

When Human Preferences Flip: An Instance-Dependent Robust Loss for RLHF

Yifan Xu,¹ Xichen Ye,² Yifan Chen,^{1,*} Qiaosheng Zhang³

¹Hong Kong Baptist University

²Fudan University

³Shanghai Artificial Intelligence Laboratory

Abstract

Quality of datasets plays an important role in large language model (LLM) alignment. In collecting human feedback, however, *preference flipping* is ubiquitous and causes corruption in data annotation; the issue necessitates the alignment algorithms with improved robustness against potential flipped pairs. To this end, this paper introduces a Flipping-Aware Direct Preference Optimization (FA-DPO) algorithm tailored to preference flipping from a reinforcement learning with human feedback (RLHF) perspective. We dissect the inherent human intention model and the preference flipping mechanism introduced by external factors as two distinct stages; in the latter, we introduce an instance-dependent flipping probability on the basis of the Bradley-Terry (BT) model. Further, by leveraging features relevant to preference annotation, we capture uncertainty in judgments and model preference flipping patterns. In practice, we design a simple yet efficient iterative optimization algorithm compatible with the original RLHF and Direct Preference Optimization (DPO) algorithms. In our experiments, we investigate the instance-dependent preference flipping model under multiple circumstances for evaluation of our proposed method, as well as other baseline methods.

Code — <https://github.com/bcxyf123/Noise-DPO>

Extended version — <https://arxiv.org/abs/2512.00709>

1 Introduction

Alignment has been identified as a crucial approach contributing to the remarkable capacity of large language models (LLMs) to understand human intentions. In general, this approach enables LLMs to produce responses that align well with human expectations and reduce toxic generations (Achiam et al. 2023; Dubey et al. 2024). Among existing alignment paradigms, a notable one is the reinforcement learning from human feedback (RLHF), which, along with its numerous variants, has attracted significant attention (Ouyang et al. 2022; Casper et al. 2023).

Despite the effectiveness of current RLHF methods, they implicitly suffer from noise in human feedback data (Zheng et al. 2023; Gao, Alon, and Metzler 2024). For example,

Gao, Alon, and Metzler (2024) reported that a 10% increase in preference flipping ratios can result in a 30% decrease in alignment performance, as measured by *the win rate* (defined in Section 5.1). As noise is inevitably introduced during data collection process, robust alignment algorithms bring the benefits of not only defending potential dataset attacks, but also reducing the costs of collecting clean data. One existing genre of robust learning approaches leverages noise information inferred from existing data (Xia et al. 2021; Song et al. 2022), and this characteristic is especially attractive in the contexts of direct alignment algorithms such as direct preference optimization (DPO; Rafailov et al. 2024), where training relies on a fixed offline dataset.

In tackling robustness challenges in RLHF, previous research focuses on a simplified scenario where preferences are randomly flipped at a *fixed* rate (Chowdhury, Kini, and Natarajan 2024; Wu et al. 2024a; Cheng et al. 2024; Liang et al. 2024). However, in practical applications, it is ungrounded to assume the possibility of annotation errors is independent of the specific content being annotated. In this work, we instead investigate a setting of “instance-dependence” (Xia et al. 2020; Liu, Cheng, and Zhang 2023), which poses greater challenges as the noise distribution can vary significantly across different samples. Therefore, in this work, we aim to answer the following question:

How can we properly model the instance-dependent preference flipping incurred during preference data annotation?

In response, we propose a statistically consistent approach to align with corrupted human feedback, which we refer to as Flipping-Aware Direct Preference Optimization (FA-DPO). This approach ensures that, given accurate estimation of the flipping model, the learned policy will achieve consistency comparable to learning the policy under clean data. We first explicitly model the flipping probability on the basis of the Bradley-Terry model, to connect with the posterior probabilities of the observed label.

Specifically, the annotation process under our framework can be viewed as two sequential stages: ❶ labeling according to true human intention and ❷ instance-dependent label contamination, where the instance-dependent preference flipping occurs after the labeling stage ❶, representing a transition from the true label to the flipped label for each

*Correspondence to: Yifan Chen (yifanc@hkbu.edu.hk).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

sample. In this way, given a corrupted human preference dataset, FA-DPO can post-train the LLM parameters robustly and produce optimal policies, as with the true labels, via correction for the original BT model loss with the posterior of the observed corrupted labels.

In particular, we further introduce an iterative optimization framework that jointly optimizes both the flipping estimation model and the LLM in post-training. To capture the transition between the original label and the observed label, we leverage a classification module on relevant features to construct the flipping probability within a preference pair. The main contributions of this paper are three-fold:

1. We investigate a challenging RLHF / DPO setting with instance-dependent preference flipping, and propose a novel probabilistic model to characterize the noise.
2. We address the instance-dependent estimation of preference flipping probabilities via a classification model that incorporates informative preference features studied in natural language processing.
3. We propose a simple yet efficient iterative optimization algorithm based on RLHF and DPO to capture the noise in the offline dataset during post-training. The whole training pipeline is validated in our experiments.

2 Related Works

We review the previous works on preference alignment and (generalized) robust RLHF in this section. Due to context limitation, we leave the complete review to Appendix B for the reader’s convenience.

Preference alignment The most well-known approach for preference alignment is Reinforcement Learning from Human Feedback (Ziegler et al. 2019; Ouyang et al. 2022, RLHF), which involves training a *reward model* to capture human preferences and then guiding LLMs to *generate high-reward responses using reinforcement learning (RL) algorithms* such as Proximal Policy Optimization (PPO) (Schulman et al. 2017). However, in practice RL-based methods can be complex and unstable during training (Rafailov et al. 2024; Wu et al. 2024c; Yuan et al. 2023). As a result, recent research has focused on simpler and more stable alternatives to RLHF, namely, direct preference alignment (Rafailov et al. 2024; Zhao et al. 2023; Ethayarajh et al. 2024; Azar et al. 2024; Meng, Xia, and Chen 2024).

RLHF against perturbations Current approaches of robust RLHF can be categorized into three main types.

❶ Noise fitting (Bukharin et al. 2024) involves making assumptions on the noise in the data and incorporating this modeling into the reward learning process, which will be jointly optimized with the parameterized reward function. ❷ Sample selection (or sample re-weighting) methods (Cheng et al. 2024) leverage this phenomenon to identify clean samples based on the loss values observed during the training process. Some approaches address corruption in input data through ❸ robust loss design (Chowdhury, Kini, and Natarajan 2024; Liang et al. 2024), focusing on constructing loss functions that are more resistant to data noise; this technique is also known as label smoothing (Mitchell 2023).

3 Preliminaries

In this section, we introduce the basics of RLHF for LLM alignment in advance of the formal proposal of our methodology in Section 4.

3.1 RLHF for LLM Alignment

Let an input (prompt) of an LLM be $x \in \mathcal{X}$ and the generated output (response) be $y \in \mathcal{Y}$. In the formulation for RLHF (Rafailov et al. 2024), the LLM is viewed as a policy $\pi_\theta(y | x)$ parameterized by θ , which outputs an action y (response) based on the state x (prompt). Preference data is further collected and annotated by human labelers, denoted as $y_w \succ y_l | x$, where y_w is the preferred response and y_l is the dispreferred one in (y_1, y_2) for the prompt x .

The pre-trained model first undergoes on round of supervised fine-tuning (SFT), resulting in a reference model π_{ref} . Then the BT model (Bradley and Terry 1952) is employed to relate the preference data $\{y_w^i \succ y_l^i\}$ to a reward model $r(x, y)$. The connection is formulated as:

$$p^*(y_w \succ y_l | x) = \sigma(r^*(x, y_w) - r^*(x, y_l)), \quad (1)$$

where σ is the standard sigmoid function, and $r^*(\cdot)$ is the optimal reward model. Using the BT model and the maximum likelihood principle, the loss for learning the reward model is:

$$\mathcal{L}_R(\phi) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\log \sigma(r_\phi(x, y_w) - r_\phi(x, y_l))], \quad (2)$$

where $r_\phi(\cdot)$ is the reward model parameterized by ϕ .

After training the reward model $r_\phi(\cdot)$ from Equation (2), RL is applied to optimize the LLM π_θ with the reward signals provided by $r_\phi(\cdot)$. The optimization objective for π_θ is formulated as:

$$\mathcal{L}_\pi(\theta) = -\mathbb{E}_{x \sim \mathbb{P}_x, y \sim \pi_\theta(\cdot | x)} [r_\phi(x, y)] + \beta \cdot \mathbb{E}_{x \sim \mathbb{P}_x} [\mathbb{D}_{\text{KL}}(\pi_\theta(\cdot | x) || \pi_{\text{ref}}(\cdot | x))], \quad (3)$$

where \mathbb{P}_x represents the marginal distribution of the prompt x . The first term corresponds to reward maximization in standard RL optimization. The second term is a KL divergence that constrains the update of π_θ to not deviate from the reference model π_{ref} , β is the coefficient that weights the KL divergence between π_θ and π_{ref} .

3.2 DPO as Efficient RLHF

Direct preference optimization (DPO) is an algorithm proposed for practical efficiency (compared to original RLHF), that merges the two stages in RLHF, reward modeling and policy optimization, into a single step. In DPO, the policy is optimized directly using the offline dataset \mathcal{D} without constructing an explicit reward model.

As shown by Peng et al. (2019), the closed-form solution of the conditional distribution $\pi(y | x)$ that minimizes Equation (3) is:

$$\pi_r(y | x) = \frac{1}{Z(x)} \pi_{\text{ref}}(y | x) \exp\left(\frac{1}{\beta} r_\phi(x, y)\right), \quad (4)$$

where $Z(x) = \sum_y \pi_{\text{ref}}(y | x) \exp(1/\beta \cdot r_\phi(x, y))$ is the partition function for normalization. By applying this result

to Equation (2), the DPO loss function that includes only the parameterized π_θ as the optimization variable is derived as:

$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}} [\log \sigma(\hat{r}_\theta(x, y_w) - \hat{r}_\theta(x, y_l))], \quad (5)$$

where $\hat{r}_\theta(x, y) = \beta \log(\pi_\theta(y|x)/\pi_{\text{ref}}(y|x))$ is the implicit reward model derived from π_θ .

4 Method

In this section, we begin by providing an overview of the motivation to model instance-dependent preference flipping. We then formulate the flipping procedure under the RLHF framework, and illustrate how this estimation can be integrated into the standard RLHF pipeline. Finally, we present a detailed model design for the instance-dependent preference flipping probability estimation in FA-DPO.

4.1 Motivations

It is widely accepted that in real-world scenarios, the intrinsic preference labeling mechanism of humans can be unified as a general human intention model, represented by the Bradley-Terry model (Bai et al. 2022). Noisy human feedback, however, usually stems from external factors rather than inherent errors in human decisions. For instance, environmental distractions may compromise an annotator’s focus, reducing labeling accuracy, and external noise can also be introduced by maliciously altering original annotations.

Consequently, this corruption process can be viewed as a post-transition applied to the initial labeling mechanism governed by the human intention model. We model this process as instance-dependent preference flipping, positing that the probability of a flip is correlated with the data content as well as the relation between each pair of responses.

We dig into the statistical structure of the flipping pattern in the preference data as opposed to the noise sparsity assumption in (Bukharin et al. 2024). The proposed modeling, as stated in Section 4.2, is supposed to exploit the relation between the noisy preference posterior and the true likelihoods. The neural network module to capture the noise is introduced in Section 4.3.

4.2 Flipping-aware Loss

We start the derivation of the flipping-aware loss from the RLHF model, and then extend the result to the DPO setting. Given a corrupted dataset $\tilde{\mathcal{D}}$, for each noisy triplet $(x, \tilde{y}_w, \tilde{y}_l)$, we denote the sample as \tilde{x} for simplicity. Following the principle of maximum likelihood, the usual loss for preference modeling is:

$$\mathcal{L}_{\text{MLE}} = -\mathbb{E}_{(x, \tilde{y}_w, \tilde{y}_l) \sim \tilde{\mathcal{D}}} [\log \mathbb{P}\{\tilde{y}_w \succ \tilde{y}_l | x\}]. \quad (6)$$

In standard RLHF, we only have the parameterized preference probability under clean data, i.e., $\mathbb{P}\{y_w \succ y_l | x\}$; the direct usage of the standard RLHF loss in Equation (6) can be sub-optimal. To ease the discussion of how to bridge the gap between preference distributions under clean and noisy data, we first formalize the preference flipping process through the following proposition:

Proposition 4.1 (Instance-dependent preference flipping). *For any input \tilde{x} , the corrupted preference probability, under the instance-dependent preference flipping setting, relates to the true preference likelihood via:*

$$\tilde{\mathbb{P}}\{\tilde{y}_w \succ \tilde{y}_l | x\} = (1 - \varepsilon_{\tilde{x}})p + \varepsilon_{\tilde{x}}(1 - p),$$

where $\varepsilon_{\tilde{x}}$ represents the instance-specific flipping probability for triplet $\tilde{x} = (x, \tilde{y}_w, \tilde{y}_l)$, and p denotes the true likelihood $\mathbb{P}\{y_w \succ y_l | x\}$ for brevity.

The proof of Proposition 4.1 is direct and omitted. We can then establish the relation between the corrupted posterior $\tilde{\mathbb{P}}\{\tilde{y}_w \succ \tilde{y}_l | x\}$ observed in noisy data and the underlying clean probability $\mathbb{P}\{y_w \succ y_l | x\}$. This enables us to recover the true preference probabilities by accounting for the instance-dependent flipping process:

$$\mathcal{L}_{\text{FA-DPO}} = -\mathbb{E}_{\tilde{x} \sim \tilde{\mathcal{D}}} [\log((1 - \varepsilon_{\tilde{x}})p + \varepsilon_{\tilde{x}}(1 - p))], \quad (7)$$

which reads the loss design in FA-DPO. For practical implementation with the BT model in reward-based RLHF, we parameterize p with ϕ and denote:

$$p_\phi = \sigma(r_\phi(x, \tilde{y}_w) - r_\phi(x, \tilde{y}_l)).$$

Similarly, for the DPO parameterization, the preference probability is given by substituting the probabilistic modeling Equation (4) into the formula above:

$$p_\theta = \sigma\left(\beta \log \frac{\pi_\theta(\tilde{y}_w | x)}{\pi_{\text{ref}}(\tilde{y}_w | x)} - \beta \log \frac{\pi_\theta(\tilde{y}_l | x)}{\pi_{\text{ref}}(\tilde{y}_l | x)}\right).$$

Comparison with cDPO and rDPO. To better understand FA-DPO, we compare it with related approaches, cDPO (Mitchell 2023) and rDPO (Chowdhury, Kini, and Natarajan 2024), from the perspective of gradients.

First, we recall standard DPO directly substitutes the parameterized preference probability $\mathbb{P}_\theta\{y_w \succ y_l | x\}$ for the corrupted posterior $\tilde{\mathbb{P}}\{\tilde{y}_w \succ \tilde{y}_l | x\}$, ignoring the preference flipping mechanism underlying the corrupted datasets. Other robust losses, such as cDPO (Mitchell 2023), instead similarly follow the principle of maximum likelihood and adopt a loss correction technique, but the ε parameter in cDPO is a hyperparameter irrelevant to each sample (while our loss considers an instance-dependent preference flipping setting). On the basis of cDPO, rDPO (Chowdhury, Kini, and Natarajan 2024) further debias the loss function.

To better understand the difference between these methods, we compare the gradient weights of these methods with that of FA-DPO.

Following the notations in Chowdhury, Kini, and Natarajan (2024), we characterize the gradient of DPO-like methods as

$$\nabla_\theta \mathcal{L} = -\zeta \cdot \beta (\nabla_\theta \log \pi_\theta(\tilde{y}_w | x) - \nabla_\theta \log \pi_\theta(\tilde{y}_l | x)). \quad (8)$$

The weighting coefficient ζ for each method can be related to the base DPO weight coefficient $\zeta_{\text{DPO}} = 1 - p_\theta$, where cDPO applies a small reduction $\zeta_{\text{cDPO}} = \zeta_{\text{DPO}} - \varepsilon$ while rDPO introduces an additive correction $\zeta_{\text{rDPO}} = \zeta_{\text{DPO}} + \frac{\varepsilon}{1-2\varepsilon}$ to adjust the weighting behaviors, respectively. We specify the coefficient $\zeta_{\text{FA-DPO}}$ in FA-DPO as follows. The derivation is deferred in Appendix A.

Lemma 4.2 (Gradient weight coefficient). *For a triplet $\tilde{\mathbf{x}} = (x, \tilde{y}_w, \tilde{y}_l)$, we have the gradient weight coefficient for FA-DPO as*

$$\zeta_{FA-DPO} = \frac{(1 - 2\varepsilon_{\tilde{\mathbf{x}}})p_\theta}{(1 - 2\varepsilon_{\tilde{\mathbf{x}}})p_\theta + \varepsilon_{\tilde{\mathbf{x}}}} \cdot \zeta_{DPO}.$$

The above lemma demonstrates that our weighting scheme constitutes a reparametrization of the DPO gradient weight, distinct from the additive correction approaches employed by cDPO and rDPO. Crucially, when no flipping occurs ($\varepsilon_{\tilde{\mathbf{x}}} = 0$), the gradient reduces exactly to that of the standard DPO. For low flipping probabilities ($\varepsilon_{\tilde{\mathbf{x}}} < 0.5$), the weight increases with the model’s confidence (p_θ), enhancing stability during convergence and improving robustness to noise compared to the fixed correction mechanisms of cDPO and rDPO. At a flipping probability of 0.5, indicating inherent ambiguity in the preference signal, the weight becomes zero, automatically filtering out these low-margin samples that could otherwise impair model performance. Most significantly, when $\varepsilon_{\tilde{\mathbf{x}}} > 0.5$, the gradient direction reverses to optimize $\mathbb{P}\{\tilde{y}_l \succ \tilde{y}_w \mid x\}$ instead of $\mathbb{P}\{\tilde{y}_w \succ \tilde{y}_l \mid x\}$ under conditions of high model uncertainty. This demonstrates self-correction for samples with high detected flipping rates—a capability absent in cDPO and rDPO which apply uniform corrections. Furthermore, the weight approaches zero when high model confidence contradicts high estimated $\varepsilon_{\tilde{\mathbf{x}}}$, yielding robust weights that are jointly determined by both the flipping probability and the policy model’s confidence.

4.3 Preference Flipping Probability Modeling

Preference flipping modeling To overcome the limitations of existing robust DPO approaches that assume either (1) a uniform flipping ratio across all samples or (2) sparse noise patterns within a given dataset, we introduce an instance-dependent preference flipping module that dynamically estimates sample-specific flipping probabilities based on instance features. We model this probability as a logistic regression function of input-dependent features:

$$\varepsilon_{\tilde{\mathbf{x}}} = \sigma(\langle \omega, h(\tilde{\mathbf{x}}) \rangle + \omega_0), \quad (9)$$

where $h : \mathcal{X} \rightarrow \mathbb{R}^d$ is the feature map and $\omega \in \mathbb{R}^d$ are learnable parameters.

Feature map construction We deliberately design the feature map $h(\cdot)$ to incorporate three concepts validated to be effective in language modeling. Notably, the features are supposed to be permutation-equivariant to the response pairs, since the order thereof can be arbitrary in the corrupted sample triplet $\tilde{\mathbf{x}} = (x, \tilde{y}_w, \tilde{y}_l)$.

Response Length. It is noted that longer responses increase cognitive load for human annotator, which may increase the error rate (Chen et al. 2024). We compute both the average lengths and the length difference within a preference sample triplet $(x, \tilde{y}_w, \tilde{y}_l)$ as the first feature,

$$h_{\text{len}}(\tilde{\mathbf{x}}) = \left[\frac{|\tilde{y}_w| + |\tilde{y}_l|}{2}, \left| |\tilde{y}_w| - |\tilde{y}_l| \right| \right]^\top.$$

Here, even we change the order of the response pair to $(x, \tilde{y}_l, \tilde{y}_w)$, the feature map $h_{\text{len}}(\cdot)$ is invariant.

Perplexity (PPL). Perplexity reflects the uncertainty or complexity of a probability distribution. High PPL often correlates with sentences that are hard for humans to comprehend, thereby increasing annotation difficulty (Kong et al. 2024).

$$h_{\text{ppl}}(\tilde{\mathbf{x}}) = \left[\frac{\log(\pi_\theta(\tilde{y}_w|x)\pi_\theta(\tilde{y}_l|x))}{2}, \left| \log \frac{\pi_\theta(\tilde{y}_w|x)}{\pi_\theta(\tilde{y}_l|x)} \right| \right]^\top.$$

Reward Margin. The reward margin implicitly quantifies the model’s confidence in distinguishing preferred from dispreferred responses, which is often used for data selection in preference learning (Wu et al. 2024b; Huang et al. 2025). The corresponding feature map is

$$h_{\text{margin}}(\tilde{\mathbf{x}}) = \left[\frac{\hat{r}_\theta(x, \tilde{y}_w) + \hat{r}_\theta(x, \tilde{y}_l)}{2}, \left| \hat{r}_\theta(x, \tilde{y}_w) - \hat{r}_\theta(x, \tilde{y}_l) \right| \right]^\top$$

where $\hat{r}_\theta(x, y) = \beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)}$ is the implicit reward function induced by DPO.

These features are then concatenated, along with a scalar 1 (for the bias term in ω), and scaled to form the final feature map:

$$h(\tilde{\mathbf{x}}) = [h_{\text{len}}(\tilde{\mathbf{x}}), h_{\text{ppl}}(\tilde{\mathbf{x}}), h_{\text{margin}}(\tilde{\mathbf{x}})]^\top \quad (10)$$

Iterative update To optimize both the preference flipping model and the LLM, we design an iterative update paradigm, which is also adopted in previous work for robust RLHF (Bukharin et al. 2024).

Regarding the concrete design, empirical studies have demonstrated that deep neural networks exhibit a consistent learning trajectory, initially capturing generalizable patterns before eventually overfitting to noisy training instances (Cheng et al. 2024). To adapt to this characteristic, in practice we utilize the learned capability of the LLM during the initial stage of training to optimize the flipping model first (this operation is referred to as *warmup*), and then we iteratively update the two models. The complete training algorithm is presented in Algorithm 1 in Appendix D.

4.4 Theoretical Analysis

This section establishes the theoretical foundations of our approach. Specifically, we demonstrate that the proposed loss function for preference flipping yields desirable statistical properties: consistency under noise and convergence guarantees for the flipping model parameters.

We first present the key result that the minimizer of our FA-DPO loss function, operating on the corrupted (flipped preference) data distribution $\tilde{\mathcal{D}}$, coincides with the minimizer of the original loss function on the underlying clean data distribution \mathcal{D} in RLHF. This *consistency* property guarantees that, asymptotically, our method recovers the same optimal model parameters as would be obtained if trained on clean preference data. We provide the detailed proof in Appendix A.

Theorem 4.3 (Consistency of p_θ). *Given both the corrupted preference data distribution $\tilde{\mathcal{D}}$ induced by the flipping process and the unobserved clean preference data distribution \mathcal{D} , the following equality holds:*

$$\arg \min_{\phi} -\mathbb{E}_{\tilde{\mathbf{x}} \sim \tilde{\mathcal{D}}} [\log \tilde{p}_\phi] = \arg \min_{\phi} -\mathbb{E}_{\mathbf{x} \sim \mathcal{D}} [\log p_\phi],$$

where $\tilde{p}_\phi = (1 - \varepsilon_{\tilde{\mathbf{x}}})p_\phi + \varepsilon_{\tilde{\mathbf{x}}}(1 - p_\phi)$ represents the predicted probability under the flipping model.

This result holds under the specific parameterization where the flipping noise is marginalized via \tilde{p}_ϕ . Furthermore, the core consistency principle holds equivalently under the DPO parameterization: replacing p_ϕ with p_θ and \tilde{p}_ϕ with \tilde{p}_θ , the same result applies to the policy θ optimized via the DPO loss adjusted for flipping.

On the other hand, through a coordinate descent perspective on the iterative updates, the optimization of the flipping model reduces to the convex logistics regression. We present the following convergence result for self-containedness.

We first make the technical assumption that the features $h(\tilde{\mathbf{x}})$ and the parameters ω for the preference flipping model are bounded, i.e., $\|h(\tilde{\mathbf{x}})\| \leq B$, $\|\omega\| \leq B_\omega$. Then, we further assume that the constructed features satisfy the following coverage assumption, which is a common assumption in robust machine learning and the implied feature diversity is critical to prevent collapse or degenerate solutions.

Assumption 4.4 (Feature coverage). *Given the corrupted data distribution $\tilde{\mathcal{D}}$ and the feature map $h : \mathcal{X} \rightarrow \mathbb{R}^d$, the population covariance matrix of features satisfies*

$$\lambda_{\min} \left(\mathbb{E}_{\tilde{\mathbf{x}} \sim \tilde{\mathcal{D}}} [h(\tilde{\mathbf{x}})h(\tilde{\mathbf{x}})^\top] \right) > 0.$$

This assumption implies that the feature vectors are not concentrated in a low-dimensional subspace, which eliminates potential collinearity. The parameters ω are thus identifiable from the data.

Under the boundedness assumption, Assumption 4.4, and the condition that the reward or policy model provides accurate predictions ($p_\phi = p^*$ or $p_\theta = p^*$), we can establish fast convergence for the estimator of the flipping parameters ω obtained via gradient descent:

Theorem 4.5 (Linear Convergence of $\hat{\omega}$). *Given loss function in Equation (7), if $p_\phi = p^*$ or $p_\theta = p^*$ and for the gradient descent update with step size $\eta > 0$:*

$$\omega^{(t+1)} = \omega^{(t)} - \eta \nabla_{\omega} \mathcal{L}_{\text{FA-DPO}}(\omega^{(t)}),$$

the sequence of parameter estimates $\{\omega^{(t)}\}$ converges Q-linearly to the optimal parameter ω^* :

$$\|\omega^{(t+1)} - \omega^*\|^2 \leq (1 - \eta\mu) \|\omega^{(t)} - \omega^*\|^2,$$

for some $\mu > 0$, the convergence holds when $0 < \eta < \frac{2}{L}$, where L is the smoothness constant for $\mathcal{L}_{\text{FA-DPO}}$.

Theorem 4.5 provides the important guarantee that, when initialized with an accurate reward model or policy (reflecting the true clean preference probability p^*), the gradient descent update on the preference flipping model parameters ω converges rapidly at a linear rate. Consequently, this justifies the iterative procedures used in our algorithm.

5 Experiments

We investigate the effectiveness of FA-DPO mainly on aligning the LLMs with preference pairs under the instance-dependent preference flipping setting.

5.1 Experiment Setup

Datasets and models We conduct experiments mainly on two preference datasets: UltraFeedback (Cui et al. 2024) and Anthropic’s HH_Golden (Anthropic 2024). For UltraFeedback, we follow the spirit of Chen et al. (2025) and first filter out samples with low score margins (≤ 0.5) to obtain a cleaner subset for noise injection. This results in 42.4k/61.1k training samples and 1.4k/2k test samples. However, in experiments without manual label flipping, we retain the full training set while using the filtered test set for evaluation. For the backbone LLM models, we first test Pythia-1B on various levels of noise, and then we scale our method on larger models, LLama-3.1-8B and Mistral-7B, to validate the model’s generation capabilities under data contamination. The detailed training setup is listed in Appendix E.

Instance-dependent flipping To validate our designs, we construct datasets with simulated instance-dependent preference flipping based on the following steps:

1. Randomly initialize a preference flipping model $\mathcal{N}_\varepsilon(\vartheta)$ (logistic model) on data features discussed in Section 4.3, and the model outputs a flip probability $\varepsilon_{\tilde{\mathbf{x}}}$ for the current preference pair.
2. To determine whether to flip a preference sample, we set a threshold τ such that if the flipping probability $\varepsilon_{\tilde{\mathbf{x}}}$ computed by the noise model is high enough to exceed this threshold, we would flip the preference. This is to maintain controlled stochasticity during preference flipping. During our experiment, we choose $\tau = 0.8$.
3. To control the total ratio of flips, we train the parameters ϑ on the training set $\mathcal{D}_{\text{train}}$ using the following loss function during noise model initialization, which forces the quantile of the flipping probability in the interval $[\tau, 1]$ to be around a pre-determined flipping ratio η .

$$\mathcal{L}_\varepsilon = \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}_{\text{train}}} \left[\eta - \frac{\mathbb{I}(\mathcal{N}_\varepsilon(x, y_w, y_l; \vartheta) \geq \tau)}{|\mathcal{D}_{\text{train}}|} \right]^2.$$

After training on the contaminated training set, we test the model’s performance on the clean test set which has no noise injected. For experiments with simulated flipping ($\eta > 0$), we intentionally restrict the capability of preference flipping model via only using length-based features, recovering the practical mismatch between estimations and real-world mechanisms. For clean datasets without added flipping ($\eta = 0$), we utilize the full feature set to model the original dataset’s inherent noise patterns.

Evaluation metrics We use *prediction accuracy* (ACC) and *win rate* (WR) as our evaluation metrics.

In particular, we evaluate each model’s *prediction accuracy* on a clean test set $\mathcal{D}_{\text{test}}$ across different noise levels. The accuracy is computed by comparing the predicted rewards for “chosen” versus “rejected” responses:

$$\text{Acc} = \mathbb{E}_{(x, y_w, y_l) \sim \mathcal{D}_{\text{test}}} [\mathbb{I}(\hat{r}_\theta(x, y_w) > \hat{r}_\theta(x, y_l))],$$

Ultrafeedback (61.1k)										
η	0%		10%		20%		30%		40%	
Methods	Acc \uparrow	WR \uparrow	Acc \uparrow	WR \uparrow	Acc \uparrow	WR \uparrow	Acc \uparrow	WR \uparrow	Acc \uparrow	WR \uparrow
DPO	68.22	67.00	61.77	60.10	58.63	56.35	55.43	64.70	51.87	64.35
SIMPO	63.64	62.60	63.73	51.20	57.81	56.50	55.53	54.90	51.77	55.80
ROPO	<u>70.75</u>	66.80	64.93	66.85	58.93	65.55	55.23	64.40	<u>54.07</u>	<u>67.25</u>
cDPO	67.20	66.65	62.57	<u>67.80</u>	59.00	<u>66.05</u>	56.67	<u>66.55</u>	53.93	67.05
rDPO	70.13	65.60	<u>65.90</u>	56.65	<u>61.70</u>	54.85	<u>56.87</u>	<u>57.85</u>	47.67	57.80
FA-DPO	73.05	<u>66.85</u>	67.20	68.45	69.77	66.90	70.97	69.80	70.77	69.80
(#Improv.)	+2.30	-0.15	+1.30	+0.65	+8.07	+0.85	+14.10	+3.25	+16.70	+2.55
HH_Golden (42.5k)										
DPO	98.89	87.00	93.50	34.05	83.53	29.95	73.30	21.05	58.63	26.85
SIMPO	98.27	65.10	93.63	29.95	82.83	24.90	73.33	22.40	60.93	27.50
ROPO	98.80	70.50	93.27	81.15	82.87	54.30	72.33	53.50	<u>62.00</u>	49.20
cDPO	<u>99.24</u>	<u>88.35</u>	93.20	56.95	82.50	<u>62.30</u>	73.30	37.75	60.87	49.05
rDPO	96.80	77.30	<u>96.67</u>	82.80	<u>87.83</u>	53.55	<u>73.73</u>	<u>58.10</u>	56.80	<u>54.10</u>
FA-DPO	99.61	88.70	99.17	88.90	98.10	82.85	99.02	87.70	98.83	78.60
(#Improv.)	+0.37	+0.35	+2.50	+6.10	+10.27	+20.55	+25.29	+29.60	+36.83	+24.50

Table 1: Model performance under different preference flipping ratios (0%-40%) for Pythia-1B on Ultrafeedback and HH_Golden. All values represent percentages, with **bold** indicating the highest score per column and underline denoting the runner-up. The (#Improv.) row quantifies absolute performance gain over the runner-up baseline.

where \mathbb{I} is the indicator function, and \hat{r}_θ is the implicit reward model induced by DPO.

To compute the *win rate*, we employ different evaluators based on model capabilities: DeepSeek-V3 for Pythia-1B, and GPT-4o for both Llama-3.1-8B and Mistral-7B. The calculation is performed as follows:

$$\text{WR} = \frac{\#(\text{Win}) + \#(\text{Tie})/2}{\#(\text{Comparisons})},$$

where $\#(\text{Win})$ and $\#(\text{Tie})$ represent the number of wins and ties compared to the *reference model* (In our case, the SFT model), respectively, and $\#(\text{Comparisons})$ denotes the total number of comparisons between the two models.

For *prediction accuracy*, we compute 5 runs for each method and report the mean value, with each run extracting randomly 1k samples from $\mathcal{D}_{\text{test}}$ for testing. For *win rate*, we extract 1k prompts from $\mathcal{D}_{\text{test}}$, and generate 1k responses from the trained policy model and the corresponding SFT model, respectively. We refer the readers to Appendix E for the detailed prompt template of LLM evaluators.

5.2 Results

Discriminative performance The prediction accuracy, derived from the policy model’s implicit reward signals, reflects the model’s discriminative performance in distinguishing between chosen and rejected responses. From Table 1 and Table 2, we can observe that, as the total flipping ratio increases, all baselines experience sharp prediction accuracy decreases across both datasets. Among the two datasets, HH_Golden exhibits a sharper decline than Ultrafeedback due to its larger chosen-rejected gap, also explaining why all methods achieve much higher accuracy on clean HH_Golden versus Ultrafeedback. Among all baselines, DPO and SIMPO show the highest sensitivity to preference flipping, while the robust methods (ROPO, cDPO,

rDPO) exhibit moderate resilience. FA-DPO demonstrates superior robustness across all flipping ratios, datasets and models, consistently outperforming other approaches.

Generative performance The generation performance shows a similar trend with the discriminative performance. As shown in Table 1, DPO and SIMPO exhibit rapid performance drops with increasing flipping ratios, revealing their inherent fragility to preference noise. This vulnerability is pronounced when chosen and rejected responses come from distinct distributions (like in HH_Golden), reinforcing the overfitting patterns in these methods. FA-DPO proves effective for preserving the generation capabilities of LLMs under instance-dependent preference flipping scenarios, with consistent performance across all conditions, as well as maintaining performance under the clean dataset.

Preference flipping model Then we compare the flipping probabilities predicted by the learned model with the ground-truth flipping distribution. Predicted flipping probabilities show close alignment with actual flipping distribution, as shown in Figure 1 (a). Moreover, Figure 1 (b) reveals a clear separation of the predicted flipping probabilities between flipped and non-flipped samples, indicating distinct decision boundaries between the two kinds of samples. We also plot the relationship between the flipping probabilities with length-based feature in Figure 1 (c) to show that our model captures the relation between noise and data features.

5.3 Ablation Studies

Hyperparameters in iterative training We conduct an ablation study on the hyperparameters in iterative training, first examining the impact of the *warmup* operation, then evaluating different iteration step combinations under a 20% flipping ratio on HH_Golden. Table 3 shows that *warmup*




Figure 1: Characterization of learned preference flipping distribution. (a) Correlation between actual and predicted noise probabilities with regression line; (b) Predicted flipping distributions separated by flipping status; (c) Pattern of predicted flipping distribution with length-based features.

LLama-3.1-8B				
η	20%		40%	
Methods	Acc \uparrow	WR \uparrow	Acc \uparrow	WR \uparrow
DPO	73.89	62.90	64.96	56.10
SIMPO	66.59	62.80	57.07	38.95
ROPO	75.82	62.75	66.59	50.25
cDPO	75.22	64.50	65.70	56.10
rDPO	73.96	59.75	63.99	55.90
FA-DPO	78.80	65.10	78.87	68.50
(#Improv.)	+2.98	+0.60	+12.28	+12.40
Mistral-7B				
DPO	71.35	45.75	62.05	35.35
SIMPO	67.04	54.75	63.77	45.00
ROPO	73.96	54.30	64.36	36.35
cDPO	72.99	47.75	63.24	34.60
rDPO	71.95	44.75	63.10	36.65
FA-DPO	78.05	61.35	78.49	59.00
(#Improv.)	+4.09	+6.60	+14.13	+14.00

Table 2: Model performance under flipping ratios (20% and 40%) for LLama-3.1-8B and Mistral-7B on Ultrafeedback.

Warmup	Iteration Batches		Metrics	
	Noise	Policy	Acc \uparrow	WR \uparrow
No	20	20	84.96	57.35
	20	50	77.00	53.20
	50	50	83.16	<u>76.85</u>
Yes	20	20	98.56	70.04
	20	50	98.40	82.90
	50	50	96.80	82.90

Table 3: Ablation studies on warmup and policy/noise iteration steps.

plays a vital role in performance improvement. We remark that the preference flipping model needs to be sufficiently trained to guide the policy learning, especially when the initial policy model is not well trained. When the flipping model performs well, training more policy steps yields improved performance.

Computational cost Although FA-DPO requires training an auxiliary flipping model, we note the overall computational cost remains comparable to—or even lower than—that of standard DPO. This is because we limit training to a single epoch, keeping the total data usage consistent across all methods. Additionally, the PPL and reward margin features required for training the preference flipping model are derived directly from the policy’s forward pass log-likelihoods, introducing no extra computation overhead.

6 Conclusion

In this paper, we tackle the challenging scenario of instance-dependent noisy human feedback through introducing a framework that simultaneously models preference flipping and post-trains the LLM, on the basis of the RLHF and the DPO frameworks. Instead of directly modeling the noise inside the BT model, our approach separates human intention from the noising process by assuming a post-transition after the BT model forming the preference with encoded probabilities. This process features the stochastic transformation from groundtruth to noisy labels, providing a more realistic representation. By integrating the MLE-based BT model with preference flipping probabilities, we can then learn a statistically consistent estimator. In more detail, our algorithm iteratively updates the noise model and fine-tunes the LLM parameters, and our implementation on DPO is achieved with little additional resource consumption. Empirically, our approach adopts relevant sequence features to model flipping ratios and yields high probability for preference flipping, as expected. In evaluations on instance-dependent noisy human preference datasets, our algorithm demonstrates higher predictive accuracy compared to vanilla DPO and other baseline methods.

Acknowledgments

This work was initiated during the first author’s internship at Shanghai Artificial Intelligence Laboratory. We sincerely thank the anonymous reviewers for their valuable feedback and constructive suggestions, which helped improve this work. The authors acknowledge funding from National Natural Science Foundation of China (NSFC) under grant 62502408, Research Grants Council (RGC) under grant 22303424, GuangDong Basic and Applied Basic Research Foundation under grant 2025A1515010259, and Guangdong and Hong Kong Universities “1+1+1” Joint Research Collaboration Scheme under grant 2025A050500.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. Anthropic’s HH_Golden.
- Azar, M. G.; Guo, Z. D.; Piot, B.; Munos, R.; Rowland, M.; Valko, M.; and Calandriello, D. 2024. A general theoretical paradigm to understand learning from human preferences. In *International Conference on Artificial Intelligence and Statistics*, 4447–4455. PMLR.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4): 324–345.
- Bukharin, A.; Hong, I.; Jiang, H.; Li, Z.; Zhang, Q.; Zhang, Z.; and Zhao, T. 2024. Robust reinforcement learning from corrupted human feedback. *Advances in Neural Information Processing Systems*, 37: 124093–124113.
- Casper, S.; Davies, X.; Shi, C.; Gilbert, T. K.; Scheurer, J.; Rando, J.; Freedman, R.; Korbak, T.; Lindner, D.; Freire, P.; et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Chen, G. H.; Chen, S.; Liu, Z.; Jiang, F.; and Wang, B. 2024. Humans or llms as the judge? a study on judgement biases. *arXiv preprint arXiv:2402.10669*.
- Chen, P.; Chen, X.; Yin, W.; and Lin, T. 2025. ComPO: Preference alignment via comparison oracles. *arXiv preprint arXiv:2505.05465*.
- Cheng, J.; Xiong, G.; Dai, X.; Miao, Q.; Lv, Y.; and Wang, F.-Y. 2024. RIME: Robust Preference-based Reinforcement Learning with Noisy Preferences. *arXiv preprint arXiv:2402.17257*.
- Chowdhury, S. R.; Kini, A.; and Natarajan, N. 2024. Provably robust dpo: Aligning language models with noisy feedback. *arXiv preprint arXiv:2403.00409*.
- Cui, G.; Yuan, L.; Ding, N.; Yao, G.; He, B.; Zhu, W.; Ni, Y.; Xie, G.; Xie, R.; Lin, Y.; et al. 2024. ULTRAFEDBACK: Boosting Language Models with Scaled AI Feedback. In *Forty-first International Conference on Machine Learning*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Ethayarajh, K.; Xu, W.; Muennighoff, N.; Jurafsky, D.; and Kiela, D. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.
- Gao, Y.; Alon, D.; and Metzler, D. 2024. Impact of preference noise on the alignment performance of generative language models. *arXiv preprint arXiv:2404.09824*.
- Huang, K.; Wu, J.; Chen, Z.; Wang, X.; Gao, J.; Ding, B.; Wu, J.; He, X.; and Wang, X. 2025. Larger or Smaller Reward Margins to Select Preferences for LLM Alignment? In *Forty-second International Conference on Machine Learning*.
- Kong, K.; Xu, X.; Wang, D.; Zhang, J.; and Kankanhalli, M. S. 2024. Perplexity-aware correction for robust alignment with noisy preferences. *Advances in Neural Information Processing Systems*, 37: 28296–28321.
- Liang, X.; Chen, C.; Wang, J.; Wu, Y.; Fu, Z.; Shi, Z.; Wu, F.; and Ye, J. 2024. Robust preference optimization with provable noise tolerance for llms. *arXiv preprint arXiv:2404.04102*.
- Liu, Y.; Cheng, H.; and Zhang, K. 2023. Identifiability of label noise transition matrix. In *International Conference on Machine Learning*, 21475–21496. PMLR.
- Meng, Y.; Xia, M.; and Chen, D. 2024. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*.
- Mitchell, E. 2023. A note on DPO with noisy preferences & relationship to IPO.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Peng, X. B.; Kumar, A.; Zhang, G.; and Levine, S. 2019. Advantage-weighted regression: Simple and scalable off-policy reinforcement learning. *arXiv preprint arXiv:1910.00177*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Song, H.; Kim, M.; Park, D.; Shin, Y.; and Lee, J.-G. 2022. Learning from noisy labels with deep neural networks: A survey. *IEEE transactions on neural networks and learning systems*, 34(11): 8135–8153.
- Wu, J.; Xie, Y.; Yang, Z.; Wu, J.; Chen, J.; Gao, J.; Ding, B.; Wang, X.; and He, X. 2024a. Towards Robust Alignment

of Language Models: Distributionally Robustifying Direct Preference Optimization. *arXiv preprint arXiv:2407.07880*.

Wu, J.; Xie, Y.; Yang, Z.; Wu, J.; Gao, J.; Ding, B.; Wang, X.; and He, X. 2024b. *beta*-DPO: Direct Preference Optimization with Dynamic *beta*. *Advances in Neural Information Processing Systems*, 37: 129944–129966.

Wu, Z.; Hu, Y.; Shi, W.; Dziri, N.; Suhr, A.; Ammanabrolu, P.; Smith, N. A.; Ostendorf, M.; and Hajishirzi, H. 2024c. Fine-grained human feedback gives better rewards for language model training. *Advances in Neural Information Processing Systems*, 36.

Xia, X.; Liu, T.; Han, B.; Gong, M.; Yu, J.; Niu, G.; and Sugiyama, M. 2021. Sample selection with uncertainty of losses for learning with noisy labels. *arXiv preprint arXiv:2106.00445*.

Xia, X.; Liu, T.; Han, B.; Wang, N.; Gong, M.; Liu, H.; Niu, G.; Tao, D.; and Sugiyama, M. 2020. Part-dependent label noise: Towards instance-dependent label noise. *Advances in Neural Information Processing Systems*, 33: 7597–7610.

Yuan, Z.; Yuan, H.; Tan, C.; Wang, W.; Huang, S.; and Huang, F. 2023. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*.

Zhao, Y.; Joshi, R.; Liu, T.; Khalman, M.; Saleh, M.; and Liu, P. J. 2023. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*.

Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36: 46595–46623.

Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.