

Preference Robustness for DPO with Applications to Public Health

Cheol Woo Kim*, Shresth Verma*, Mauricio Tec*, Milind Tambe

Harvard University

Abstract

We study an LLM fine-tuning task for designing reward functions for sequential resource allocation problems in public health, guided by human preferences expressed in natural language. This setting presents a challenging testbed for alignment due to complex and ambiguous objectives and limited data availability. We propose DPO-PRO, a robust fine-tuning algorithm based on Direct Preference Optimization (DPO), which accounts for uncertainty in the preference distribution using a lightweight Distributionally Robust Optimization (DRO) formulation. Unlike prior DRO-based variants, DPO-PRO focuses solely on uncertainty in preferences, avoiding unnecessary conservatism and incurring negligible computational overhead. We evaluate DPO-PRO on a real-world maternal mobile health program operated by the non-profit organization ARMMAN, as well as on standard alignment benchmarks. Experimental results demonstrate that our method consistently improves robustness to noisy preference signals compared to existing DPO variants. Moreover, DPO-PRO achieves comparable performance to prior self-reflection-based baseline for reward function design, while requiring significantly lower inference-time cost.

1 Introduction

Limited resource allocation is a core challenge in public health, where decision-makers must prioritize individuals for care or intervention under tight financial and operational constraints. In maternal healthcare, for instance, preventative care awareness programs have been shown to significantly reduce maternal mortality (Helpmum 2025; Jahan et al. 2018; ARMMAN 2025), yet are often implemented by non-profit organizations that must serve large populations with limited resources. To support such efforts, it is crucial to develop allocation policies that make effective sequential decisions based on individual-level attributes such as age, income, or health status.

Reinforcement learning (RL) offers a powerful framework for learning allocation policies in complex, dynamic environments. A widely used model in this context is the Restless Multi-Armed Bandit (RMAB) (Niño-Mora 2023), a class of Markov Decision Processes (MDPs) well-suited

for sequential decision-making under resource constraints. However, in RL, it is well-recognized that designing an appropriate reward function is critical (Singh, Lewis, and Barto 2009; Booth et al. 2023). In public health applications, decision-makers’ priorities (e.g., which populations to prioritize) must be encoded into the reward function, so that resulting policy outcomes align with long-term public health goals. Crafting a reward function that accurately captures these priorities is often labor-intensive, demanding manual tuning, repeated simulation, and domain expertise. As public health goals shift over time, this design process must be repeated, limiting the scalability and adaptability of RL-based solutions in dynamic, real-world environments.

Recently, several works explored using large language models (LLMs) to automate reward function design for sequential decision-making problems (Mirchandani, Karamcheti, and Sadigh 2021; Goyal, Niekum, and Mooney 2019; Carta et al. 2022; Ma et al. 2024; Behari et al. 2024; Verma et al. 2025). In these works, an LLM interprets a user’s high-level objective (e.g., desired policy outcomes) in natural language and outputs a reward function that guides RL policy training. This enables efficient workflow for developing human-aligned RL policies. Recent studies demonstrated that LLMs can effectively generate reward functions for RMAB problems in public health domains (Behari et al. 2024; Verma et al. 2025).

Existing approaches often rely on inference-time optimization via self-reflection (Shinn et al. 2023), where the LLM iteratively refines its outputs based on simulated feedback. Although this strategy provides flexibility, it incurs substantial inference-time costs due to repeated rounds of simulation and reasoning. In many non-profit settings, data sensitivity constraints further limit the use of large proprietary models, pushing organizations toward open-source alternatives, which are often smaller and less capable. Because self-reflection is highly dependent on model strength, its effectiveness can be diminished in these environments. As a result, fine-tuning open-source models offers a more practical path: it directly improves their capabilities while avoiding the high inference costs and privacy issues in large-scale public health deployments (Wu et al. 2024). This makes fine-tuning particularly suitable for resource-constrained domains such as public health (Zhao et al. 2025).

Thus, we aim to move beyond inference-time methods

*These authors contributed equally.

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

and instead fine-tune an LLM to more effectively translate natural language preferences into high-performing reward functions for RMABs. A promising approach is Direct Preference Optimization (DPO) (Rafailov et al. 2023), which enables LLM fine-tuning from pairwise preference annotations. Recent work has also leveraged LLM-as-a-judge to scale annotation as a proxy for human feedback (Zhang et al. 2025; Tan et al. 2024; Cui et al. 2024; Zhu et al. 2024). However, applying these techniques to the public health domain introduces significant challenges, particularly due to the presence of noise in the preference signal.

First, the objectives in the public health domain are often ambiguous or subjective (e.g., “prioritize older individuals”). *Second*, reward design in RMABs demands complex reasoning, as the quality of a proposed reward function is not evaluated in isolation, but rather by the policy it induces, increasing the complexity of assessing preference. *Third*, real-world public health datasets are typically smaller, increasing the risk of misalignment due to overfitting and reward hacking, which is a common problem with reinforcement learning for human feedback methods such as DPO (Gheshlaghi Azar et al. 2024; Amini, Vieira, and Cotterell 2024; Yang et al. 2024; Xiao et al. 2025). However, public health is a high-stakes domain, where poorly aligned policies can lead to serious real-world consequences.

To tackle these challenges, we introduce DPO-PRO, a lightweight and distributionally robust enhancement of DPO. DPO-PRO incorporates uncertainty in the preference distribution using an efficient yet effective Distributionally Robust Optimization (DRO) formulation (Rahimian and Mehrotra 2022), ensuring that the fine-tuned policy remains reliable under imperfect preference annotation. Unlike prior DRO-based DPO methods (Wu et al. 2025; Xu et al. 2025; Mandal, Sasnauskas, and Radanovic 2025), our approach avoids excessive conservatism and only adds negligible computational overhead. We further show that the resulting robust loss can be interpreted as a regularized DPO loss that penalizes model overconfidence and weak preference signals.

We evaluate DPO-PRO on the UltraFeedback benchmark (Cui et al. 2024) and on a simulated maternal health task based on ARMMAN’s real-world mobile health program (ARMMAN 2025). In this setting, health workers can place only a limited number of live service calls each week to boost beneficiary engagement. Prior work models this problem using RMABs, where each beneficiary’s engagement evolves over time and the objective is to allocate a limited number of calls to maximize long-term engagement. Our experiments show that DPO-PRO is more robust to noisy preference annotation than vanilla DPO and prior DRO-based approaches, both on UltraFeedback and across a range of public health objectives. We further show that DPO-PRO achieves performance comparable to the Decision Language Model (DLM), a self-reflection-based method used in the ARMMAN setting (Behari et al. 2024), while requiring significantly lower inference-time cost (See Figure 1 in Section 5.2).

2 Related Works

A growing body of work has explored robustness in RLHF and DPO, particularly by leveraging DRO. Wu et al. (2025) introduce a DRO-based DPO formulation using a KL-divergence-based ambiguity set. Xu et al. (2025) extend this approach to both KL and Wasserstein distances, while Mandal, Sasnauskas, and Radanovic (2025) explore robustness under total variation distance, applying it to both DPO and standard two-stage RLHF.

In contrast to these DRO-based approaches, other works assume explicit corruption models. For example, Chowdhury, Kini, and Natarajan (2024) assume that preference labels are flipped according to a known noise rate, and Bukharin et al. (2024) model preferences using a Bradley–Terry (BT) framework (Bradley and Terry 1952) with added noise on the reward differences.

Hong et al. (2024) use a DRO framework to develop adaptive loss functions for reward model learning, but focus is not on distributional robustness. Zhan et al. (2024) introduce robustness by constructing confidence sets over learned reward functions, followed by pessimistic policy optimization.

Broadly, existing robustness techniques fall into two categories: (1) methods relying on strong assumptions about noise models, which may be unrealistic in practice, and (2) DRO-based methods that avoid explicit corruption modeling but hedge against worst-case distributions. As discussed in Section 3.2, standard DRO approaches tend to be overly conservative and often require solving challenging min–max optimization problems. For tractability, prior methods rely on heuristic approximations (Wu et al. 2025; Xu et al. 2025), potentially undermining theoretical guarantees or practical robustness. In contrast, we provide efficient approach DPO-PRO, with theoretical guarantees, and simultaneously illustrate its applicability in real world public health settings.

3 Preliminaries

3.1 Data Distribution

In DPO, each data point is a tuple (x, y_1, y_2, c) . A prompt $x \in \mathcal{X}$ is drawn from a distribution μ . Given x , responses $y_1, y_2 \in \mathcal{Y}$ are sampled independently from a policy $\pi(\cdot|x)$. A label $c \in \{1, -1\}$ is a Bernoulli random variable indicating whether response y_1 is preferred over y_2 , with preference probability given by $p^*(y_1 \succ y_2|x)$. The distribution p^* represents the ground-truth preference distribution that the user intends the language model to align with.

In summary, the data-generating process involves three components: (1) the prompt distribution, (2) the response distribution, and (3) the preference distribution. This yields the following joint distribution: $P(x, y_1, y_2, c) = \mu(x) \pi(y_1|x) \pi(y_2|x) [\mathbb{I}_{\{c=1\}} p^*(y_1 \succ y_2 | x) + \mathbb{I}_{\{c=-1\}} p^*(y_2 \succ y_1 | x)]$. We often omit the dependence of p^* (and other preference distributions) on (x, y_1, y_2) and simply write p^* instead of $p^*(y_1 \succ y_2|x)$ when the meaning is clear from context.

3.2 (Distributionally Robust) DPO

In standard RLHF, the first step is to learn a reward function R_ϕ from a fixed dataset \mathcal{D} . This is typically done by

minimizing the negative log-likelihood:

$$-\mathbb{E}_{(x, y_1, y_2, c) \sim \mathcal{D}} [\log \sigma(c \cdot (R_\phi(x, y_1) - R_\phi(x, y_2)))] ,$$

where σ is the sigmoid function.

Once the reward model is trained, a policy π_θ is optimized to maximize the expected reward while remaining close to a reference policy π_{ref} , typically via Proximal Policy Optimization (PPO) (Schulman et al. 2017):

$$\max_{\pi_\theta} \mathbb{E}_{(x, y) \sim \pi_\theta} [R_\phi(x, y)] - \beta \text{KL}(\pi_\theta \parallel \pi_{ref}),$$

where β is a regularization parameter controlling the strength of the KL penalty.

DPO simplifies this pipeline by combining the two stages into a single objective that directly optimizes the policy from pairwise preference data. The DPO loss $\mathcal{L}_{\text{DPO}}(\pi_\theta)$ is given by:

$$-\mathbb{E}_{(x, y_1, y_2, c) \sim \mathcal{D}} [\log \sigma(c \beta (\log \frac{\pi_\theta(y_1|x)}{\pi_{ref}(y_1|x)} - \log \frac{\pi_\theta(y_2|x)}{\pi_{ref}(y_2|x)})))].$$

For brevity, we define the per-sample loss as:

$$\ell_\theta(x, y_1, y_2, c) = -\log \sigma(c \beta (\log \frac{\pi_\theta(y_1|x)}{\pi_{ref}(y_1|x)} - \log \frac{\pi_\theta(y_2|x)}{\pi_{ref}(y_2|x)})) .$$

Using this expression, $\mathcal{L}_{\text{DPO}}(\pi_\theta)$ is also equivalent to

$$\mathbb{E}_{(x, y_1, y_2) \sim \mathcal{D}} [p^* \ell_\theta(x, y_1, y_2, 1) + (1 - p^*) \ell_\theta(x, y_1, y_2, -1)].$$

When the context is clear that the triplet (x, y_1, y_2) and the parameter θ is fixed, we also use ℓ_1 and ℓ_{-1} to denote $\ell_\theta(x, y_1, y_2, 1)$ and $\ell_\theta(x, y_1, y_2, -1)$, respectively.

As discussed earlier, DPO is vulnerable to overfitting and noise in the data. To address this, recent work has proposed distributionally robust objectives of the form $\max_{P \in \mathcal{Q}(\mathcal{D})} \mathcal{L}_P(\pi_\theta)$, where $\mathcal{Q}(\mathcal{D})$ denotes an ambiguity set centered around the empirical distribution (data) \mathcal{D} and the loss \mathcal{L}_P is the DPO loss evaluated under the perturbed distribution P rather than the original data distribution (Xu et al. 2025; Wu et al. 2025).

Conservatism and Computational Costs The above DRO formulation hedges against shifts in the entire joint distribution on (x, y_1, y_2, c) . This allows the adversary to assign weights to highly unlikely prompts or responses. The outer minimization must then optimize against losses in these practically irrelevant regions, which might make the update overly conservative and slow improvement on the data the model actually sees. From a computational standpoint, solving the resulting min-max problem can be intractable in practice, and existing methods often rely on approximations that deviate from the original DRO formulation.

3.3 Reward Function Design Tasks for RMABs

In our setting, the prompt x represents a desired policy outcome specified by a human prompt, for example, Focus on the young mothers by age and also focus on those with low income. The responses y_1 and y_2 correspond to candidate reward functions for RMABs. The role of the LLM-annotator is to judge which of the two reward functions, y_1 or y_2 , better aligns with the objective expressed in the prompt.

A key challenge in this task is that the quality of a reward function cannot be assessed in isolation. Rather, it must be evaluated based on the policy it induces when optimized, i.e., how effectively the resulting policy fulfills the user’s (often ambiguous) intent. This requires the annotator to infer the downstream effects of each reward function. Given the complexity and ambiguity involved, the resulting preference signals are inevitably noisy and should not be treated as fully reliable.

4 DPO with Preference Robustness

In this section, we introduce DPO-PRO, a distributionally robust version of DPO that specifically targets distributional shifts or other forms of noise in the preference distribution. The complete proofs for all theorems are provided in the extended version of this paper (Kim et al. 2025), which includes the full technical appendix.

4.1 Uncertainties in the Preference Distribution

In each data point (x, y_1, y_2, c) , the triplet (x, y_1, y_2) represents observed content, which typically comes from a well-understood and controllable data collection process. In general, x is drawn from a pool of candidate inputs, and the response pair (y_1, y_2) is generated by the reference policy π_{ref} (usually a supervised-fine tuned model) for every sampled prompt x . As more data is collected, the uncertainty in this part of the data is expected to diminish. Noise in the prompts and responses becomes even less of an issue under iterative DPO methods, which are increasingly adopted in practice (Xiong et al. 2024; Cen et al. 2025; Xie et al. 2025).

In contrast, we argue that the preference distribution $c \sim p^*$ is inherently noisy in practice, and collecting more data is unlikely to resolve the underlying uncertainty in human preference. Even when humans are directly annotating, various sources of noise remain, including human subjectivity, irrational or inconsistent behavior, and temporal variability (e.g., the same annotator providing different judgments at different times). This noise is persistent and cannot be easily addressed through better data curation or increased data volume. In other words, the ground-truth preference distribution p^* is an idealized and inaccessible object in practice. As mentioned in the introduction, this challenge is particularly pronounced in public health domains, where objectives are often ambiguous or subjective (e.g., “prioritize older individuals”), and reward design demands complex reasoning about long-term policy consequences that are difficult to assess accurately. Moreover, real-world public health datasets are typically smaller, further exacerbating preference uncertainty and increasing the risk of misalignment. Therefore, we argue that uncertainty in the preference distribution should be the primary motivation for incorporating robustness into DPO.

We assume that for each prompt and response pair (x, y_1, y_2) , we have access to a (potentially noisy) preference distribution $q(y_1 \succ y_2|x)$. We apply DRO adjustment, assuming worst-case deviation from q within a pre-specified chi-squared divergence ball.

Formally, we define the DRO loss $\mathcal{L}_{\text{DRO}}(\pi_\theta)$ as:

$$\mathbb{E}_{(x, y_1, y_2) \sim \mathcal{D}} \max_{p \in Q(x, y_1, y_2, \rho)} \mathbb{E}_{c \sim p} [\ell_\theta(x, y_1, y_2, c)] \quad (1)$$

where $Q(x, y_1, y_2, \rho) = \{p : \chi^2(p \| q(y_1 \succ y_2 | x)) \leq \rho\}$. Our formulation avoids unnecessary robustification over the joint distribution on (x, y_1, y_2) , and assumes that this part of the data is reliable. Furthermore, unlike standard DRO methods that can add significant computational burden (Rahimian and Mehrotra 2022), our formulation introduces negligible additional cost as shown next.

4.2 Computing Worst-case Distribution

Consider a single data point (x, y_1, y_2) . We assume access to an estimate of the probability that one response is preferred over the other: $q(y_1 \succ y_2 | x) \in (0, 1)$.

Given this probability q , we define the following per-sample optimization problem:

$$\begin{aligned} \max_{p \in [0, 1]} \quad & p \ell_\theta(x, y_1, y_2, 1) + (1 - p) \ell_\theta(x, y_1, y_2, -1) \\ \text{s.t.} \quad & \frac{(p - q)^2}{q(1 - q)} \leq \rho \end{aligned} \quad (2)$$

This problem finds the worst-case distribution within the chi-squared ambiguity set centered at q , $\{p : \chi^2(p \| q(y_1 \succ y_2 | x)) \equiv \frac{(p - q)^2}{q(1 - q)} \leq \rho\}$. Let $\hat{p}(y_1 \succ y_2 | x)$ denote the optimal solution to this problem. Using this, the robust loss $\mathcal{L}_{\text{DRO}}(\pi_\theta)$ can be equivalently expressed as

$$\mathbb{E}_{(x, y_1, y_2) \sim \mathcal{D}} [\hat{p} \ell_\theta(x, y_1, y_2, 1) + (1 - \hat{p}) \ell_\theta(x, y_1, y_2, -1)].$$

Eq (2) is a one-dimensional optimization problem with a linear objective, which admits a simple closed-form solution:

$$\hat{p}(y_1 \succ y_2 | x) = \begin{cases} \min \left\{ 1, q + \sqrt{\rho q(1 - q)} \right\}, & \text{if } \ell_1 \geq \ell_{-1}, \\ \max \left\{ 0, q - \sqrt{\rho q(1 - q)} \right\}, & \text{if } \ell_1 < \ell_{-1}. \end{cases}$$

Consider the case where $\ell_1 \geq \ell_{-1}$, which implies that under the current policy θ , the model assigns a higher likelihood to y_2 being preferred over y_1 . To increase the loss, the adversary seeks to shift the preference probability $q(y_1 \succ y_2 | x)$ in the opposite direction. That is, it tries to increase q to emphasize a preference for y_1 , which contradicts the model's belief. Consequently, the worst-case distribution becomes $q + \sqrt{\rho q(1 - q)}$. Since probabilities must remain within the unit interval, we clip the value at 1, yielding the final expression $\min \left\{ 1, q + \sqrt{\rho q(1 - q)} \right\}$.

4.3 Efficient Optimization of the Loss

The worst-case probability obtained from the optimization problem above, \hat{p} , replaces the preference probability in the per-sample DPO gradient as well:

$$\hat{p} \nabla_\theta \ell_\theta(x, y_1, y_2, 1) + (1 - \hat{p}) \nabla_\theta \ell_\theta(x, y_1, y_2, -1). \quad (3)$$

Proposition 4.1. *Eq (3) provides an unbiased gradient estimate of the DRO loss in Eq (1).*

Note that we do not formally differentiate through the inner maximization in Eq (1) with respect to θ . However, Danskin's theorem (Bertsekas 1999) justifies our approach. For a fixed θ , we may solve the inner maximization and directly substitute the resulting worst-case \hat{p} into the gradient expression.

We emphasize the resulting DRO gradient is both exact and computationally efficient. Unlike prior work, our approach does not introduce any approximation or heuristic in minimizing the DRO objective. Moreover, our distributional robustness is applied only to the preference distribution q , rather than the full data-generating distribution on (x, y_1, y_2, c) . Hence, the resulting method is significantly less conservative than earlier DRO-based DPO approaches, and matching more accurately the type of distribution shift we expect in the true human preferences in our public health application as seen later.

4.4 Analysis of the Loss Function

In this section, we analyze the robust loss \mathcal{L}_{DRO} and show that it is equivalent to the original DPO loss augmented with a regularization term. This interpretation provides insight into how the DRO formulation affects model learning.

Proposition 4.2. *The DRO loss \mathcal{L}_{DRO} is equivalent to regularizing the original DPO loss \mathcal{L}_{DPO} as:*

$$\mathcal{L}_{\text{DRO}} = \begin{cases} \mathcal{L}_{\text{DPO}} + \min\{1 - q, \sqrt{\rho q(1 - q)}\}(\ell_1 - \ell_{-1}), & \text{if } \ell_1 \geq \ell_{-1}, \\ \mathcal{L}_{\text{DPO}} + \min\{q, \sqrt{\rho q(1 - q)}\}(\ell_{-1} - \ell_1), & \text{if } \ell_1 < \ell_{-1}. \end{cases}$$

We now analyze the behavior of the regularization term, assuming that $\ell_1 \geq \ell_{-1}$. The opposite case follows symmetric logic.

(1) Uncertainty-Weighted Coefficient. The first factor $\min \left\{ 1 - q, \sqrt{\rho q(1 - q)} \right\}$ penalizes cases where the preference signal is uncertain (i.e., when q is close to 0.5). This term is largest when q is near 0.5 and diminishes as $q \rightarrow 0$ or $q \rightarrow 1$. Its magnitude is modulated by the robustness radius ρ , with larger ρ leading to a stronger penalty.

(2) Model Confidence. The second factor $(\ell_1 - \ell_{-1})$ reflects the log-odds of the current model's preference for y_2 over y_1 . A large value indicates strong model confidence. This term increases when the model becomes more certain about the preference, regardless of the ground-truth label.

(3) Combined Effect. The regularization term penalizes the model for being overly confident in its preferences when the preference signal q is ambiguous. In contrast, when the signal is strong (i.e., q is close to 0 or 1), the penalty diminishes, allowing the model to express stronger preferences. In effect, the DRO loss encourages calibrated learning: the model is allowed to be confident only when the preference signal is also confident.

4.5 Obtaining the Soft Score q in Practice

The DRO adjustment described above requires a soft score $q(y_1 \succ y_2|x)$, rather than a binary annotation for each prompt and response pair (x, y_1, y_2) . Soft scores can be obtained from LLM judges in several ways: (1) producing multiple binary preference judgments through repeated queries and averaging the outcomes; (2) directly outputting a numerical preference score via appropriate prompting; or (3) extracting the log-probabilities assigned to tokens representing the preferences (e.g., 1 for $y_1 \succ y_2$ and -1 for $y_2 \succ y_1$), and applying a softmax transformation to estimate a preference distribution (Lee et al. 2024). When using a off-the-shelf reward model, the scalar outputs of the reward model can be converted into a pairwise preference probability using the BT model. We provide further discussion on annotation formats and the connection to robustness in (Kim et al. 2025).

5 Experiments

We validate DPO-PRO in two different settings. First, we validate the methodology on a standardized benchmark for preference-based LLM finetuning. Next, we adapt this experiment for the ARMMAN maternal health dataset (ARMMAN 2025) in the following section. The complete experimental details are provided in the extended version of this paper (Kim et al. 2025).

5.1 Benchmarks on General Alignment Datasets

We use the recent UltraFeedback alignment benchmark dataset (Cui et al. 2024), consisting of 60,000 top-quality human pairwise preference annotations that cover a wide variety of tasks and instruction types, rendering it a broad benchmark for examining alignment research.

Noise in the training data Our primary objective is to evaluate how effectively preference-based LLM fine-tuning methods handle shifts in preference distributions between training and evaluation data. We use q^* denote the true preference distribution used at test time. For our experiments, we define q^* using the state-of-the-art preference reward model Eurus 7B (Yuan et al. 2025), which has been trained for the UltraFeedback dataset (Cui et al. 2024). We compute q^* using the BT model.

To simulate noise during training, we construct a noisy preference distribution q_α for training as the mixture

$$q_\alpha := q^*(1 - \alpha) + (1 - q^*)\alpha. \quad (4)$$

Intuitively, q_α “flips” the preference distribution with probability α . We refer to this type of noise interchangeably as *label switching*. During training, we have access to q_α but not to q^* , with α being unknown.

Training details We base our experiments on the recent Phi 3-mini 3B model without instruction (Abdin et al. 2024). We initially train for two epochs of supervised fine tuning (SFT), followed by one epoch of DPO, with the baselines discussed below. All experiments are performed on a Nvidia A100 GPU 40 GB with batch size 2.

Baselines We compare DPO-PRO against the standard DPO (Rafailov et al. 2023) and the prior distributionally robust variant DrDPO (Wu et al. 2025) using their best hyperparameters suggested by the authors. We evaluate DPO-PRO under several values of the robustness parameter ρ based on the χ^2 -divergence. Models are trained under two levels of label-flip noise, $\alpha = 0.3$ (low) and $\alpha = 0.6$ (high), as well as the noiseless setting $\alpha = 0$.

Evaluation After the fine-tuning step, we generate a response $y_g \sim \pi_\theta(x)$ for every prompt x in the evaluation dataset and compute its reward $R_\phi(x, y_g)$. We consider two evaluation metrics. The first is the *win rate*, computed by comparing the generated response with the chosen response y_c in the dataset using an LLM-as-a-judge approach with GPT-4o-mini (OpenAI 2024). To mitigate position bias, we randomly flip the order of the two responses presented to the judge. The second metric, *evaluation reward*, is computed as the average reward score defined as $\frac{1}{N_{\text{eval}}} \sum_{i=1}^{N_{\text{eval}}} R_\phi(x, y_g)$.

Results Table 1 summarizes the results. For the LLM-as-a-judge evaluation, the different DPO-PRO variants (each corresponding to different ρ value) tend to underperform the baselines in the low-noise setting, but their relative performance improves as noise increases. Under high noise, all DPO-PRO variants outperform both vanilla DPO and DrDPO. The effect of ρ follows the expected pattern: smaller ρ values perform well when noise is low but degrade more quickly as noise increases, while larger ρ values exhibit more stability. Overall, the DPO-PRO family shows improved robustness as the preference noise becomes larger.

For metrics based on the reward model, DPO-PRO outperforms the baselines even in the noiseless setting, except for the case with $\rho = 0.1$. As expected, performance declines for all methods as noise increases. In the low-noise regime, however, the DPO-PRO variants perform slightly worse than the baseline, but they regain and ultimately surpass baseline performance as the noise becomes even larger. Note that this does not contradict the DRO objective, because the true optimization target for fine-tuning is reward plus KL regularization, while our evaluation here relies solely on reward-model scores. Moreover, these scores are only a proxy for quality, particularly for out-of-distribution responses, and are limited by the reward model’s own generalization capabilities.

5.2 Benchmarks on Maternal Health Data

We apply our proposed method in the context of mobile health program run by ARMMAN (2025). We use real-world anonymized beneficiary listenership data from a quality-improvement study conducted by ARMMAN¹. The data consists of registration information of beneficiaries (mothers) containing their socio-demographic features of age, income education as well as their preferences on call slot times and language. Additionally, we have information on beneficiaries’ interaction with automated voice calls and live service calls made by health-workers. We model the

¹The authors have been access to this restricted dataset under a data usage agreement.

(a) Win rate (%) based on LLM-as-a-judge evaluation (\uparrow).

Method	No noise	Low noise	High noise
DPO-PRO ($\rho = 0.008$)	43.24	35.17	28.75
DPO-PRO ($\rho = 0.03$)	41.71	33.33	28.98
DPO-PRO ($\rho = 0.1$)	40.33	33.54	27.97
DPO	43.77	36.19	26.97
DrDPO	44.85	38.84	25.51

(b) Evaluation reward (\uparrow).

Method	No noise	Low noise	High noise
DPO-PRO ($\rho = 0.008$)	923.70	713.00	446.65
DPO-PRO ($\rho = 0.03$)	912.65	695.00	499.59
DPO-PRO ($\rho = 0.1$)	781.28	656.32	499.36
DPO	877.01	729.12	400.44
DrDPO	897.67	753.77	419.05

Table 1: Results on UltraFeedback under label-flip noise.

problem of allocating live service calls as a RMAB problem. For all our experiments, we consider a population with 2100 beneficiaries (i.e., 2100 arms) and the number of live service calls that can be made every week is limited to 210.

We use the Phi 3-Mini 3B model, and the noise-injection and training procedures follow the same setup as in Section 5.1. Here, we experiment only with the noiseless and low-noise settings to better reflect realistic conditions, rather than injecting artificially large noise.

Preference Dataset Construction We generate a preference dataset of health worker prioritization commands, preferred, and rejected reward functions. Since it is costly to generate this dataset through human annotation, we use an LLM judge with ChatGPT 4o-mini (OpenAI 2024) as follows: i) query LLM-judge to obtain 20 candidate reward functions that align with the prioritization command, ii) for each candidate reward function, solve RMAB problem using Whittle index method (Whittle 1988) and generate trajectory outputs (see (Niño-Mora 2023) or the Appendix in (Kim et al. 2025) for the explanation on the Whittle index), iii) sample 50 pairs of reward functions from this set and query the LLM-judge to select preferred reward function, iv) perform this query 10 times to estimate the LLM’s uncertainty over preferences. Finally, we obtain a dataset of 9500 preferred and rejected reward function responses over 190 prioritization prompts.

Baselines In addition to the DPO variants, we also compare against the Decision Language Model (DLM), a self-reflection-based inference-time approach introduced in (Behari et al. 2024) using Gemini 2.0. Note that DLM is not fine-tuned on the training data, so this method has no notion of noise. For each prompt, we simply apply the out-of-the-box Gemini model and run the self-reflection procedure described in (Behari et al. 2024). We use Gemini 2.0, as in the original work, rather than the open-source Phi-3B model used for fine-tuning. This is to assess how our method compares against a state-of-the-art closed model that a typical user would rely on, further strengthened by self-reflection.

Evaluation To evaluate the quality of each generated reward function, we follow a procedure parallel to the dataset construction. For every prompt in the evaluation set, we generate a reward function using a given fine-tuning method and obtain the corresponding RMAB policy using the Whittle Index. We repeat the same procedure using the chosen reward function associated with that prompt in the evaluation data, producing a reference policy for comparison.

We compare these policies using two complementary methods. First, because this task provides access to ground-truth human-designed reward functions from Behari et al. (2024), we can directly evaluate RMAB policy quality by computing the achieved reward under this ground-truth function. We compare the average ground-truth reward obtained by the policy induced by the fine-tuned LLM with that obtained by the policy from the evaluation data. Second, we also perform an LLM-as-a-judge evaluation using GPT-4o-mini to compute win-rate as in Section 5.1. The judge compares the resulting trajectories of the RMAB policies, i.e., the resource allocation outcomes across the demographic data, to determine which policy is preferred.

Results We report the results in Table 2. For the ground-truth reward metric, all variants of DPO-PRO outperform both DrDPO and Vanilla DPO in both the noiseless and high-noise settings. Their performance is slightly lower than DLM, but this gap must be interpreted in context: DLM requires (1) access to a high-end proprietary LLM API at deployment time and (2) substantially higher inference-time cost, as we show below.

For the win-rate metric, the results are more mixed. The DPO-PRO variants generally outperform DrDPO and Vanilla DPO in the low-noise setting, with the exception of the $\rho = 0.03$ variant. In particular, DPO-PRO ($\rho = 0.1$) outperforms both DrDPO and Vanilla DPO in both the noiseless and low-noise settings, while DPO-PRO ($\rho = 0.3$) performs similarly to these baselines under no noise and clearly exceeds them under low noise. As with the ground-truth reward metric, all DPO-PRO variants perform below DLM.

One interesting observation is that, unlike the clean trends in Section 5.1, the performance across different ρ values and noise levels in the ARMMAN task does not yield a simple or monotonic pattern. For example, larger ρ values are not consistently more conservative, and performance even improves as noise is injected. Although pinpointing the exact cause is difficult, several factors likely contribute to this behavior.

First, unlike the curated benchmark data in Section 5.1, the ARMMAN dataset is substantially noisier. Preference labels generated by LLM judges are difficult to obtain reliably for this task (see Section 3), reducing the quality of the signal available during fine-tuning. In other words, the noiseless setting here does not necessarily correspond to a clean preference signal, since the underlying annotations are already highly noisy. Second, DPO and related off-policy preference-based methods are known to be sensitive to overfitting and reward hacking when trained on fixed datasets (Gheshlaghi Azar et al. 2024; Amini, Vieira, and Cotterell 2024; Yang et al. 2024; Xiao et al. 2025). In such settings, injected noise can act as a regularizer: it discourages the model




Figure 1: Inference-time comparison of DPO-PRO and DLM across different population sizes (i.e., number of arms) in the underlying RMAB.

from fitting spurious or inconsistent preference patterns and promotes more stable decision rules. Similar trends, where performance improves under moderate noise, have been reported in (Wu et al. 2025; Bukharin et al. 2024). This interpretation is also consistent with recent theoretical work showing that controlled noise injection can improve generalization by implicitly encouraging simpler and more robust decision boundaries (Li et al. 2025).

Inference Time Comparison Figure 1 compares the inference time of DPO-PRO (with $\rho = 0.1$) and DLM. Unlike DLM, whose inference time scales linearly with the number of arms due to repeated simulation of resource-allocation policy and self-reflection, DPO-PRO incurs a constant inference cost regardless of problem size. This efficiency gap becomes increasingly significant in large-scale settings, making DPO-PRO more scalable. While overall performance is comparable or mixed across evaluation metrics, minimal inference-time overhead makes DPO-PRO especially well-suited for real-world deployment.

Examples To further illustrate the differences between vanilla DPO and DPO-PRO (with $\rho = 0.1$), we present response outputs from both models on two representative tasks from the evaluation dataset (Table 3). Many tasks contain attributes that can be interpreted in multiple ways. For instance, the concept of “young beneficiaries” can be identified directly through the age feature, but also indirectly via lower education level or income. Similarly, a task that expresses a preference for “midday calls” could refer explicitly to the 12:30–3:00 PM time slot, but might also include the 10:30–12:30 PM window, depending on interpretation.

In such cases, we observe that vanilla-DPO favors uncommon interpretations of ambiguous attributes based on spurious correlations in noisy preference data. In contrast, DPO-PRO adopts a cautious strategy by selecting narrower definition, aiming to be robust to uncertainty in the preference signal at evaluation.

6 Conclusion

Motivated by reward-function design in RMABs for public health, we introduced DPO-PRO, a robust fine-tuning algorithm that extends DPO by explicitly modeling uncertainty in preference distributions through a lightweight

(a) Average Ground Truth Reward (\uparrow).

Method	No noise	Low noise
DLM	6073.04	NA
DPO	5465.36	5707.09
DrDPO	5480.06	5675.17
DPO-PRO ($\rho = 0.008$)	5530.58	5832.66
DPO-PRO ($\rho = 0.03$)	5557.9	5790.8
DPO-PRO ($\rho = 0.1$)	5677.76	5807.83
DPO-PRO ($\rho = 0.3$)	5797.1	5904.7

(b) Win rate (% , \uparrow).

Method	No noise	Low noise
DLM	24.82	NA
DPO	15.38	13.84
DrDPO	15.89	12.13
DPO-PRO ($\rho = 0.008$)	14.18	14.87
DPO-PRO ($\rho = 0.03$)	13.5	13.5
DPO-PRO ($\rho = 0.1$)	17.7	14.35
DPO-PRO ($\rho = 0.3$)	15.7	17.8

Table 2: Results on ARMMAN environment under no-noise and low-noise settings.

Task 1: Prefer both young and elderly beneficiaries

Vanilla-DPO	$s + 3 * (\text{youngest_age or second_youngest_age or oldest_age}) + 2 * (\text{lowest_education or second_lowest_education or third_lowest_education}) + (\text{lowest_income or second_lowest_income or third_lowest_income})$
DPO-PRO	$s + 3 * (\text{youngest_age or second_youngest_age or oldest_age}) + 2 * (\text{lowest_education or second_lowest_education or third_lowest_education})$

Task 2: Prioritize midday calls and NGO-registered users

Vanilla-DPO	$s + 3 * (\text{10_30-12_30pm and NGO_registered}) + 2 * (\text{12_30-3pm and NGO_registered})$
DPO-PRO	$s + 3 * (\text{12_30-3pm and NGO_registered})$

Table 3: Examples of reward functions produced by Vanilla-DPO and DPO-PRO when trained under noise. DPO-PRO produces more conservative preference interpretations.

DRO formulation. Our method is particularly well-suited for high-stakes decision-making settings such as public health, where preference annotations are inherently noisy and often limited. Compared to existing DRO-based approaches, DPO-PRO avoids excessive conservatism while offering a theoretical interpretation as a regularized variant of DPO. Through extensive evaluation on real-world public health data and standard alignment benchmarks, we showed that DPO-PRO improves robustness and achieves performance competitive with inference-time methods at a fraction of their computational overhead. Moreover, compared with the previous inference-time approach, DPO-PRO eliminates the need for external LLM APIs at deployment, an essential property when handling sensitive or private data. These results highlight the potential of DPO-PRO as a scalable and reliable tool for learning aligned reward functions in real-world policy learning tasks for public health.

Acknowledgements

This work was supported by ONR MURI N00014-24-1-2742 and based upon work supported by the AI Research Institutes Program funded by the National Science Foundation under the AI Institute for Societal Decision Making (NSF AI-SDM), Award No. 2229881.

References

- Abdin, M.; Aneja, J.; Awadalla, H.; Awad; et al. 2024. Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone. Technical Report MSR-TR-2024-12, Microsoft. Introduces `phi-3-mini`, a 3.8B-parameter model trained on 3.3T tokens.
- Amini, A.; Vieira, T.; and Cotterell, R. 2024. Direct Preference Optimization with an Offset. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 9954–9972. Bangkok, Thailand: Association for Computational Linguistics.
- ARMMAN. 2025. ARMMAN: Advancing Reduction in Mortality and Morbidity of Mothers, Children, and Neonates.
- Behari, N.; Zhang, E.; ZHAO, Y.; Taneja, A.; Nagaraj, D. M.; and Tambe, M. 2024. A Decision-Language Model (DLM) for Dynamic Restless Multi-Armed Bandit Tasks in Public Health. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Bertsekas, D. P. 1999. *Nonlinear programming*. Belmont, Mass.: Athena Scientific, 2nd ed. edition.
- Booth, S.; Knox, W. B.; Shah, J.; Niekum, S.; Stone, P.; and Allievi, A. 2023. The Perils of Trial-and-Error Reward Design: Misdesign through Overfitting and Invalid Task Specifications. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(5): 5920–5929.
- 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. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Carta, T.; Oudeyer, P.-Y.; Sigaud, O.; and sylvain lamprier. 2022. EAGER: Asking and Answering Questions for Automatic Reward Shaping in Language-guided RL. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.
- Cen, S.; Mei, J.; Goshvadi, K.; Dai, H.; Yang, T.; Yang, S.; Schuurmans, D.; Chi, Y.; and Dai, B. 2025. Value-Incentivized Preference Optimization: A Unified Approach to Online and Offline RLHF. In *The Thirteenth International Conference on Learning Representations*.
- Chowdhury, S. R.; Kini, A.; and Natarajan, N. 2024. Provably robust DPO: aligning language models with noisy feedback. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Cui, G.; Yuan, L.; Ding, N.; Yao, G.; He, B.; Zhu, W.; Ni, Y.; Xie, G.; Xie, R.; Lin, Y.; Liu, Z.; and Sun, M. 2024. UL-TRAFEEEDBACK: boosting language models with scaled AI feedback. In *Proceedings of the 41st International Conference on Machine Learning*. JMLR.org.
- Gheshlaghi Azar, M.; Daniel Guo, Z.; Piot, B.; Munos, R.; Rowland, M.; Valko, M.; and Calandriello, D. 2024. A General Theoretical Paradigm to Understand Learning from Human Preferences. In Dasgupta, S.; Mandt, S.; and Li, Y., eds., *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, 4447–4455. PMLR.
- Goyal, P.; Niekum, S.; and Mooney, R. J. 2019. Using Natural Language for Reward Shaping in Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 2385–2391. International Joint Conferences on Artificial Intelligence Organization.
- Helpmum. 2025. Improving Maternal and Infant Healthcare in Africa with AI and low cost Innovations.
- Hong, I.; Li, Z.; Bukharin, A.; Li, Y.; Jiang, H.; Yang, T.; and Zhao, T. 2024. Adaptive Preference Scaling for Reinforcement Learning with Human Feedback. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Jahan, R.; Zou, P.; Huang, Y.; and Jibb, L. 2018. Impact of MomConnect Program in South Africa: A Narrative Review. *Online Journal of Nursing Informatics*.
- Kim, C. W.; Verma, S.; Tec, M.; and Tambe, M. 2025. Lightweight Robust Direct Preference Optimization. In *NeurIPS 2025 Workshop: Reliable ML from Unreliable Data*.
- Lee, H.; Phatale, S.; Mansoor, H.; Mesnard, T.; Ferret, J.; Lu, K.; Bishop, C.; Hall, E.; Carbune, V.; Rastogi, A.; and Prakash, S. 2024. RLAIIF vs. RLHF: scaling reinforcement learning from human feedback with AI feedback. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Li, J.; Pan, J.; Toh, K.-C.; and Zhou, P. 2025. Towards Understanding Why Data Augmentation Improves Generalization. *arXiv preprint arXiv:2502.08940*.
- Ma, Y. J.; Liang, W.; Wang, G.; Huang, D.-A.; Bastani, O.; Jayaraman, D.; Zhu, Y.; Fan, L.; and Anandkumar, A. 2024. Eureka: Human-Level Reward Design via Coding Large Language Models. In *The Twelfth International Conference on Learning Representations*.
- Mandal, D.; Sasnauskas, P.; and Radanovic, G. 2025. Distributionally Robust Reinforcement Learning with Human Feedback. *arXiv:2503.00539*.
- Mirchandani, S.; Karamcheti, S.; and Sadigh, D. 2021. ELLA: exploration through learned language abstraction. In *Proceedings of the 35th International Conference on Neural Information Processing Systems, NIPS '21*. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781713845393.
- Niño-Mora, J. 2023. Markovian Restless Bandits and Index Policies: A Review. *Mathematics*, 11(7).

- OpenAI. 2024. GPT-4o mini: Advancing cost-efficient intelligence.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 53728–53741. Curran Associates, Inc.
- Rahimian, H.; and Mehrotra, S. 2022. Frameworks and Results in Distributionally Robust Optimization. *Open Journal of Mathematical Optimization*, 3: 1–85.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. arXiv:1707.06347.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K. R.; and Yao, S. 2023. Reflexion: language agents with verbal reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Singh, S.; Lewis, R. L.; and Barto, A. G. 2009. Where Do Rewards Come From? In *Proceedings of the 31st Annual Meeting of the Cognitive Science Society*, 2601–2606. Amsterdam, The Netherlands.
- Tan, Z.; Li, D.; Wang, S.; Beigi, A.; Jiang, B.; Bhattacharjee, A.; Karami, M.; Li, J.; Cheng, L.; and Liu, H. 2024. Large Language Models for Data Annotation and Synthesis: A Survey. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 930–957. Miami, Florida, USA: Association for Computational Linguistics.
- Verma, S.; Boehmer, N.; Kong, L.; and Tambe, M. 2025. Balancing Act: Prioritization Strategies for LLM-Designed Restless Bandit Rewards. arXiv:2408.12112.
- Whittle, P. 1988. Restless Bandits: Activity Allocation in a Changing World. *Journal of Applied Probability*, 25: 287–298.
- Wu, J.; Xie, Y.; Yang, Z.; Wu, J.; Chen, J.; Gao, J.; Ding, B.; Wang, X.; and He, X. 2025. Towards Robust Alignment of Language Models: Distributionally Robustifying Direct Preference Optimization. In *The Thirteenth International Conference on Learning Representations*.
- Wu, Y.; Sun, Z.; Li, S.; Welleck, S.; and Yang, Y. 2024. Inference scaling laws: An empirical analysis of compute-optimal inference for problem-solving with language models. arXiv preprint arXiv:2408.00724.
- Xiao, W.; Wang, Z.; Gan, L.; Zhao, S.; Li, Z.; Lei, R.; He, W.; Tuan, L. A.; Chen, L.; Jiang, H.; Zhao, Z.; and Wu, F. 2025. A Comprehensive Survey of Direct Preference Optimization: Datasets, Theories, Variants, and Applications. arXiv:2410.15595.
- Xie, T.; Foster, D. J.; Krishnamurthy, A.; Rosset, C.; Awadallah, A. H.; and Rakhlin, A. 2025. Exploratory Preference Optimization: Harnessing Implicit Q^* -Approximation for Sample-Efficient RLHF. In *The Thirteenth International Conference on Learning Representations*.
- Xiong, W.; Dong, H.; Ye, C.; Wang, Z.; Zhong, H.; Ji, H.; Jiang, N.; and Zhang, T. 2024. Iterative Preference Learning from Human Feedback: Bridging Theory and Practice for RLHF under KL-constraint. In *Forty-first International Conference on Machine Learning*.
- Xu, Z.; Vemuri, S.; Panaganti, K.; Kalathil, D.; Jain, R.; and Ramachandran, D. 2025. Distributionally Robust Direct Preference Optimization. arXiv:2502.01930.
- Yang, C.; Jia, R.; Gu, N.; Lin, Z.; Chen, S.; Pang, C.; Yin, W.; Sun, Y.; Wu, H.; and Wang, W. 2024. Orthogonal Finetuning for Direct Preference Optimization. arXiv:2409.14836.
- Yuan, L.; Cui, G.; Wang, H.; Ding, N.; Wang, X.; Deng, J.; Shan, B.; Chen, H.; Xie, R.; Lin, Y.; Liu, Z.; Zhou, B.; Peng, H.; Liu, Z.; and Sun, M. 2025. Advancing LLM Reasoning Generalists with Preference Trees. In *International Conference on Learning Representations (ICLR) 2025*.
- Zhan, W.; Uehara, M.; Kallus, N.; Lee, J. D.; and Sun, W. 2024. Provable Offline Preference-Based Reinforcement Learning. In *The Twelfth International Conference on Learning Representations*.
- Zhang, C.; Chong, D.; Jiang, F.; Tang, C.; Gao, A.; Tang, G.; and Li, H. 2025. Aligning Language Models Using Follow-up Likelihood as Reward Signal. In *AAAI*, 25832–25841.
- Zhao, Y.; Boehmer, N.; Taneja, A.; and Tambe, M. 2025. Towards Foundation-model-based Multiagent System to Accelerate AI for Social Impact. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025)*, AAMAS '25, 2901–2907. Detroit, MI, USA: International Foundation for Autonomous Agents and Multiagent Systems. Blue Sky Ideas Track.
- Zhu, B.; Frick, E.; Wu, T.; Zhu, H.; Ganesan, K.; Chiang, W.-L.; Zhang, J.; and Jiao, J. 2024. Starling-7B: Improving Helpfulness and Harmlessness with RLAIIF. In *First Conference on Language Modeling*.