Mutual-Taught for Boosting Policy and Reward Models

Anonymous ACL submission

Abstract

Preference optimization has proven effective in aligning large language models (LLMs) with human objectives. However, as LLMs evolve during training, distribution shifts arise between the model-generated samples and the reward model, which diminishes the reward model's effectiveness and, in turn, constrains the policy's performance. In this paper, we propose to enhance both the policy and reward models collaboratively without relying on external supervision. We present **Mutual-Taught**, a self-training technique based on the Expectation-Maximization (EM) algorithm. In the *E-step*, the policy model is fine-tuned using feedback from the reward model to better approximate the optimal preference distribution. In the *M-step*, the reward model is updated by creating training data based on the outputs of the policy model before and after the update. Experimental results demonstrate that this iterative process steadily improves both models' performance. Our 8B policy model, LLaMA3-8B-Instruct-MT, achieves a length-controlled win rate of 52.0% on AlpacaEval-2. Moreover, our 8B reward model, FsfairX-LLaMA3-RM-MT, performs comparably to the strong GPT-4o-2024-08-06 on RewardBench.

1 Introduction

As large language models (LLMs) are fine-tuned to align with human preferences using techniques such as Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022) and Direct Preference Optimization (DPO) (Rafailov et al., 2024), distributional shifts may arise between the model-generated outputs and the reward model's assessments. As the model evolves, the distribution of the text it produces may gradually diverge from the distribution on which the reward model or preference dataset was originally trained. This mismatch between the model's evolving outputs and the preferences encoded in the reward model or preference data can create a feedback loop that amplifies errors and misalignments over time. For example, in RLHF, the model might generate outputs that receive high reward scores not because they genuinely capture human preferences, but because they exploit particular weaknesses or biases in the reward model—a phenomenon known as "reward hacking" (Gao et al., 2023; Zheng et al., 2023b). Similarly, in DPO, shifts in data distribution may hinder the effectiveness of preference learning. Addressing these challenges is essential for maintaining robust and reliable alignment with human preferences as models continue to evolve.

One approach to addressing this issue is to collect additional human preference annotations for newly generated samples and update the reward model to account for the shifted distribution (Touvron et al., 2023). However, this method is not scalable due to its heavy reliance on human annotations. An alternative strategy is to use LLM-as-a-Judge prompting (Yuan et al., 2024; Wu et al., 2024a), where the language model itself provides rewards during training. These techniques employ iterative DPO to improve both instruction-following and judgment capabilities, but they typically require strong base models or pre-training on judgment datasets to establish effective judgment skills.

In this paper, we investigate whether both the policy and reward models can be enhanced automatically, without relying on external supervision signals. Our central research question is: *How can we automatically acquire high-quality feedback for LLM self-training and effectively guide reward model optimization?* To address this, we propose a self-training framework based on the Expectation-Maximization (EM) algorithm, named **Mutual-Taught**. The Mutual-Taught framework iteratively updates both the policy and reward models in a collaborative manner. During the E-step (policy model update), we fine-tune the policy model using feedback from the current reward model, which

seeks to estimate the optimal preference distribution under the policy. In the M-step (reward model update), the reward model is updated by constructing training data through contrastive comparisons between the outputs of the policy model before and after the update. These pseudo-preference pairs capture the improvements made by the policy and enable the reward model to adapt to the policy's distribution without requiring external annotations.

A key insight of our approach is that the distribution shift arising from policy updates can be leveraged to automatically generate contrastive examples for updating the reward model, thereby ensuring alignment between the policy and reward models throughout the training process. Experimental results demonstrate that Mutual-Taught outperforms previous methods by continuously improving both the policy and reward models without relying on human annotations. Notably, the reward model exhibits strong generalization, effectively guiding the optimization of various policy models.

2 Related Work

Offline preference optimization Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al., 2022) has emerged as a pivotal approach of preference optimization. However, it depends on reinforcement learning techniques such as Proximal Policy Optimization (PPO) (Schulman et al., 2017), which are challenging to implement and often unstable during training. To simplify and stabilize the RLHF process, Direct Preference Optimization (DPO) (Rafailov et al., 2024) was proposed. DPO trains a policy model directly from human-annotated preference pairs using a simple classification loss derived from the Bradley-Terry model (Bradley and Terry, 1952). Besides DPO, various preference optimization objectives have been proposed to improve performance and simplify training, including SLiC-HF (Zhao et al., 2023), KTO (Ethayarajh et al., 2024), RSO (Liu et al., 2023), ORPO (Hong et al., 2024), and SimPO (Meng et al., 2024). However, the absence of an explicit reward model in these methods limits their ability to adapt to the evolving policy distribution and to generate new preference pairs effectively.

Iterative preference optimization With an explicit reward model, the preference optimization methods mentioned above can be applied repeatedly over multiple rounds. In each round, new data is generated by the policy obtained from the previous round, annotated with a reward score and then used to train a stronger policy. For example, Xu et al. (2023) apply PCO and DPO to an iterative manner by annotating new preferences with a fixed reward model. ReST^{EM} (Singh et al., 2023) proposed an Expectation-Maximization (EM) framework, where in each round, the new policy is obtained by minimizing the reward-weighted negative log-likelihood loss on data generated by the old policy. SELM (Zhang et al., 2024b) and XPO (Xie et al., 2024) augment the DPO objective with a novel and principled exploration bonus, enabling the algorithm to explore beyond the initial model and human feedback data. SAIL (Ding et al., 2024) generates preference data by combining initial offline preferences with model-generated preferences, where the model itself estimates the likelihood of one response being preferred over another using its learned policy. SPIN (Chen et al., 2024), DNO (Rosset et al., 2024) and SPPO (Wu et al., 2024b) leverage self-play mechanism to iteratively refine the policy towards achieving Nash equilibrium by optimizing general preferences. However, these approaches typically employ a fixed reward model, which can constrain further policy improvements due to distribution shifts. To address this issue, Ouyang et al. (2022) collects preference data on the current best policy, which is then annotated by humans and used to train a new reward model, this process consumes significant annotation resources. ReST-MCTS* (Zhang et al., 2024a) uses the policy to perform a modified Monte Carlo Tree Search to generate solutions and evaluates them against the ground truth for process reward model training. However, its reliance on ground truth restricts the method's applicability to only a few specific domains. Another line of work uses the policy model to judge its own responses in an LLM-as-a-Judge mechanism (Zheng et al., 2023a), which eliminates the reward model and simultaneously updates the knowledge of the judge. Self-rewarding (Yuan et al., 2024) and Meta-rewarding (Wu et al., 2024a) language models generate responses to prompts using the current policy and then assign scores to each response themselves to create preference data for the next iteration of training.

3 Preliminaries

3.1 Reward Modeling

In RLHF, a reward model $r(y; x)$ is first trained to predict human preference scores for responses y

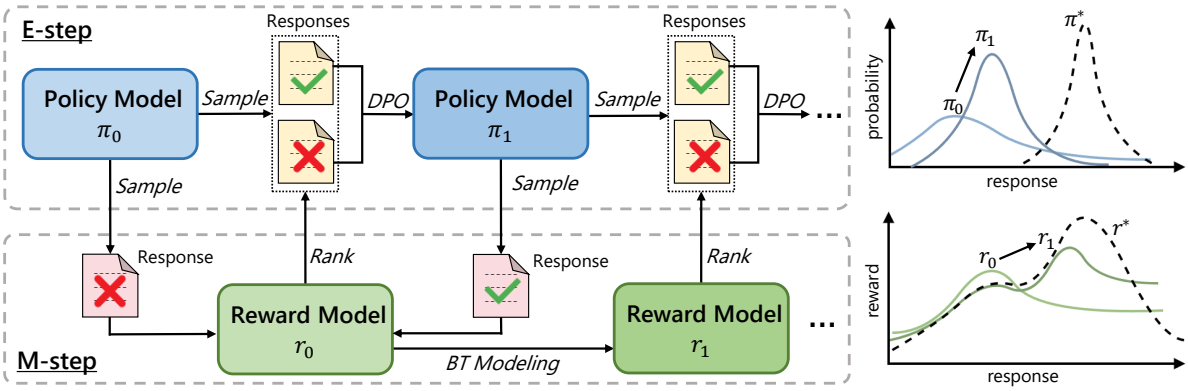

Figure 1: Overview of the Mutual-Taught framework: The process alternates between policy model updates (E-step) and reward model updates (M-step). In the E-step, the policy model is fine-tuned using feedback from the current reward model to align the policy with the optimal preference distribution. In the M-step, the reward model is updated through contrastive comparisons of policy model outputs before and after updates, allowing it to adapt to the evolving policy without requiring additional human annotations. This iterative process mitigates distributional shifts and improves both models synergistically.

given prompts x. The reward model is typically trained using human-annotated preference pairs (x, y_w, y_l), where y_w is preferred over y_l.

Bradley-Terry reward model (Bradley and Terry, 1952) is commonly used to model the probability that one response is preferred over another:

$$P(y_w \succ y_l \mid x) = \sigma(r(y_w; x) - r(y_l; x)) = \frac{\exp(r(y_w; x))}{\exp(r(y_w; x)) + \exp(r(y_l; x))}, \quad (1)$$

The reward model is then trained by maximizing the log-likelihood of observed preferences: $\log P(y_w \succ y_l \mid x)$.

3.2 Direct Preference Optimization (DPO)

DPO (Rafailov et al., 2024) simplifies the training process by combining the two-step procedure of PPO (Schulman et al., 2017) into a single unified objective. Specifically, DPO conducts a closed-form solution for the reward function, yieding the following loss formulation:

$$\mathcal{L}_{\text{DPO}} = -\log \sigma \left(\beta \log \frac{\pi_\theta(y_w \mid x)}{\pi_{\text{ref}}(y_w \mid x)} - \beta \log \frac{\pi_\theta(y_l \mid x)}{\pi_{\text{ref}}(y_l \mid x)} \right). \quad (2)$$

While DPO offers enhanced stability and ease of optimization, its offline nature and the absence of an explicit reward model limit its ability to adapt to the evolving policy distribution effectively.

4 Mutual-Taught Preference Optimization

In conventional iterative preference learning frameworks, the reward model is often assumed to represent the ground-truth preference distribution and is kept fixed during policy optimization. However, as the policy model improves and generates responses that differ from the data used to train the reward model, a distribution shift occurs (Touvron et al., 2023; Cheng et al., 2024). This mismatch can lead to suboptimal guidance from the reward model, constraining the performance upper bound of the policy model. To address this limitation, we introduce Mutual-Taught, a self-training framework that leverages a modified Expectation-Maximization (EM) algorithm to iteratively improve both the policy and reward models without requiring additional human feedback.

4.1 Objective of Mutual-Taught

Our key idea is to treat the policy and reward models as mutually reinforcing components that can be improved together through iterative updates. Specifically, we aim to maximize the likelihood of the observed data under a joint model of the policy and reward models. Let π denote the policy model, and r denote the reward model. Our objective is to find the optimal policy π^* and reward model r^* that maximize the expected reward:

$$\pi^*, r^* = \arg\max_{\pi, r} \mathbb{E}_{x \sim \mathcal{D}, y \sim \pi(\cdot \mid x)} \left[r(y; x) \right]. \quad (3)$$

However, directly optimizing this objective is challenging due to the interdependence between π and r. To overcome this, we adopt an EM-like iterative optimization procedure.

Our Mutual-Taught algorithm operates in two

Algorithm 1 Mutual-Taught Preference Optimization

Input: Initial policy model π_0, initial reward model r_0, dataset $\mathcal{D}$, number of iterations T
Data Partitioning: Split $\mathcal{D}$ into $T+1$ disjoint subsets: $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_T, \mathcal{D}_R$
for each iteration $t = 1, 2, \ldots, T$ **do**
  E-step (Policy model update):
  Generate response samples $y_t^1, y_t^2, ...y_t^K \sim \pi_{t-1}(\cdot \mid x)$ for each prompt $x \sim \mathcal{D}_t$.
  Compute reward $r_{t-1}(y_t^k; x)$ using the current reward model.
  Update the policy model π_t by maximizing the expected reward, e.g., via DPO.
  M-step (Reward model update):
  Generate comparison pair (x, y_t, y_{t-1}) by sampling from π_t and π_{t-1} for each prompt $x \sim \mathcal{D}_R$.
  Update the reward model r_t by training on the generated comparison pairs $\{(x, y_t, y_{t-1})\}$.
end for
Output: Optimized policy model π_T, optimized reward model r_T

main steps at each iteration, analogous to the E-step and M-step in the EM algorithm:

E-step (Policy model update): With the reward model r_{t-1} fixed, we update the policy model π_t by maximizing the expected reward:

$$\pi_t = \arg\max_{\pi} \mathbb{E}_{x\sim\mathcal{D}, y\sim\pi(\cdot|x)} \left[r_{t-1}(y; x) \right]. \quad (4)$$

This step involves generating responses using the current policy and updating the policy model to increase the likelihood of high-reward responses according to r_{t-1}.

M-step (Reward model update): With the policy model π_t fixed, we update the reward model r_t by maximizing the likelihood of the policy's improvement over the previous iteration:

$$r_t = \arg\max_{r} \mathbb{E}_{x\sim\mathcal{D}, y_t\sim\pi_t(\cdot|x), y_{t-1}\sim\pi_{t-1}(\cdot|x)} \left[\log P_r(y_t \succ y_{t-1} \mid x) \right], \quad (5)$$

where $P_r(y_t \succ y_{t-1} \mid x)$ is the preference probability estimated by the reward model r.

By comparing the outputs of the current and previous policy models, we generate preference pairs (x, y_t, y_{t-1}) without human annotations. The reward model is updated to assign higher scores to the improved responses from π_t. This iterative process allows the policy and reward models to co-evolve, with the policy improving under the guidance of the reward model, and the reward model adapting to the policy's improvements.

4.2 Algorithm Details

The Mutual-Taught Preference Optimization algorithm is summarized in Algorithm 1.

In conventional EM algorithms, the E-step involves computing the expected value of the latent variables given the current parameters, and the M-step involves maximizing the expected complete-data log-likelihood. In our Mutual-Taught framework, we adapt the EM concept to jointly optimize two models by treating the policy model's outputs as estimates of the latent preference distribution, which in turn updates the reward model. A critical aspect of our approach is leveraging the distribution shift between policy iterations. While distribution shift is often considered a challenge in RLHF, we turn it into an advantage by using it to generate high-quality training data for the reward model. By utilizing the policy's outputs across iterations, we create preference pairs that reflect policy improvements, enabling the reward model to adapt to the evolving policy distribution.

5 Experiments

5.1 Experimental Setup

Base model and dataset We use LLaMA3-8B-Instruct (Dubey et al., 2024) as our base policy model and FsfairX-LLaMA3-RM-v0.1 (Xiong et al., 2024) as the base reward model, which is one of the top models on RewardBench (Lambert et al., 2024) and provides open-source code, facilitating iterative training. Following previous work, we use the UltraFeedback dataset (Cui et al., 2024) with approximately 60,000 prompts from diverse sources. To prevent overfitting during fine-tuning, we divide the dataset into three parts: two parts for policy model iteration and one for reward model iteration. In each iteration, we generate $K = 5$ responses per prompt with a temperature of 0.8 and top-p $= 0.95$. Duplicate generations are removed.

Evaluation benchmarks Since Mutual-Taught aims to automatically improve both the policy

model and the reward model, we evaluate the performance of each component separately.

Policy model instruction following: We utilize two widely recognized automatic evaluation benchmarks where GPT-4 acts as the judge: AlpacaEval-2 (Li et al., 2023) and Arena-Hard (Li et al., 2024). Each benchmark targets different aspects of model performance. AlpacaEval-2 assesses chat capabilities using 805 instructions spanning a wide range of prompts, evaluated through length-controlled (LC) win rate and raw win rate (WR) metrics. Arena-Hard presents more challenging tasks, including 500 well-defined technical problem-solving queries.

Reward model judgment accuracy: We assess the reward model's accuracy using RewardBench (Lambert et al., 2024), which measures performance across four different categories: Chat, Chat-Hard, Safety, and Reasoning.

Baselines We compare our approach against the following baseline models and methods:

- Base policy model and base reward model: We use LLaMA3-8B-Instruct, an instruction-following LLM developed by Meta, as our base policy model. For the reward model, we employ FsfairX-LLaMA3-RM-v0.1, a high-performing reward model fine-tuned from LLaMA3-8B-Instruct.

- Offline preference optimization methods: We implement DPO (Rafailov et al., 2024), IPO (Gheshlaghi Azar et al., 2024) and SimPO (Meng et al., 2024). Preference pairs are derived from multiple responses generated by the base policy model, with scores provided by the base reward model.

- Iterative preference optimization methods: We implement SPPO (Wu et al., 2024b) and Meta-rewarding (Wu et al., 2024a). Since these methods do not update the reward model, we use all three portions of the dataset for policy training, performing three iterations.

To ensure a fair comparison, the sampling settings used in the above experiments align with those applied in the Mutual-Taught. Further implementation details for these baselines can be found in Appendix A.

Implementation details We conducted Mutual-Taught between the policy and reward models for two iterations. In each iteration, both models were trained for one epoch using a cosine learning rate schedule with a warmup ratio of 0.1. For each policy model iteration, we initialized the model from the previous round and generated responses using the current policy. Preference data was then derived using the reward model at the current iteration. The policy model was optimized via Direct Preference Optimization (DPO) with a beta of 0.01, a batch size of 128, a maximum sequence length of 2,048 tokens, and a learning rate of 7×10^{-7}. To mitigate the risk of overfitting on the same prompts across iterations, each reward model iteration started from the base reward model. The reward model was trained on preference pairs consisting of chosen and rejected responses sampled from the current and preceding policy models. We used a batch size of 512, a maximum sequence length of 2,048 tokens, and an empirically set learning rate of 2×10^{-6}. All experiments were conducted on 8 NVIDIA A100 GPUs. Further details are provided in Appendix A.

5.2 Main Results

Iterative performance improvement on policy model We first report the performance of Mutual-Taught and baseline methods on the instruction-following benchmarks AlpacaEval-2 and Arena-Hard in Table 1. For AlpacaEval-2 and Arena-Hard, Mutual-Taught delivers substantial improvements to the LLaMA-3-8B-Instruct model, achieving a 28.9% increase in length-controlled (LC) win rate and a 16.9% increase in win rate, respectively. Notably, the proposed method provides consistent improvements in each iteration, validating the robustness of our approach. Compared to other baselines, our method demonstrates clear superiority on AlpacaEval-2 and Arena-Hard. Additionally, when comparing iterative preference optimization methods, we observe that methods using a reward model for preference feedback (e.g., SPPO and Mutual-Taught) perform significantly better than methods relying on LLM-as-a-judge feedback (e.g., Meta-rewarding). This indicates that a reward model, fine-tuned through supervised training, offers stronger initial judgment capabilities than the policy model itself and provides more effective guidance throughout the iterative process.

It is noteworthy that our method employs only two-thirds of the available datasets for updating the policy model, reserving the remaining third for updating the reward model. Despite using less data for policy model iterations compared to other

Model	AlpacaEval-2			Arena-Hard	
	LC Win Rate	Win Rate	Avg. Len	Win Rate	Avg. Len
Base Policy Model					
LLaMA-3-8B-Instruct	23.1	23.1	1899	20.6	585
Offline Preference Optimization Methods					
SimPO	47.9	46.3	1934	32.5	552
IPO	43.7	42.1	1899	34.5	569
DPO	44.7	42.7	1945	33.1	557
Iterative Preference Optimization Methods					
Meta-rewarding Iter1	34.2	32.6	1893	27.7	531
Meta-rewarding Iter2	36.4	34.5	1876	27.0	530
Meta-rewarding Iter3	37.5 (↑ 14.4)	35.2	1868	27.9 (↑ 7.3)	530
SPPO Iter1	39.4	39.5	2021	30.6	570
SPPO Iter2	41.0	44.4	2396	34.4	653
SPPO Iter3	46.4 (↑ 23.3)	48.5	2128	33.6 (↑ 13.0)	542
DPO Iter1	33.6	33.8	1989	30.3	559
DPO Iter2	43.4	42.3	1961	33.3	587
DPO Iter3	47.2 (↑ 24.1)	48.7	1930	34.7 (↑ 14.1)	571
Our Methods					
Mutual-Taught Iter1	37.1	36.5	1957	33.5	553
Mutual-Taught Iter2	**52.0** (↑ 28.9)	**56.0**	2214	**37.5** (↑ 16.9)	692

Table 1: Overall results of our proposed Mutual-Taught method with LLaMA-3-8B-Instruct as the policy model, compared against various baseline methods on AlpacaEval-2 and Arena-Hard. The improvement is calculated relative to LLaMA-3-8B-Instruct. Text in bold indicates the best performance.

Model	Chat	Chat Hard	Safety	Reasoning	Average
GPT-4o-2024-08-06	96.10	76.10	88.10	86.60	86.70
FsfairX-LLaMA3-RM-v0.1	99.40	65.10	87.80	86.40	84.70
Mutual-Taught Iter1	98.32	62.61	84.86	96.60	85.60
Mutual-Taught Iter2	98.32	65.90	87.26	95.69	86.80

Table 2: Out-of-distribution (OOD) evaluation results of the reward models on RewardBench.

iterative training baselines, we achieve significantly better performance on AlpacaEval-2 and Arena-Hard. This outcome highlights the importance of synchronously updating both the policy and reward models during the iterative training process. We believe that enhancing the reward model can yield greater benefits than simply increasing the amount of data used to train the policy model.

Iterative performance improvement on reward model To assess the effectiveness of Mutual-Taught in improving the reward model, we evaluate the reward models obtained during the iterative process from two perspectives:

In-distribution (ID): This test evaluates whether the reward model (RM) improves in selecting optimal responses after each iteration. Specifically, we assess the RM's performance using the iteration data employed during RM training. For each evaluation prompt, the policy model (after two iterations) generates five candidate responses. The base RM, along with the RMs from the first and second iterations, then selects the best response from these candidates. To measure performance differences between the RMs, GPT-4 serves as a judgment model, conducting pairwise comparisons of the responses selected by the different RMs. As shown in Figure 2, the resulting RM achieves a progressively higher win rate against the base RM as iterations advance. This highlights that the RM's ability to identify high-quality responses improves with each iteration, contributing to the enhancement of the policy model in subsequent iterations.

Out-of-distribution (OOD): We further evaluate the RM's generalization ability under OOD conditions using RewardBench. As shown in Table 2, the

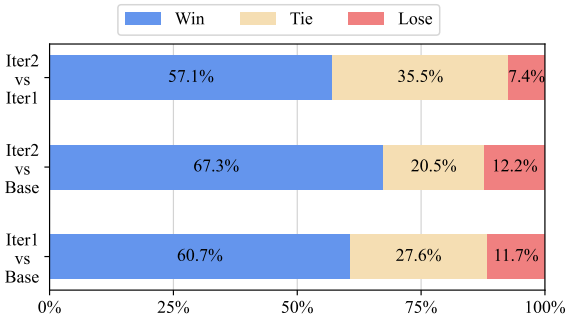

Figure 2: In-distribution (ID) evaluation results of the reward models. We compare the reward models at different iterations and show the win, tie, lose rates.

RM achieves consistent improvements after each iteration, with an average score increase of 2.1 points after two iterations, approaching the performance of GPT-4o-2024-08-06. Notably, the most significant contribution to the improvement comes from the enhancement in reasoning capabilities, with FsfairX-iter-2 achieving a 9.29 points increase in reasoning compared to the base RM. We attribute this to the strong reasoning ability of LLaMA3-8B-Instruct, which provides high-quality feedback on reasoning prompts to guide the RM effectively.

5.3 Ablation Study

Our main hypothesis is that the updated policy distribution is, on average, superior to the previous policy distribution. This improvement enables the reward model to learn a better preference distribution from the current data. To understand the underlying reasons for the effectiveness of Mutual-Taught, we conduct an ablation study focusing on two aspects: the impact of updating the reward model on policy iteration, and the effect of data synthesis strategies on reward model iteration.

As illustrated in Figure 3, if the reward model is not updated, the overall iterative training degenerates into iterative DPO, leading to a significant decline in the policy model's performance after iteration. This observation underscores the effectiveness of Mutual-Taught in optimizing the reward model by leveraging comparisons between the policy model's outputs before and after updates. In our implementation, to prevent excessive knowledge forgetting during the optimization of the reward model's distribution, each iteration of RM training incorporates two types of data: **D (PM)**: Data generated by the current policy model (self-training data). **D ($\mathbf{PM_{new}}$, $\mathbf{PM_{old}}$)**: Contrastive data comparing outputs from the new and old policy models. Both data types are unseen by the reward model, and using either type alone could independently improve RM performance. Notably, using only D (PM) is akin to self-training for the RM. To investigate the impact of each data type on RM performance and the overall iterative process, we conduct experiments by replacing the training data with only a single data type at each RM iteration.

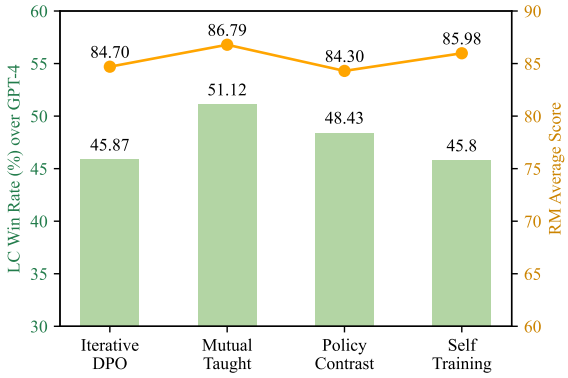

Figure 3: The impact of different reward model data synthesis strategies on the performance of Mutual-Taught.

As shown in Figure 3, the policy model's performance declines in both data-type ablation scenarios. Specifically, when only self-training data is used, the policy model's performance drops by 5.32 points, though the RM performance does not show a significant decline. When using only policy comparison data, the RM performance declines slightly, but the policy model's performance is relatively less affected.

We hypothesize that self-training data tends to reflect the original distribution of the RM, which helps prevent catastrophic forgetting but struggles to model better preference distributions, making it less effective in guiding the policy model for the next iteration. In contrast, using comparison data between the updated and previous policy models aligns more closely with the iterative optimization goal, allowing the RM to approach a better preference distribution and provide more effective feedback for the next policy model update. The combination of both data types in Mutual-Taught achieves a balance between preventing knowledge forgetting and modeling improved preference distributions, leading to better iterative performance than using either data type alone.

5.4 Further Analysis

Generalization of the iterated reward model A key question is whether the iterated reward model

(RM) can generalize effectively to guide the optimization of other policy models, given that the RM training data is based on outputs from LLaMA3-8B-Instruct. To address this, we use the RM obtained after two iterations and apply it to train a different policy model, Mistral-7B-Instruct-v0.2 (Jiang et al., 2023), via a single round of DPO training on the UltraFeedback dataset. As shown in Table 3, the iterated RM improves the performance of the Mistral-7B-Instruct-v0.2 on AlpacaEval-2 by up to 3.77 points. This demonstrates that the RM, after undergoing Mutual-Taught iterations, remains effective at guiding the optimization of other policy models under the current data distribution. We believe this is due to the RM's enhanced ability to capture general preference distributions, which can benefit different models.

Model	AlpacaEval-2	
	LC Win Rate	Win Rate
Mistral-7B-Instruct-v0.2	19.39	15.75
w/ RM-Base	41.96	42.77
w/ RM-Iter1	44.81	43.29
w/ RM-Iter2	**45.73**	**50.02**

Table 3: Results from using reward models at different iterations in the main experiment to guide DPO training of Mistral-7B-Instruct-v0.2.

Compatibility with different preference objectives In the main experiments, we primarily used DPO to optimize the policy model during the E-step. We also explored integrating Mutual Taught with different preference optimization objectives, specifically SimPO and IPO.

As shown in Table 4, Mutual Taught consistently provides significant improvements in preference optimization across both scenarios, highlighting its strong compatibility with various preference learning objectives. The pairwise win rates, measured by FsfairX-LLaMA3-RM-v0.1 (Xiong et al., 2024), are presented in Figure 4. In all preference optimization objectives, updated models consistently outperformed previous ones. However, SimPO surpassed IPO in the final iteration. Interestingly, IPO performed better than SimPO on standard benchmarks evaluated by GPT-4 against ground-truth answers. We attribute SimPO's final iteration advantage to its tendency to generate longer sequences, exploiting the length bias in FsfairX-LLaMA3-RM-v0.1, which favors longer outputs.

Model	AlpacaEval-2		
	LC Win Rate	Win Rate	Avg. Len
SimPO	47.94	46.25	1934
IPO	43.72	42.11	1899
SimPO-MT	49.78	49.61	2147
IPO-MT	50.73	50.29	2029

Table 4: Experimental results with different preference optimization objectives in Mutual-Taught E-step. SimPO-MT and IPO-MT represent iterative training with SimPO and IPO, respectively.

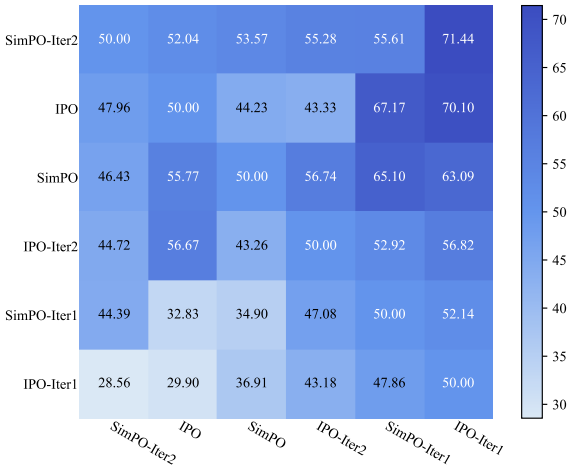

Figure 4: Pairwise evaluation of models with different preference optimization objectives in the Mutual-Taught framework on Alpaca-Eval 2 using FsfairX-LLaMA3-RM-v0.1.

6 Conclusion

This paper introduces Mutual-Taught Preference Optimization, an approach to automatically improve policy and reward models without relying on external supervision signals. The method follows an expectation-maximization (EM)-based iterative process, where in each iteration, the policy model is improved using preference feedback from the reward model to provide better observations for training the reward model. Then, comparisons between the policy model's observations before and after updates are leveraged to optimize the reward model's distribution. We demonstrated that this iterative process can continuously enhance both the policy and reward models. The resulting policy model achieves significant improvements over existing methods, such as DPO, SPPO, and Meta-rewarding, across multiple benchmarks, including AlpacaEval-2 and Arena-Hard. Moreover, the iterated reward model achieves performance comparable to GPT-4o-2024-08-06 on RewardBench.

Limitations

While Mutual-Taught demonstrates promising results, it relies on the assumption that the policy model's improvements can be effectively captured and utilized by the reward model through self-generated data. In scenarios where the policy model does not improve significantly across iterations, the effectiveness of this method may be limited. Additionally, the approach requires careful tuning of hyperparameters to balance the updates between the policy and reward models.

Ethics Statement

All experiments in this study were conducted using publicly available datasets that do not contain any private information. Our work does not involve the analysis or utilization of identity characteristics, and we do not engage in any form of gender or racial discrimination.

References

Ralph Allan Bradley and Milton E. Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39:324.

Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. In *International Conference on Machine Learning*.

Pengyu Cheng, Yifan Yang, Jian Li, Yong Dai, Tianhao Hu, Peixin Cao, Nan Du, and Xiaolong Li. 2024. Adversarial preference optimization: Enhancing your alignment via rm-llm game. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3705–3716.

Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2024. UltraFeedback: Boosting language models with high-quality feedback. In *International Conference on Machine Learning*.

Mucong Ding, Souradip Chakraborty, Vibhu Agrawal, Zora Che, Alec Koppel, Mengdi Wang, Amrit Bedi, and Furong Huang. 2024. Sail: Self-improving efficient online alignment of large language models. *arXiv preprint arXiv:2406.15567*.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model alignment as prospect theoretic optimization. *arXiv preprint arXiv:2402.01306*.

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

Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. A general theoretical paradigm to understand learning from human preferences. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4447–4455. PMLR.

Jiwoo Hong, Noah Lee, and James Thorne. 2024. Reference-free monolithic preference optimization with odds ratio. *arXiv preprint arXiv:2403.07691*.

AQ Jiang, A Sablayrolles, A Mensch, C Bamford, DS Chaplot, D de las Casas, F Bressand, G Lengyel, G Lample, L Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, et al. 2024. Openassistant conversations-democratizing large language model alignment. *Advances in Neural Information Processing Systems*, 36.

Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, et al. 2024. Rewardbench: Evaluating reward models for language modeling. *arXiv preprint arXiv:2403.13787*.

Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gonzalez, and Ion Stoica. 2024. From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline. *arXiv preprint arXiv:2406.11939*.

Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Alpacaeval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.

Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and Jialu Liu. 2023. Statistical rejection sampling improves preference optimization. *arXiv preprint arXiv:2309.06657*.

Yu Meng, Mengzhou Xia, and Danqi Chen. 2024. Simpo: Simple preference optimization with a reference-free reward. *arXiv preprint arXiv:2405.14734*.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al.

2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.

Corby Rosset, Ching-An Cheng, Arindam Mitra, Michael Santacroce, Ahmed Awadallah, and Tengyang Xie. 2024. Direct nash optimization: Teaching language models to self-improve with general preferences. *arXiv preprint arXiv:2404.03715*.

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

Avi Singh, John D Co-Reyes, Rishabh Agarwal, Ankesh Anand, Piyush Patil, Peter J Liu, James Harrison, Jaehoon Lee, Kelvin Xu, Aaron Parisi, et al. 2023. Beyond human data: Scaling self-training for problem-solving with language models. *arXiv preprint arXiv:2312.06585*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Shengyi Huang, Kashif Rasul, Alvaro Bartolome, Alexander M. Rush, and Thomas Wolf. The Alignment Handbook. https://github.com/huggingface/alignment-handbook.

Tianhao Wu, Weizhe Yuan, Olga Golovneva, Jing Xu, Yuandong Tian, Jiantao Jiao, Jason Weston, and Sainbayar Sukhbaatar. 2024a. Meta-rewarding language models: Self-improving alignment with llm-as-a-meta-judge. *arXiv preprint arXiv:2407.19594*.

Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. 2024b. Self-play preference optimization for language model alignment. *arXiv preprint arXiv:2405.00675*.

Tengyang Xie, Dylan J Foster, Akshay Krishnamurthy, Corby Rosset, Ahmed Awadallah, and Alexander Rakhlin. 2024. Exploratory preference optimization: Harnessing implicit q*-approximation for sample-efficient rlhf. *arXiv preprint arXiv:2405.21046*.

Wei Xiong, Hanze Dong, Chenlu Ye, Ziqi Wang, Han Zhong, Heng Ji, Nan Jiang, and Tong Zhang. 2024. Iterative preference learning from human feedback: Bridging theory and practice for rlhf under kl-constraint. In *Forty-first International Conference on Machine Learning*.

Jing Xu, Andrew Lee, Sainbayar Sukhbaatar, and Jason Weston. 2023. Some things are more cringe than others: Preference optimization with the pairwise cringe loss. *arXiv preprint arXiv:2312.16682*.

Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Sainbayar Sukhbaatar, Jing Xu, and Jason Weston. 2024. Self-rewarding language models. *arXiv preprint arXiv:2401.10020*.

Dan Zhang, Sining Zhoubian, Yisong Yue, Yuxiao Dong, and Jie Tang. 2024a. Rest-mcts*: Llm self-training via process reward guided tree search. *arXiv preprint arXiv:2406.03816*.

Shenao Zhang, Donghan Yu, Hiteshi Sharma, Ziyi Yang, Shuohang Wang, Hany Hassan, and Zhaoran Wang. 2024b. Self-exploring language models: Active preference elicitation for online alignment. *arXiv preprint arXiv:2405.19332*.

Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. 2023. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv preprint arXiv:2305.10425*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023a. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

Rui Zheng, Wei Shen, Yuan Hua, Wenbin Lai, Shihan Dou, Yuhao Zhou, Zhiheng Xi, Xiao Wang, Haoran Huang, Tao Gui, et al. 2023b. Improving generalization of alignment with human preferences through group invariant learning. *arXiv preprint arXiv:2310.11971*.

A Experiments Details

In our experiments, we use the Alignment Handbook framework (Tunstall et al.) for policy model iteration and the RLHF-Reward-Modeling [1] framework for reward model iteration.

Mutual taught training We follow SimPO (Meng et al., 2024) to set the policy sampling and training parameters. Specifically, for policy sampling: temperature is set to 0.8, $K = 5$, and top-p to 0.95. For policy training: learning rate is set to 7×10^{-7}, batch size to 128, and warmup ratio to 0.1. These settings remain consistent across both iterations. For reward model iteration, we use the default parameter settings from RLHF-Reward-Modeling. In both iterations, we use the same settings: learning rate of 2×10^{-6}, batch size of 512, and weight decay of 0.001. Reward model

[1] RLHF-Reward-Modeling at https://github.com/RLHFlow/RLHF-Reward-Modeling

training data is constructed by sampling from the policy distributions before and after updates, with a temperature of 0.8 and top-p of 0.95.

Baselines In Offline Preference Optimization Methods, we maintain the same sampling and training parameters as Mutual-Taught. For Iterative Preference Optimization Methods, in iterative DPO, we observed performance degradation in the final iteration with a larger learning rate, so we lowered it to 5×10^{-7}. For SPPO, we use the default training parameters provided by the method. For Meta-Rewarding, we first built Evaluation Fine-Tuning (EFT) data from the Open Assistant (Köpf et al., 2024) dataset to boost the initial judgment ability of the model before self-training iterations. During the construction of EFT data, we prompt GPT-4o to generate judgments with high quality instead of the SFT baseline in Yuan et al. (2024). During self-training iterations, we use prompts from the UltraFeedback dataset instead of those generated by LLaMA2-70B-Chat to align with Mutual-Taught.

Length-control To prevent length explosion, we implement a length-control mechanism for selecting preference data. For each prompt, we first select responses with above-average reward scores, then choose the shortest one as the chosen response. The response with the lowest score is selected as the rejected one. This length control mechanism is applied to all experiments except for Meta-Rewarding, where we use the length control mechanism proposed by the original method.