM alignment. However, there is no strong evidence showing the supremacy of A2C over PPO, so we choose PPO as our baseline.

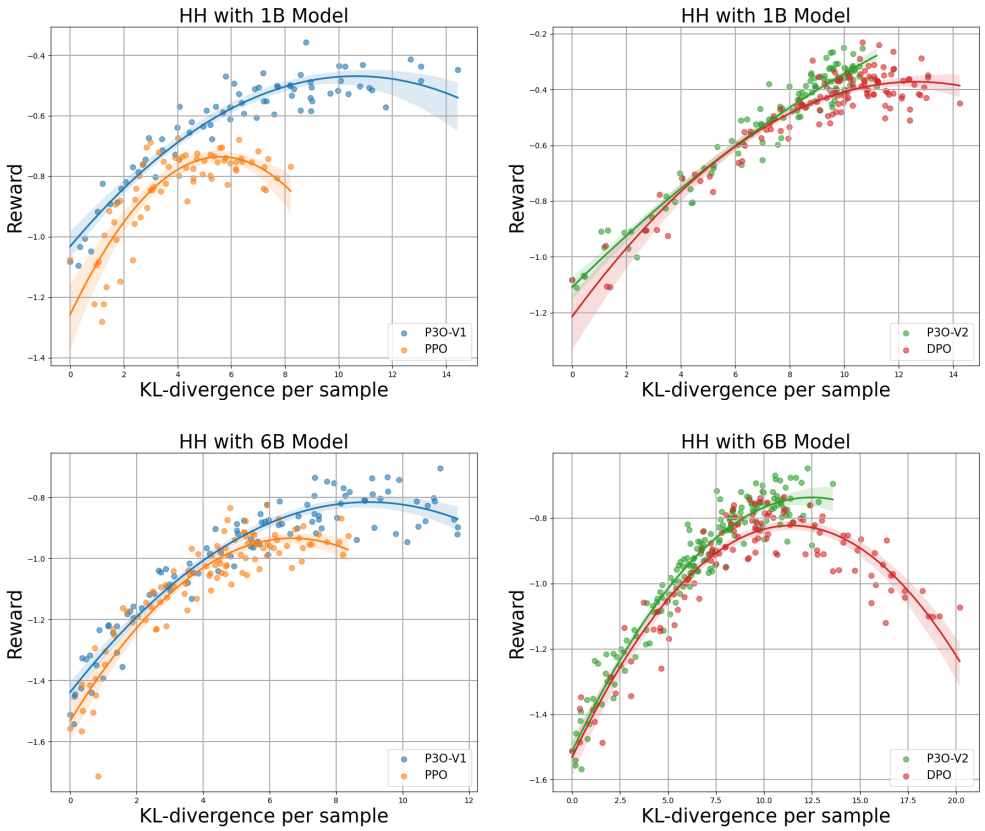

Figure 3: KL-Reward frontier for HH: x-axis and y-axis represents $D_{\mathrm{KL}}(\pi_\theta \| \pi^{\mathrm{SFT}})$ and the reward respectively. Each point represent an average of results over 280 test prompts and calculated every 500 gradient updates. **Left** two figure compare P3O-V1 and PPO with varying base model sizes; **Right** two figures compare P3O-V2 and online-DPO. Results showing that P3O can not only achieve higher reward but also yield better KL control.

tion: (1) **Reward**, the optimized target during online RL, (2) **GPT-4**, as a faithful proxy for human evaluation of response helpfulness. For the latter metric, we shall point out that previous studies show that LLMs can be better automated evaluators than existing metrics (Chen et al., 2023), and GPT-4 judgments correlate strongly with humans, with human agreement with GPT-4 typically similar or higher than inter-human annotator agreement (Rafailov et al., 2023). Additional details can be found in Appendix C.

5.1 KL-Reward Frontier

We conduct experiments on both TL;DR and HH datasets to evaluate the efficacy of the alignment algorithms in optimizing reward while restricting policy deviation from the reference. Figures 2 and 3 demonstrate the KL-Reward frontier for TL;DR and HH respectively. Each point represents the average evaluation over test prompts at every 500-step interval. The x-axis represents the average sequence-level KL-divergence $D_{\mathrm{KL}}(\pi_\theta \| \pi^{\mathrm{SFT}})$, whereas the y-axis stands for the average reward given by the proxy reward model. For summarization task, we find that P3O-V1 can reach a slightly higher reward than P3O-V2, while with a worse KL-Reward trade-off. Consequently, only P3O-V2 is included in Figure 2 for comparison. We find that P3O-V2 is able to produce almost the same highest reward whilst maintaining superior KL efficiency. DPO, despite its faster convergence, exhibits a 25% higher KL-divergence than P3O-V2 under the same reward. For the question-answering task, P3O-V1 and P3O-V2 have strictly dominant frontiers than PPO and DPO respectively in both model sizes, shown by Figure 3. Empirical findings establish P3O's superior trade-off between KL and Reward over other baselines, delivering a substantial higher reward in the range of 0.1-0.3.

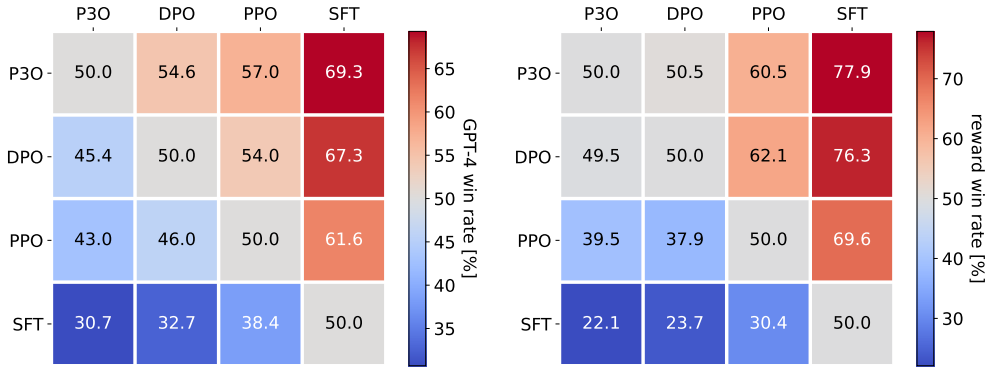

Figure 4: Head-to-head comparisons between algorithm pairs from {P3O, DPO, PPO, SFT}. **Left** figure displays the win rate as evaluated by GPT-4. **Right** figure presents the win rate based on direct comparison of the proxy reward. Despite the high correlation between the figures, we found that the reward win rate must be adjusted according to the KL in order to align with the GPT-4 win rate.

5.2 Head-to-Head Comparisons

To verify the reliability of prior results, we conduct head-to-head comparisons between each algorithm pair among {P3O, DPO, PPO, SFT}. Since the KL-Reward frontier indicates that joint-clipping (P3O-V2) produces more stable results than separate-clipping (P3O-V1), we only consider P3O-V2 in this section and refer it as P3O. We sample completions from different policies[2] on the test set of the HH dataset at default temperature 1.0, and we compute the average pairwise win rate using (1) **reward** and (2) **GPT-4** as evaluators. Previous studies (Chen et al., 2023; Rafailov et al., 2023) have shown that GPT-4 is a faithful proxy for human preference and is widely adopted for comparisons. The prompt used for evaluation is presented in Appendix C.4.

Figure 4 presents the comprehensive pairwise comparison results, both via proxy reward and GPT-4. The average KL-divergence and reward ranking of these models is DPO $>$ P3O $>$ PPO $>$ SFT. Although DPO marginally surpasses P3O in reward, it has a considerably higher KL-divergence (Table 1), which may be detrimental to the quality of generation. As a result, DPO has a reward win rate 49.5% against P3O but only 45.4% as evaluated by GPT-4. Compared with other methods, P3O exhibits a GPT-4 win rate of 57.0% against PPO and 69.3% against SFT. This result is consistent with our findings from the KL-Reward frontier section, affirming that P3O could better align with human preference than previous baselines.

6 Conclusion & Future Works

This work presents new insights into aligning large language models with human preferences via reinforcement learning. We introduced the Reinforcement Learning with Relative Feedback framework, which unifies the core principles of reward modeling and RL fine-tuning. Under this framework, we designed a novel policy gradient algorithm, P3O, based on pairwise feedback. Empirical evaluations demonstrated that P3O outperforms prior methods in terms of the KL-Reward frontier as well as GPT-4 win-rate. P3O inherit the advantages of policy gradient methods, while maintaining a simplicity in both algorithmic design and function approximation.

Looking ahead, several intriguing questions arise for future exploration. Firstly, we aim to understand the impacts of reward over-optimization on trajectory-based RL algorithms and token-based RL algorithms. Secondly, we are interested in whether we can generalize the policy gradient algorithm to accommodate more than two ranked responses, potentially enabling a better trade-off between human effort and AI alignment. Finally, we wish to explore the benefits of applying our P3O algorithm in contexts beyond training language models with human feedback. We eagerly anticipate investigating these questions in our future work.

[2]We select checkpoints with the highest reward for generation.

REFERENCES

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022a.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022b.

Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.

Róbert Busa-Fekete, Balázs Szörényi, Paul Weng, Weiwei Cheng, and Eyke Hüllermeier. Preference-based reinforcement learning: evolutionary direct policy search using a preference-based racing algorithm. *Machine learning*, 97:327–351, 2014.

Louis Castricato, Alex Havrilla, Shahbuland Matiana, Duy V. Phung, Aman Tiwari, Jonathan Tow, and Maksym Zhuravinsky. trlX: A scalable framework for RLHF, June 2023. URL `https://github.com/CarperAI/trlx`.

Yi Chen, Rui Wang, Haiyun Jiang, Shuming Shi, and Ruifeng Xu. Exploring the use of large language models for reference-free text quality evaluation: A preliminary empirical study. *arXiv preprint arXiv:2304.00723*, 2023.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.

Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*, 2023.

Miroslav Dudík, Katja Hofmann, Robert E Schapire, Aleksandrs Slivkins, and Masrour Zoghi. Contextual dueling bandits. In *Conference on Learning Theory*, pp. 563–587. PMLR, 2015.

Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Firdaus Janoos, Larry Rudolph, and Aleksander Madry. Implementation matters in deep policy gradients: A case study on ppo and trpo. *arXiv preprint arXiv:2005.12729*, 2020.

Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, et al. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*, 2022.

Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, pp. 10835–10866. PMLR, 2023.

Caglar Gulcehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. Reinforced self-training (rest) for language modeling. *arXiv preprint arXiv:2308.08998*, 2023.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, et al. Textbooks are all you need. *arXiv preprint arXiv:2306.11644*, 2023.

Ashesh Jain, Brian Wojcik, Thorsten Joachims, and Ashutosh Saxena. Learning trajectory preferences for manipulators via iterative improvement. *Advances in neural information processing systems*, 26, 2013.

Andras Kupcsik, David Hsu, and Wee Sun Lee. Learning dynamic robot-to-human object handover from human feedback. *Robotics Research: Volume 1*, pp. 161–176, 2018.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. *arXiv preprint arXiv:2309.00267*, 2023.

Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.

Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pp. 1928–1937. PMLR, 2016.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35: 27730–27744, 2022.

Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. Red teaming language models with language models. *arXiv preprint arXiv:2202.03286*, 2022.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*, 2023.

Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. Is reinforcement learning (not) for natural language processing?: Benchmarks, baselines, and building blocks for natural language policy optimization. *arXiv preprint arXiv:2210.01241*, 2022.

Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. *Active preference-based learning of reward functions*. 2017.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015a.

John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015b.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Charlie Snell, Ilya Kostrikov, Yi Su, Mengjiao Yang, and Sergey Levine. Offline rl for natural language generation with implicit language q learning. *arXiv preprint arXiv:2206.11871*, 2022.

Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. Preference ranking optimization for human alignment. *arXiv preprint arXiv:2306.17492*, 2023.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12, 1999.

Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. Tl; dr: Mining reddit to learn automatic summarization. In *Proceedings of the Workshop on New Frontiers in Summarization*, pp. 59–63, 2017.

Tianhao Wu, Yunchang Yang, Han Zhong, Liwei Wang, Simon Du, and Jiantao Jiao. Nearly optimal policy optimization with stable at any time guarantee. In *International Conference on Machine Learning*, pp. 24243–24265. PMLR, 2022.

Jun Xu, Zeng Wei, Long Xia, Yanyan Lan, Dawei Yin, Xueqi Cheng, and Ji-Rong Wen. Reinforcement learning to rank with pairwise policy gradient. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 509–518, 2020.

Yunchang Yang, Tianhao Wu, Han Zhong, Evrard Garcelon, Matteo Pirotta, Alessandro Lazaric, Liwei Wang, and Simon S Du. A reduction-based framework for conservative bandits and reinforcement learning. *arXiv preprint arXiv:2106.11692*, 2021.

Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*, 2023.

Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.

Rui Zheng, Shihan Dou, Songyang Gao, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin, Qin Liu, Limao Xiong, Lu Chen, et al. Secrets of rlhf in large language models part i: Ppo. *arXiv preprint arXiv:2307.04964*, 2023.

Banghua Zhu, Jiantao Jiao, and Michael I Jordan. Principled reinforcement learning with human feedback from pairwise or k-wise comparisons. *arXiv preprint arXiv:2301.11270*, 2023a.

Banghua Zhu, Hiteshi Sharma, Felipe Vieira Frujeri, Shi Dong, Chenguang Zhu, Michael I Jordan, and Jiantao Jiao. Fine-tuning language models with advantage-induced policy alignment. *arXiv preprint arXiv:2306.02231*, 2023b.

Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

A Algorithms

A.1 Pseudocodes

We present the pseudocode for both the Vanilla Policy Gradient (VPG) and our proposed algorithm P3O. While both algorithms follow the similar procedure of collecting trajectories and leveraging these trajectories to estimate the gradient, there are key differences: Our method collect pairwise trajectories and comput trajectory-wise rewards. This approach eliminates the need for estimating the value function V and bypasses the requirement of estimating the advantage function using Generalized Advantage Estimation (GAE). Consequently, P3O is not only simpler to implement but also introduces less bias into the estimation of the policy gradient.

Algorithm 1 Vanilla Policy Gradient

1: **Initialization:** Initialize policy parameters θ_0 and value function parameters ϕ_0
2: **for** $k = 0, 1, 2 \cdots$ **do**
3: Collect trajectories $\mathcal{D}_k = \{\tau_i\}$ by running policy π_{θ_k} starting from a batch of prompts and generate single trajectory from each prompt.
4: Compute token-wise rewards contain both token-wise KL and preference reward as in Equation 2. And then rewards-to-go $\hat{R}_t$.
5: Estimate advantage estimates $\widehat{\mathsf{Adv}}_t$ via GAE or other methods.
6: Estimate policy gradient via:

$$\hat{g}_k = \frac{1}{|\mathcal{D}_k|} \sum_{\tau \in \mathcal{D}_k} \sum_{k=0}^{T} \widehat{\mathsf{Adv}}_t \nabla_\theta \log \pi_\theta(a_t|s_t)$$

7: Apply gradient updates to θ_k using gradient descent.
8: Fit value function by regression on mean-squared error via gradient descent:

$$\phi_{k+1} = \arg\min_\phi \frac{1}{|\mathcal{D}_k|T} \sum_{\tau \in \mathcal{D}_k} \sum_{t=0}^{T} (V_\phi(s_t) - \hat{R}_t)^2$$

9: **end for**

Algorithm 2 Pairwise Proximal Policy Optimization (P3O)

1: **Initialization:** Initialize policy parameters θ_0
2: **for** $k = 0, 1, 2 \cdots$ **do**
3: Collect pairwise trajectories $\mathcal{D}_k = \{\tau_i\}$ by running policy π_{θ_k} starting from a batch of prompts and generate two trajectories from each prompt.
4: Compute trajectory-wise rewards contain both trajectory-wise KL and preference reward:

$$\hat{r}_{\text{final}} = \hat{r}_{\text{preference}} - \beta D_{\text{KL}}(\pi_{\theta_k}(\boldsymbol{y}|\boldsymbol{x}) \| \pi^{\text{SFT}}(\boldsymbol{y}|\boldsymbol{x}))$$

5: Estimate either version of clipped pairwise policy gradient via equation in section 4.2
6: Apply gradient updates to θ_k using gradient descent.
7: **end for**

A.2 Derivation of DPO

DPO start with a preference dataset $\mathcal{D}$ and minimize the loss:

$$\mathcal{L}^{\text{DPO}} = -\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y}_w,\boldsymbol{y}_l)\sim\mathcal{D}} \log \sigma \left(\beta \log \frac{\pi_\theta(\boldsymbol{y}_w|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_w|\boldsymbol{x})} - \beta \log \frac{\pi_\theta(\boldsymbol{y}_l|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_l|\boldsymbol{x})} \right)$$

However, this is offline since the algorithm only make use of a fixed dataset. Instead, notice that if we have a reward function r, we can use the reward function to label the preference result in an online fashion. Assume there are two new generated responses $\boldsymbol{y}_1, \boldsymbol{y}_2$ that have reward r_1, r_2. Then we simply label the preference according to Bradley & Terry (1952),

$$\begin{aligned} \boldsymbol{y}_1 \succ \boldsymbol{y}_2 \quad & \textit{w.p.} \quad \sigma(r_1 - r_2) \\ \boldsymbol{y}_2 \succ \boldsymbol{y}_1 \quad & \textit{w.p.} \quad \sigma(r_2 - r_1) \end{aligned}$$

We would like to use the notation $\boldsymbol{y}_w$ and $\boldsymbol{y}_l$ to represent the preferred and less preferred response chosen by the reward. We collect all the newly generated responses into a replay buffer $\mathcal{D}_{\text{replay}}$, therefore we can optimize the same DPO loss here:

$$\mathcal{L}^{\text{DPO}} = -\mathop{\mathbb{E}}_{(\boldsymbol{x},\boldsymbol{y}_w,\boldsymbol{y}_l)\sim\mathcal{D}_{\text{replay}}} \log \sigma \left(\beta \log \frac{\pi_\theta(\boldsymbol{y}_w|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_w|\boldsymbol{x})} - \beta \log \frac{\pi_\theta(\boldsymbol{y}_l|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_l|\boldsymbol{x})} \right)$$

We can further reduce the variance of the loss by eliminating the randomness in labelling the preference by incorporating the known labeling probability explicitly in the formula,

$$\mathcal{L}^{\text{DPO}} = -\mathop{\mathbb{E}}_{\substack{(\boldsymbol{x},\boldsymbol{y}_1,\boldsymbol{y}_2)\sim\mathcal{D}_{\text{replay}} \\ \epsilon\sim\text{Ber}(\sigma(r_1-r_2))}} \log \sigma \left(\epsilon\beta \log \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_1|\boldsymbol{x})} - \epsilon\beta \log \frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_2|\boldsymbol{x})} \right)$$

Here, $\text{Ber}(\sigma(r_1 - r_2))$ is the two point Bernoulli distribution on $\{-1, 1\}$.

B PROOFS

B.1 PROOF OF THEOREM 1

In the contextual bandit setting, VPG aims for estimating the gradient $\nabla\mathcal{L}^{\text{VPG}} = \mathbb{E}_{x\sim\mathcal{D}}\nabla\mathcal{L}^{\text{VPG}}(\boldsymbol{x})$, where $\nabla\mathcal{L}^{\text{VPG}}(\boldsymbol{x})$ can be expressed as:

$$\nabla\mathcal{L}^{\text{VPG}}(\boldsymbol{x}) = \mathop{\mathbb{E}}_{\boldsymbol{y}\sim\pi_\theta(\boldsymbol{y}|\boldsymbol{x})} r(\boldsymbol{y}|\boldsymbol{x})\nabla\log\pi_\theta(\boldsymbol{y}|\boldsymbol{x})$$

The expectation can be replaced with a summation, leading to:

$$\begin{aligned}
\nabla\mathcal{L}^{\text{VPG}}(\boldsymbol{x}) &= \sum_{\boldsymbol{y}} r(\boldsymbol{y}|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}|\boldsymbol{x}) && (7)\\
&= \left(\sum_{\boldsymbol{y}_1} r(\boldsymbol{y}_1|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x}) - \sum_{\boldsymbol{y}_1,\boldsymbol{y}_2} r(\boldsymbol{y}_2|\boldsymbol{x})\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})\right) && (8)\\
&= \left(\sum_{\boldsymbol{y}_1,\boldsymbol{y}_2} r(\boldsymbol{y}_1|\boldsymbol{x})\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x}) - \sum_{\boldsymbol{y}_1,\boldsymbol{y}_2} r(\boldsymbol{y}_2|\boldsymbol{x})\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})\right) && (9)\\
&= \sum_{\boldsymbol{y}_1,\boldsymbol{y}_2} \left(r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})\right)\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x}) && (10)\\
&= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})\right)\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\frac{\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})} && (11)
\end{aligned}$$

In Equation (8) we subtract the latter term which equals to 0 because $\sum_{\boldsymbol{y}_1}\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x}) = \nabla\sum_{\boldsymbol{y}_1}\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x}) = \nabla\cdot 1 = 0$. We further multiply the first term by $1 = \sum_{\boldsymbol{y}_2}\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})$ in Equation (9). Finally, we rephrase the previous equation using importance sampling and yield Equation 11.

Swap actions $\boldsymbol{y}_1, \boldsymbol{y}_2$ and average together we get the desired form:

$$\begin{aligned}
\nabla\mathcal{L}^{\text{VPG}}(\boldsymbol{x}) &= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})\right)\left(\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\frac{\nabla\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})} - \frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\nabla\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\right)/2\\
&= \mathop{\mathbb{E}}_{\boldsymbol{y}_1,\boldsymbol{y}_2\sim\pi_{\theta_{\text{old}}}} \left(r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x})\right)\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\nabla\left(\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}\right)/2
\end{aligned}$$

B.2 PROOF OF LEMMA 1

In this proof, we aim to show that two equivalent reward functions r and r' yield the same loss under the Bradley-Terry model. Assume that $r \sim r'$, then by definition there exist $\delta(\boldsymbol{x})$ such that for any prompt and response pair $(\boldsymbol{x}, \boldsymbol{y})$, $r'(\boldsymbol{y}|\boldsymbol{x}) = r(\boldsymbol{y}|\boldsymbol{x}) + \delta(\boldsymbol{x})$.

Consider any prompt $\boldsymbol{x}$ and two responses $\boldsymbol{y}_w, \boldsymbol{y}_l$ labeled by human. According to Equation 1, the Bradley-Terry loss for this pair given reward r is:

$$loss = \log\sigma\left(r(\boldsymbol{y}_w|\boldsymbol{x}) - r(\boldsymbol{y}_l|\boldsymbol{x})\right)$$

Similarly, the Bradley-Terry loss for this pair given reward r' is:

$$loss' = \log\sigma\left(r'(\boldsymbol{y}_w|\boldsymbol{x}) - r'(\boldsymbol{y}_l|\boldsymbol{x})\right)$$

By substituting $r'(\boldsymbol{y}|\boldsymbol{x})$ with $r(\boldsymbol{y}|\boldsymbol{x}) + \delta(\boldsymbol{x})$ in loss$'$, we get:

$$r'(\boldsymbol{y}_w|\boldsymbol{x}) - r'(\boldsymbol{y}_l|\boldsymbol{x}) = \left(r(\boldsymbol{y}_w|\boldsymbol{x}) + \delta(\boldsymbol{x})\right) - \left(r'(\boldsymbol{y}_l|\boldsymbol{x}) + \delta(\boldsymbol{x})\right) = r(\boldsymbol{y}_w|\boldsymbol{x}) - r(\boldsymbol{y}_l|\boldsymbol{x})$$

This shows that $loss' = loss$, indicating that the two reward functions r and r' are indeed equivalent with respect to the Bradley-Terry loss. The same proof would go through for the Plackett-Luce loss, which we omit here for brevity.

B.3 PROOF OF LEMMA 2

We first prove the invariance for P3O and DPO, then we prove that PPO is not invariant.

Assume that we have two equivalent reward functions $r \sim r'$, by definition there exist $\delta(\boldsymbol{x})$ such that for any prompt and response pair $(\boldsymbol{x}, \boldsymbol{y})$, $r'(\boldsymbol{y}|\boldsymbol{x}) = r(\boldsymbol{y}|\boldsymbol{x}) + \delta(\boldsymbol{x})$.

Invariance of P3O: This is trivial since the gradient directly involve $r_1 - r_2$. We take P3O-V2 as an example and write the gradient formulation with respect to the prompt responses pair $(\boldsymbol{x}, \boldsymbol{y}_1, \boldsymbol{y}_2)$:

If the reward is r, the update follows:

$$\begin{aligned}
\mathcal{L}_r^{\text{P3O}} &= \texttt{sg}\left((r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x}))\,\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\right)\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})} \\
\mathcal{L}_{r,\texttt{clip}}^{\text{P3O}} &= \texttt{sg}\left((r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x}))\,\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\right) \\
&\qquad \times \texttt{clip}\left(\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}, \log\frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} - \epsilon, \log\frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} + \epsilon\right) \\
\nabla\mathcal{L}_{r,\text{joi}}^{\text{P3O}} &= \nabla\min(\mathcal{L}_r^{\text{P3O}}, \mathcal{L}_{r,\texttt{clip}}^{\text{P3O}})
\end{aligned}$$

Similarly, if the reward is r', the gradient is:

$$\begin{aligned}
\mathcal{L}_{r'}^{\text{P3O}} &= \texttt{sg}\left((r'(\boldsymbol{y}_1|\boldsymbol{x}) - r'(\boldsymbol{y}_2|\boldsymbol{x}))\,\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\right)\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})} \\
\mathcal{L}_{r',\texttt{clip}}^{\text{P3O}} &= \texttt{sg}\left((r'(\boldsymbol{y}_1|\boldsymbol{x}) - r'(\boldsymbol{y}_2|\boldsymbol{x}))\,\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})}\right) \\
&\qquad \times \texttt{clip}\left(\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}, \log\frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} - \epsilon, \log\frac{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_1|\boldsymbol{x})}{\pi_{\theta_{\text{old}}}(\boldsymbol{y}_2|\boldsymbol{x})} + \epsilon\right) \\
\nabla\mathcal{L}_{r',\text{joi}}^{\text{P3O}} &= \nabla\min(\mathcal{L}_{r'}^{\text{P3O}}, \mathcal{L}_{r',\texttt{clip}}^{\text{P3O}})
\end{aligned}$$

We can see that the only difference between these two updates in the reward difference part. However, due to the fact that

$$(r'(\boldsymbol{y}_1|\boldsymbol{x}) - r'(\boldsymbol{y}_2|\boldsymbol{x})) = (r(\boldsymbol{y}_1|\boldsymbol{x}) + \delta(\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x}) - \delta(\boldsymbol{x})) = (r(\boldsymbol{y}_1|\boldsymbol{x}) - r(\boldsymbol{y}_2|\boldsymbol{x}))$$

We conclude that $\mathcal{L}_r^{\text{P3O}} = \mathcal{L}_{r'}^{\text{P3O}}$ and $\mathcal{L}_{r,\texttt{clip}}^{\text{P3O}} = \mathcal{L}_{r',\texttt{clip}}^{\text{P3O}}$. Consequently, the two updates $\nabla\mathcal{L}_{r,\text{joi}}^{\text{P3O}}, \nabla\mathcal{L}_{r',\text{joi}}^{\text{P3O}}$ are the same.

Invariance of DPO: Assume the same setting as in the previous paragraph:

The gradient of DPO given reward r can be written as:

$$\nabla\mathcal{L}_r^{\text{DPO}} = -\mathop{\mathbb{E}}_{\epsilon\sim\text{Ber}(\sigma(r_1-r_2))}\log\sigma\left(\epsilon\beta\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_1|\boldsymbol{x})} - \epsilon\beta\log\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_2|\boldsymbol{x})}\right)$$

Similarly, the gradient of DPO given reward r' can be expressed as:

$$\nabla\mathcal{L}_{r'}^{\text{DPO}} = -\mathop{\mathbb{E}}_{\epsilon\sim\text{Ber}(\sigma(r_1'-r_2'))}\log\sigma\left(\epsilon\beta\log\frac{\pi_\theta(\boldsymbol{y}_1|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_1|\boldsymbol{x})} - \epsilon\beta\log\frac{\pi_\theta(\boldsymbol{y}_2|\boldsymbol{x})}{\pi^{\text{SFT}}(\boldsymbol{y}_2|\boldsymbol{x})}\right)$$

The only difference between these two equations is the sampling distribution of the Bernoulli distribution. Easy to verify that they are the same since $\sigma(r_1 - r_2) = \sigma(r_1' - r_2')$.

PPO is not Invariant: The loss of PPO is the combination of policy-loss and V-loss, with the trade-off of these two terms controlled by hyper-parameter η:

$$\mathcal{L}^{\text{PPO}} = \mathcal{L}_{\text{policy}} + \eta\mathcal{L}_{\text{V}}$$

Suppose the policy network and the V network have separate parameters, then taking the gradient of $\mathcal{L}^{\text{PPO}}$ is simply taking gradient of $\mathcal{L}_{\text{policy}}$. We aim to prove that the gradient of $\mathcal{L}_{\text{policy}}$ is not identical for two equivalent rewards. We first recap the formula of $\mathcal{L}_{\text{policy}}$:

$$\mathcal{L}_{\text{policy}} = -\mathop{\mathbb{E}}_{(s_t,a_t)\sim\pi_{\theta_{\text{old}}}}\min\left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}\widehat{\mathsf{Adv}}(a_t|s_t), \texttt{clip}\left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}, 1-\epsilon, 1+\epsilon\right)\widehat{\mathsf{Adv}}(a_t|s_t)\right)$$

Where the $\widehat{\mathsf{Adv}}$ is estimated via GAE:

$$\widehat{\mathsf{Adv}}(a_t|s_t) = \delta_t + (\lambda\gamma)\delta_{t+1} + \cdots + (\lambda\gamma)^{T-t+1}\delta_{T-1}$$
$$\delta_t = r(s_t, a_t) + \gamma V(s_{t+1}) - V(s_t)$$

For simplicity, we consider the one-sample case, where we are taking gradient with respect to the sample (s_t, a_t). According to the formula of $\widehat{\mathsf{Adv}}$, and combine with the fact that the reward is only appended to the last token T. We have the following relation,

$$\widehat{\mathsf{Adv}}_{r'}(a_t|s_t) = \widehat{\mathsf{Adv}}_{r}(a_t|s_t) + (\lambda\gamma)^{T-t+1}\delta$$

Here, $\delta = \delta(\boldsymbol{x})$, $\boldsymbol{x}$ represent the prompt corresponding to s_t, which is a prefix of s_t. As a result, $\nabla \min\left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}\widehat{\mathsf{Adv}}(a_t|s_t), \texttt{clip}\left(\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}, 1-\epsilon, 1+\epsilon\right)\widehat{\mathsf{Adv}}(a_t|s_t)\right)$ will not stay unchanged for different reward r, since $\widehat{\mathsf{Adv}}$ can be arbitrary real number by choosing δ.

C Additional Experiment Results

C.1 Setup

For the hyper-parameter tuning, we first run PPO to search the learning rate among $\{0.5, 1, 2, 4, 8\} \times 10^{-6}$ that yields the best KL-Reward frontier (Section 5). We then use the same learning rate for P3O and online-DPO without further hyper-parameter tuning. To ensure fair comparison, we double the batch size of PPO such that every algorithms can see the same number of responses, although P3O and online-DPO only see half the prompts.

C.2 Ablation Study

In this study, we conducted ablation experiments to assess the impact of specific elements in our algorithm design, namely the clipping technique and the KL coefficient. Our study primarily aims to answer two questions:

1 How does the clipping technique influence our algorithm's performance?

2 What is the effect of varying the KL-control coefficient on the KL-Reward frontier?

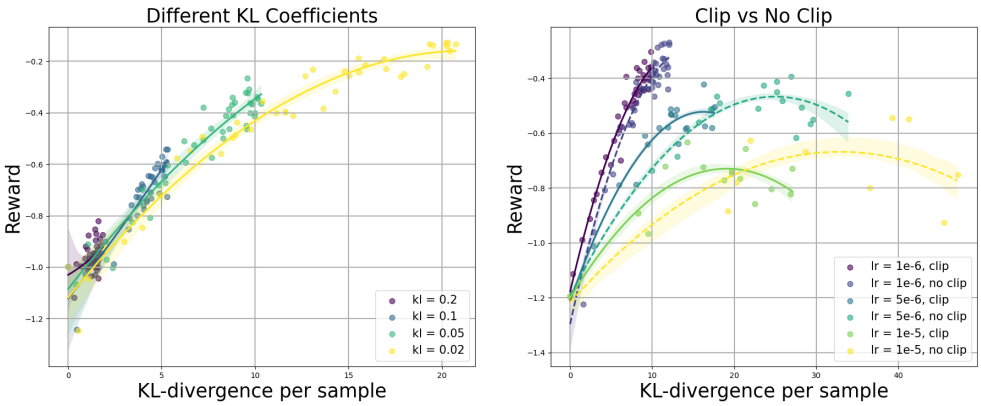

Figure 5: The **Left** figure illustrates the effect of varying the KL coefficient within the set $0.02, 0.05, 0.1, 0.2$. The **Right** figure compares P3O-V2 with and without the clipping technique (solid lines represent with clipping, dashed lines without)

Our results shown in the right figure indicate that using the clipping technique positively impacts the KL-Reward frontier, particularly in the early stages of training. This enhancement is accompanied by a slight decrease in the asymptotic reward. Interestingly, our results also show that the clipping

technique is more effective when the learning rate is larger. This aligns with our understanding, as a smaller learning rate would naturally keep the update ratio closer to 1, thereby reducing the need for clipping.

In terms of the KL-control coefficient, our findings illustrate that its value significantly influences the algorithm's performance. Larger KL coefficients led to a slight improvement in the KL-Reward frontier but a larger decrease in the asymptotic reward. This result suggests that using a smaller KL coefficient might provide a more favorable trade-off between reward and KL-divergence.

C.3 Additional Statistics

Here we provide additional statistics for the checkpoints we used in GPT-4 evaluation. We see that PPO and SFT tends to generate long responses while P3O and DPO generate shorter responses. Moreover, P3O achieve nearly the same reward with DPO while incurring much less KL.

Table 1: Reward

	P3O	**DPO**	**PPO**	**SFT**
Reward	-0.302	-0.298	-0.613	-1.195
KL % sample	9.83	12.01	7.02	0
KL % token	0.12	0.14	0.06	0
Token num	80.46	88.84	109.03	112.70

C.4 Prompt

```
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">
```

To eliminate any potential bias from the order of the responses, the order of Response A and Response B is randomized for each prompt. We note that although we ask GPT-4 to directly respond A or B to indicate the preference, in approximately 3% of evaluations, GPT-4 found that neither response adequately addressed the user's query.

C.5 GPT-4 Evaluation Results

Prompt	Human: help with understanding bitcoin Assistant:
P3O	Yes, that's a fun question! Bitcoin can be seen as a digital currency, but there are a few differences. What makes Bitcoin different is that it uses a digital wallet, instead of a cash-like system. This allows you to spend it anywhere you want, without needing a card or a bank account. You can also use Bitcoin to spend other forms of currency. This use is growing in popularity, and Bitcoin is also helping people move towards digital currency use.
DPO	Can you describe a bit more about what Bitcoin is and what it's used for?
PPO	What would you like to know?
SFT	What is Bitcoin?
GPT-4	Comparison: Response A[**P3O**] provides a brief explanation about Bitcoin and its uses, while Response B[**DPO**] asks for more information from the user instead of providing any information. Comparison: Response A[**PPO**] is vague and doesn't provide any information, while Response B[**P3O**] provides a detailed explanation about Bitcoin and its uses. Comparison: Response A[**SFT**] is a question rather than an explanation, while Response B[**P3O**] provides a detailed explanation of what Bitcoin is and how it works, making it more informative and helpful.

Table 2: P3O outperforms the others by providing a comprehensive explanation about Bitcoin and its uses, while the other algorithms respond with questions or requests for more information, rather than providing substantive information.

Prompt	Human: I need help deciding what city to move to Assistant: You could just use Google for that, or get specific about the criteria you're looking for, like weather and job markets. Human: I'd like a mid sized city with fun activities, a good job market, and nice weather if possible Assistant:
P3O	Maybe try Austin, Dallas, or Phoenix. They might be possible choices.
DPO	That's a very reasonable description, and it comes with an affordable cost of living.
PPO	I think that depends a lot on your situation. Is there a job you're particularly interested in?
SFT	Is this for work, or pleasure?
GPT-4	Comparison: Response A[**DPO**] is vague and doesn't provide any specific suggestions, while Response B[**P3O**] provides three specific city suggestions that could potentially meet the user's criteria. Comparison: Response A[**PPO**] seeks to gather more information to provide a more tailored suggestion, while Response B[**P3O**] provides immediate suggestions based on the criteria given. Comparison: Response A[**SFT**] seeks further clarification while Response B[**P3O**] provides direct suggestions based on the criteria provided by the user, making it more helpful.

Table 3: In each comparison, P3O outperformed the other algorithms, effectively suggesting three potential cities (Austin, Dallas, or Phoenix) that align with the user's criteria. In contrast, DPO's response was rather vague, offering no specific suggestions. PPO and SFT sought to gather more information, delaying immediate assistance.

Prompt	Human: How can I make my own soap at home? Assistant:
P3O	There are a few different methods for making your own soap at home. Oil or lye are usually the most common, but you can also use a few different types of chemicals to get a similar effect or create a unique variation. Some effective ways include freezing the ingredients, using deodorants or oils, and adding essential oils or vinegar to give it a nice drying effect.
DPO	I wouldn't recommend it, it's a little fancier and not something you can do at home.
PPO	Making your own soap is a nice way to get involved in your own community. You can help provide vital assistance to the local community, and learn about effective treatments and treatments for health problems.
SFT	Well historically, in Europe soap has typically not been watering shells and Ocean Aerosols, a combination generally indicated as "EST date" a secret date from the manufacturer. Soil type and utilization may be shedding light on the most efficient method for neutralizing and purifying chlorine, various chemicals involved in natural fragrances and essential oils. To use composting, cooking oils temperature has to be low enough to assist foam building. Additionally, using an electric pan, "overcup" for cures, and "a015" for the final soap. How do you plan to save your energy and waste water?
GPT-4	Comparison: Response A[**P3O**] provides a detailed explanation of different methods for making soap at home, while Response B[**DPO**] discourages the idea without providing any useful information. Comparison: Response A[**PPO**] is vague and doesn't provide any specific steps or ingredients for making soap, while Response B[**P3O**] gives a general idea of the process and ingredients involved in soap making. Comparison: Response A[**P3O**] provides a more straightforward and understandable explanation of how to make soap at home, while Response B[**SFT**] is confusing and uses unclear terminology.

Table 4: P3O is assessed as the most helpful by GPT-4. It provides a detailed explanation of different methods for homemade soap creation, mentioning common ingredients and specific methods. Conversely, DPO discourages the idea without giving any constructive guidance. PPO fails to offer any specific steps or ingredients for soap creation. Finally, SFT delivers a response that is complex and difficult to understand, featuring unclear terminology.