

# HAF-RM: A Hybrid Alignment Framework for Reward Model Training

Anonymous ACL submission

## Abstract

The reward model has become increasingly important in alignment, assessment, and data construction for large language models (LLMs). Most existing researchers focus on enhancing reward models through data improvements, following the conventional training framework for reward models that directly optimizes the predicted rewards. In this paper, we propose a hybrid alignment framework HAF-RM for reward model training by introducing an additional constraint on token-level policy probabilities in addition to the reward score. It can simultaneously supervise the internal preference model at the token level and optimize the mapping layer of the reward model at the sequence level. Theoretical justifications and experiment results on five datasets show the validity and effectiveness of our proposed hybrid framework for training a high-quality reward model. By decoupling the reward modeling procedure and incorporating hybrid supervision, our HAF-RM framework offers a principled and effective approach to enhancing the performance and alignment of reward models, a critical component in the responsible development of powerful language models. We release our code at <https://haf-rm-anonymized.github.io>.

## 1 Introduction

Recent periods have witnessed a continuous evolution of Large Language Model (LLM) techniques, especially pre-training (Devlin et al., 2019; Radford et al., 2019; Brown et al., 2020) and instruction tuning (Wei et al., 2021; Wang et al., 2022; Yue et al., 2023). Researchers start to shift their focus from generating correct responses to aligning responses more closely with human preferences (Russell, 2014). As an efficient alternative to human feedback, the reward model for generative language models emerges, facilitating scalable alignment in training (Christiano et al., 2017; Stiennon et al., 2020), response generation (Gao et al., 2023; Mud-

gal et al., 2024; Jinnai et al., 2024), data construction (Yuan et al., 2023) etc.

Despite the availability of numerous sophisticated reward models (Kopf et al., 2023; Zhu et al., 2023), these exist several key limitations. First, most reward models originate from industry and are not open-source, making further training and transfer impossible. Second, prior studies have highlighted incorrect and ambiguous preferences within the training data of these reward models (Bai et al., 2022; Pitis, 2023). These two issues both limit the quality and generalizability of existing reward models, necessitating further enhancement either from the data perspective or the training process. While recent researches mainly focus on enriching data sources for better reward models, including utilizing external tools or information sources to enhance generalization (Li et al., 2023a; Sun et al., 2023) or leveraging fine-grained signals (Wu et al., 2023; Cao et al., 2024) and their combinations (Go et al., 2023; Lai et al., 2024), we focus on the training framework of reward models in this work.

A reward model is typically structured with two components: a transformer-based model (referred to as the internal preference model) that outputs preference vectors for each token, and a projection module called “reward layer” (usually a linear layer with normalization) that maps these vectors to sequence-level rewards. The standard practice for training the reward model involves utilizing the ranking loss of paired rewards. However, optimizing both two components using such a single sequence-level objective may cause insufficient supervision for token-level preference modeling. We argue that hybrid optimization of the two components of the reward model with corresponding token-level and sequence-level objectives will lead to more consistent improvement.

Since a policy model is also based on an internal preference model to predict the expected reward

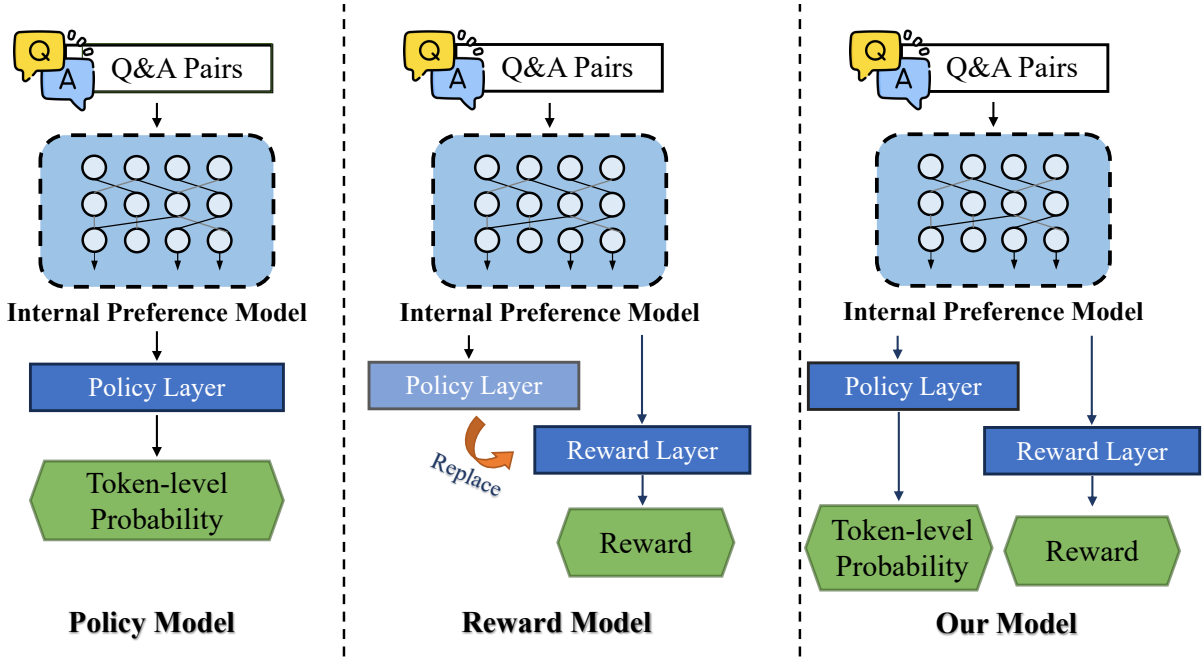


Figure 1: The standard reward model substitutes the policy layer from the policy model, while our HAF model retains the policy layer. By optimizing the model’s two outputs, we achieve a better alignment process for the reward model with little additional training overhead.

for each action/token, essentially acting as a  $Q$ -function under token-level supervision (Rafailov et al., 2024), we propose a Hybrid Alignment Framework (HAF). This framework jointly optimizes the reward model and policy model with a shared internal preference model. With the policy loss, we can directly supervise the internal preference model at the token level while simultaneously optimizing the mapping layer of the reward model using the reward loss, enabling more effective alignment of the reward model.

We provide both theoretical justifications and empirical experiments to demonstrate the effectiveness of our HAF. In the experiment section, we compare the performance of reward models trained using our framework against those resulting from baseline approaches across four public datasets. The results highlight the advantage of HAF with different policy losses integrated. Further analysis reveals that using additional policy loss can improve the performance of policy model calibration, which opens a new horizon for training high-quality reward models.

## 2 Preliminary

The objective of our framework is to train the reward model  $r$  based on a pairwise comparison dataset (also known as “preference dataset”)  $\mathcal{D}$ ,

following typical reward model training settings.

### 2.1 Notation

- $\mathcal{D} = \{(x_i, y_i, y'_i)\}_{i=1}^n$  represents the dataset used to train the reward model, where  $x_i$ ,  $y_i$  and  $y'_i$  are the query, preferred and non-preferred responses respectively.
- $\mathcal{P} = \{(x, y) \mid (x, y, y') \in \mathcal{D}\} \cup \{(x, y') \mid (x, y, y') \in \mathcal{D}\}$  is the set of query-response pairs from the dataset  $\mathcal{D}$ .
- $r$  is the **reward model** which can be split into two parts as  $r(x, y) = F \circ \phi(x, y)$ , to output the reward of a response  $y$  given a query  $x$ . Here,  $\phi(\cdot, \cdot)$  denotes the model’s internal preference model, while  $F$  serves as the reward prediction layer mapping the model’s internal preference to the final reward. We use the symbol  $\circ$  to signify function nesting, i.e.,  $F \circ \phi(x, y) = F(\phi(x, y))$ .
- $\pi$  is the **policy model**, and  $\pi(x, y)$  is the generation probability of  $y$  given  $x$ . It can also be divided into two parts as  $\pi(x, y) = K \circ \phi(x, y)$  where the policy prediction layer  $K$  maps the model’s internal preference to the generation probability.
- The **oracle (optimal) value** is denoted as

the corresponding letter with an asterisk such as  $\mathbf{r}^*$  (oracle reward model),  $\phi^*$  (optimal model preference),  $F^*$  (optimal reward prediction layer) and  $K^*$  (optimal policy prediction layer).

## 2.2 Training Loss

We use  $D_1$  to represent the distribution discrepancy between the reward model’s output and the oracle reward model’s output, and  $D_2$  for the outputs of the policy model and the oracle policy model.

**Reward Loss** The standard reward loss  $\mathcal{L}_s$  considers the precision of rewards alone, being a simple and direct metric to quantify the quality of a reward model.

$$\mathcal{L}_s := \mathbb{E}_{d \sim \mathcal{P}} [D_1(\mathbf{r}(d), \mathbf{r}^*(d))] \quad (1)$$

For notational convenience, we use  $d$  to denote  $(x, y)$  and use  $\operatorname{argmin}_{\mathbf{r}} \mathcal{L}_s$  or  $\operatorname{argmin}_{F, \phi} \mathcal{L}_s$  to represent the model training with the standard reward loss.

**Policy Loss** Similar to the reward loss, standard policy loss aims to measure the error of the policy model.

$$\mathcal{L}_P := \mathbb{E}_{d \sim \mathcal{P}} [D_2(\boldsymbol{\pi}(d), \boldsymbol{\pi}^*(d))] \quad (2)$$

**Hybrid Alignment Loss** To fully leverage the similarity between the reward model and the policy model, we incorporate an additional supervising term  $D_2$  on the policy model into the loss function. By calibrating the shared preference space, we effectively align the model in a hybrid manner:

$$\begin{aligned} \mathcal{L}_H &:= \mathbb{E}_{d \sim \mathcal{P}} [D_1(\mathbf{r}(d), \mathbf{r}^*(d)) \\ &\quad + \alpha \cdot D_2(\boldsymbol{\pi}(d), \boldsymbol{\pi}^*(d))] \\ &= \mathbb{E}_{d \sim \mathcal{P}} [D_1(F \circ \phi(d), F^* \circ \phi^*(d)) \\ &\quad + \alpha \cdot D_2(K \circ \phi(d), K^* \circ \phi^*(d))] \end{aligned} \quad (3)$$

where  $\alpha$  is a hyperparameter to balance losses from the reward and policy model,  $\phi$  is the shared internal preference model which receives gradients from both loss terms. Similarly,  $\operatorname{argmin}_{F, K, \phi} \mathcal{L}_H$  and  $\operatorname{argmin}_{\mathbf{r}, \boldsymbol{\pi}} \mathcal{L}_H$  represent the model training with our hybrid alignment loss.

## 3 Hybrid Alignment Framework

### 3.1 Model Implementation

The most commonly used decoder-only LLM consists of stacked transformer blocks (Vaswani et al., 2017) or similar structures, and a linear layer for policy projection. In the reward model, only the shape of the final linear layer is adjusted to match the format of the reward value output compared to the policy model. We retain two linear layers for our model, enabling it to output rewards and probabilities simultaneously.

To significantly reduce the resources required for training, it is standard practice to initialize the internal preference module of the reward model with a fine-tuned language model as it retains the model’s language modeling capabilities.

### 3.2 Loss Calculation

There is consensus on the specific calculation method for the **reward loss**. In avoiding the issue of uncertain reward values, the Bradley-Terry model (Christiano et al., 2017) is used to transform the reward modeling problem into a probability optimization problem. Treating the problem as a binary classification task yields the popular form of reward loss function:

$$\begin{aligned} \mathcal{L}_s &= \mathbb{E}_{d \sim \mathcal{P}} [D_1(\mathbf{r}(d), \mathbf{r}^*(d))] \\ &= \mathbb{E}_{(x, y, y') \sim \mathcal{D}} [-\log \sigma(\mathbf{r}(x, y) - \mathbf{r}(x, y'))] \end{aligned} \quad (4)$$

where  $\sigma(\cdot)$  is the sigmoid function.

Given the preference data, there currently does not exist a universally optimal **policy loss**. However, since the derivation of the DPO loss is based on assumptions similar to those made for the reward loss (as detailed in Appendix C.2), we choose to use the DPO loss as the method for calculating the policy loss.

$$\begin{aligned} \mathcal{L}_P &= \mathbb{E}_{d \sim \mathcal{P}} [D_2(\boldsymbol{\pi}(d), \boldsymbol{\pi}^*(d))] \\ &= \mathbb{E}_{(x, y, y') \sim \mathcal{D}} [-\log \sigma(\tau(pd_{win} - pd_{lose}))] \end{aligned} \quad (5)$$

where

$$pd_{win} = \log \frac{\pi(x, y)}{\pi_{ref}(x, y)}, \quad pd_{lose} = \log \frac{\pi(x, y')}{\pi_{ref}(x, y')}.$$

$\pi_{ref}$  is the reference policy model and  $\tau$  is the hyperparameter set to 0.1.

Combining the two losses, we have our HAF loss calculate in the following manner:

$$\mathcal{L}_H = \mathcal{L}_s + \alpha \cdot \mathcal{L}_P \quad (6)$$

We will elaborate in Appendix C.1 on why Eq. 4 and Eq. 5 hold and why there is no optimal model on the right-hand side.

### 3.3 Theoretical Analysis

In this subsection, we present several properties of HAF that are independent of the specific calculation methods of the two loss functions. We will start from Section 2.2.

In practice, functions such as  $F$  and  $\phi$  are represented by parameterized models with finite parameters, and thus cannot precisely model arbitrary distributions. Here we show that under certain assumptions, using the hybrid alignment loss can yield a better solution than simply using the standard reward loss.

**Proposition 1.** *Unless  $K$  can exactly fit  $K^*$ , there exists  $\epsilon > 0$ , such that*

$$\begin{aligned} & \mathbb{E}_{d \sim \mathcal{P}} [D_2(K_H \circ \phi_H(d), K^* \circ \phi^*(d))] \\ & \leq \min_K \mathbb{E}_{d \sim \mathcal{P}} [D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))] - \frac{\epsilon}{\alpha} \end{aligned}$$

holds for all  $\alpha \in (0.1, 2)$ , where  $K_H, \phi_H = \operatorname{argmin}_{K, \phi} \mathcal{L}_H$  in Equation 3 and  $\phi_s = \operatorname{argmin}_{\phi} \mathcal{L}_s$  in Equation 4.

Here we use  $\operatorname{argmin}$  to represent the best models optimized with the corresponding loss functions, so  $\phi_H$  and  $\phi_s$  are not equal to  $\phi^*$  although  $\phi^*$  is the minimum mathematically. Intuitively this indicates that *the model learned from the joint calibrated loss outperforms the one learned solely from the preference space using the standard reward loss.*

**Proposition 2.** *Assume that  $\phi^*$  is unique,  $K^*$  is locally Lipschitz continuous, , and  $0.1 < \alpha < 2$ , there exists  $k, \delta > 0$ , such that*

$$\begin{aligned} & \mathbb{E}_{d \sim \mathcal{P}} [|\phi_H(d) - \phi^*(d)| - |\phi_s(d) - \phi^*(d)|] < \\ & \frac{g_{\max} - g_{\min}}{g_{\min}} \mathbb{E}_{d \sim \mathcal{P}} |\phi_s(d) - \phi^*(d)| + 2\delta - \frac{\epsilon}{\alpha \cdot k} \end{aligned}$$

The detailed derivations for both propositions are provided in Appendix D. Here we obtain an upper bound on the model preference error. By tuning the hyperparameter  $\alpha$ , the right term can be strictly negative. In other words, *model preference space trained with our calibrated loss can be strictly closer to the true preference space compared to the standard reward loss.* (In practice, there is no need for an exhaustive search, we find  $\alpha = 0.2$  already yields satisfactory results. We give a discussion about this in B)

Name	Size	Words/QA	Tokens/QA
Harmless	12,915	42.9	61.5
Helpful	13,543	54.3	77.2
BS	47,625	69.3	88.5
AHP	8,722	59.6	81.9
CA	19,466	165.5	257.6

Table 1: Statistics of the Training Datasets

## 4 Experiment setup

### 4.1 Datasets

We comprehensively assess the performance of our framework using five public datasets, namely Anthropic-HH-Harmless (HH-harmless) (Bai et al., 2022), Anthropic-HH-Helpful (HH-Helpful) (Bai et al., 2022), Beaver Safe (BS) (Ji et al., 2023), Alpaca Human Pref (AHP) (Dubois et al., 2023) and Chatbot Arena (CA) (Zheng et al., 2023). Note that AHP and CA do not have original data split for evaluation, we randomly extract 10% from the original data as a test set, the details of the used datasets are shown in Tab 1.

### 4.2 Comparative Models

**Baseline** We compare our framework with the standard training approach, in which the reward model only has a reward layer for reward prediction and is optimized via Eq. 4.

**DPO** Although DPO loss (Eq. 5) is typically used for training policy models rather than reward models, it can implicitly convert the model’s outputs into reward values (Rafailov et al., 2023). Therefore, the DPO model can also be considered a reward model (Rafailov et al., 2024). Following the work of Lambert et al. (2024), we also evaluate the model trained with DPO loss.

**HAF** Under our framework, the reward model has both the reward and policy layer for predicting sequence-level rewards and providing token-level probabilities.

In our implementation, we use Phi-2-2.7B and Mistral-7B-Instruct-v0.2 as our base model. We train Phi-2 and Mistral-7B using full-parameter and Low-rank Adaptation (LoRA) (Hu et al., 2022), respectively. More experiment setup can be found in Appendix A.

Method	Helpful	Harmless	CA	BS	AHP	Avg
DPO(Phi-2)	<u>69.70</u>	66.30	66.80	<b>87.80</b>	52.60	68.64
Baseline(Phi-2)	64.30	<u>69.50</u>	<b>79.30</b>	76.00	<u>58.40</u>	69.50
HAF (Phi-2)	<b>76.40</b>	<b>70.40</b>	<u>79.00</u>	<u>84.00</u>	<b>60.80</b>	<b>74.12</b>
DPO(Mistral)	74.29	70.30	<b>81.90</b>	<b>92.70</b>	<u>60.30</u>	75.90
Baseline(Mistral)	<b>76.20</b>	<u>72.70</u>	79.80	80.80	56.30	73.16
HAF (Mistral)	<u>75.80</u>	<b>73.10</b>	<b>81.90</b>	<u>88.70</u>	<b>63.10</b>	<b>76.52</b>

Table 2: Overall results on each dataset for accuracy, which denotes the proportion that the better response is scored higher. The best performance is highlighted in boldface and the suboptimal result is underlined.




Figure 2: Comparison of models trained with HAF/baseline/DPO methods on the mixed dataset.

## 5 Experiment Results

### 5.1 Intrinsic performance of Reward Models

The primary function of a reward model is to evaluate the quality of responses to a given question, which involves accurately comparing two answers to the same question. Using judgment accuracy as the evaluation metric, we conduct several experiments to assess the effectiveness of HAF in training the reward model.

#### 5.1.1 Overall Performance

Firstly we compare the performance of HAF with the baseline and two judging models across the five datasets. Table 2 presents the overall results. HAF has higher accuracy than the baseline in most cases, indicating that the model can more sensitively identify whether an answer is good and give a more accurate high (or low) score. At the same time, those results worse than the baseline or DPO are generally only slightly worse, indicating that our method is basically not weaker than the baseline under various circumstances.

#### 5.1.2 Mixed Data

For the mixed data setting, we construct two datasets by sampling and combining examples

from multiple sources: Anthropic-HH (Anthropic Helpful + Anthropic Harmless) and Mixed (evenly sampled from each of the five datasets in our corpus). As shown in Figure 2, our proposed hybrid alignment framework achieves the best generalization performance across all reward models when evaluated on these mixed data distributions. This suggests our approach can better learn the diversity present in the combined datasets for generalization.

#### 5.1.3 OOD Data

Data within the same dataset often exhibits certain distributional similarities due to similar or even identical data cleaning and processing methods. To simulate a distribution shift in real-world application, we also evaluate generalization to entirely held-out OOD datasets. Specifically, we train models on one dataset and evaluate on the remaining four. Although different datasets have distinct distributions, their main preferences can be generalized as “overall better” (AHP, CA and Helpful) and “safer” (BS and Harmless), which we use **rAcc** (“**r**” stands for “relevant”) to represent model’s generalization ability within similar preferences.

The results are detailed in Table 3. We can tell from the table that the **rAcc** of HAF is basically higher than that of both Baseline and DPO, indi-

Acc(%)	AHP	CA	Helpful	BS	Harmless	AVG	rAcc
Phi-2-2.7B							
AHP	*	67.40 <sup>(1.00↓)</sup> <sub>(30.30↑)</sub>	67.60 <sup>(3.40↑)</sup> <sub>(17.10↑)</sub>	39.80 <sup>(0.20↑)</sup> <sub>(14.80↓)</sub>	41.90 <sup>(5.40↓)</sup> <sub>(9.30↓)</sub>	54.18 <sup>(0.70↓)</sup> <sub>(5.83↑)</sub>	67.50 <sup>(1.20↑)</sup> <sub>(23.70↑)</sub>
CA	60.20 <sup>(0.50↑)</sup> <sub>(8.20↑)</sub>	*	64.70 <sup>(3.20↓)</sup> <sub>(15.00↑)</sub>	37.60 <sup>(0.80↑)</sup> <sub>(13.00↓)</sub>	42.10 <sup>(5.60↑)</sup> <sub>(9.40↓)</sub>	51.15 <sup>(0.92↑)</sup> <sub>(0.20↑)</sub>	62.45 <sup>(1.35↓)</sup> <sub>(11.60↑)</sub>
Helpful	60.20 <sup>(2.90↑)</sup> <sub>(6.90↑)</sub>	72.00 <sup>(1.10↓)</sup> <sub>(32.70↑)</sub>	*	36.20 <sup>(1.40↓)</sup> <sub>(10.30↓)</sub>	38.50 <sup>(6.90↓)</sup> <sub>(0.30↓)</sub>	51.73 <sup>(1.62↓)</sup> <sub>(7.25↑)</sub>	66.10 <sup>(0.90↑)</sup> <sub>(19.80↑)</sub>
BS	47.90 <sup>(0.20↓)</sup> <sub>(1.20↓)</sub>	41.00 <sup>(2.50↑)</sup> <sub>(9.20↑)</sub>	35.70 <sup>(1.40↓)</sup> <sub>(9.30↓)</sub>	*	70.60 <sup>(5.60↑)</sup> <sub>(4.60↑)</sub>	48.80 <sup>(1.62↑)</sup> <sub>(0.82↑)</sub>	70.60 <sup>(5.60↑)</sup> <sub>(4.60↑)</sub>
Harmless	43.80 <sup>(1.30↑)</sup> <sub>(6.20↓)</sub>	29.40 <sup>(0.50↑)</sup> <sub>(5.70↓)</sub>	32.60 <sup>(0.80↑)</sup> <sub>(9.10↓)</sub>	76.90 <sup>(1.50↑)</sup> <sub>(8.60↑)</sub>	*	45.67 <sup>(1.02↑)</sup> <sub>(3.10↓)</sub>	76.90 <sup>(1.50↑)</sup> <sub>(8.60↑)</sub>
Mistral-7B-Instruct							
AHP	*	75.50 <sup>(6.20↑)</sup> <sub>(17.90↑)</sub>	68.90 <sup>(10.60↑)</sup> <sub>(7.60↑)</sub>	55.70 <sup>(7.20↑)</sup> <sub>(5.90↑)</sub>	48.00 <sup>(1.20↑)</sup> <sub>(1.40↓)</sub>	62.02 <sup>(6.30↑)</sup> <sub>(7.50↑)</sub>	72.20 <sup>(8.40↑)</sup> <sub>(12.75↑)</sub>
CA	60.80 <sup>(0.20↓)</sup> <sub>(6.80↑)</sub>	*	65.80 <sup>(1.20↓)</sup> <sub>(12.50↑)</sub>	38.50 <sup>(6.60↓)</sup> <sub>(3.00↓)</sub>	36.80 <sup>(4.00↓)</sup> <sub>(6.80↓)</sub>	50.47 <sup>(3.00↓)</sup> <sub>(2.37↑)</sub>	63.30 <sup>(0.70↓)</sup> <sub>(9.65↑)</sub>
Helpful	60.90 <sup>(1.00↓)</sup> <sub>(8.20↑)</sub>	73.90 <sup>(0.60↑)</sup> <sub>(20.30↑)</sub>	*	36.00 <sup>(9.30↓)</sup> <sub>(5.40↓)</sub>	37.50 <sup>(0.00)</sup> <sub>(2.10↑)</sub>	52.08 <sup>(2.42↓)</sup> <sub>(6.30↑)</sub>	67.40 <sup>(0.20↓)</sup> <sub>(14.25↑)</sub>
BS	52.90 <sup>(4.00↑)</sup> <sub>(1.10↓)</sub>	55.20 <sup>(9.50↑)</sup> <sub>(12.30↑)</sub>	43.80 <sup>(3.40↓)</sup> <sub>(8.10↓)</sub>	*	71.90 <sup>(1.40↑)</sup> <sub>(3.00↑)</sub>	55.95 <sup>(2.87↑)</sup> <sub>(1.52↑)</sub>	71.90 <sup>(1.40↑)</sup> <sub>(3.00↑)</sub>
Harmless	46.50 <sup>(1.00↑)</sup> <sub>(0.40↓)</sub>	38.30 <sup>(4.60↑)</sup> <sub>(10.50↓)</sub>	32.40 <sup>(0.50↑)</sup> <sub>(2.10↓)</sub>	76.70 <sup>(2.40↑)</sup> <sub>(5.70↑)</sub>	*	48.48 <sup>(2.13↑)</sup> <sub>(1.82↓)</sub>	76.70 <sup>(2.40↑)</sup> <sub>(5.70↑)</sub>

Table 3: Results for the OOD experiment. The results in the same row are derived from the same backbone and the same training dataset, while the columns represent different test datasets. The displayed accuracies are for HAF, with superscripts and subscripts indicating the performance differences relative to the baseline and DPO, respectively.  $\uparrow$  denotes an improvement with HAF, whereas  $\downarrow$  indicates a decline. **rAcc** is the average accuracy among grey blocks.

cating HAF possesses a strong ability to learn preferences and effectively generalize them to similar preference distributions, despite great differences in language style and topic. Touvron et al. (2023) noted that RLHF involves distributional shifts in the policy model during training, necessitating iterative training for the reward model. The robustness of HAF against such distributional shifts could potentially be a key factor in alleviating this problem.

Comparing Table 3 and Table 2, we can observe that models trained using CA or Helpful datasets outperform those directly trained on AHP dataset when the test set is AHP. This suggests two things: firstly, there is a certain similarity in preferences across the three datasets, and secondly, the amount of AHP data may be insufficient to support the complete training of the reward model, as shown in Table 1. Consequently, the model’s preference learning is incomplete, which results in low test outcomes for AHP in Table 2.

On this observation, it can be noted that the HAF-mistral model, when fully trained using CA or Helpful datasets, performs worse on BS and

Harmless compared to Baseline. However, when insufficiently trained using AHP dataset, its test results are better than the baseline. This might indicate that during the training process of the reward model, the learning of reward mapping precedes the learning of preferences. When the model is not fully trained, HAF’s advantage in learning speed enables it to outperform. Yet, once fully trained, the baseline’s weaker preference learning ability might allow it to exhibit some degree of cross-preference generalization. This hypothesis requires further validation in future work.

One easily overlooked result is that nearly all the test outcomes of the DPO model converge to approximately 50% in a highly exaggerated manner, indicating a complete loss of modeling capability for out-of-distribution data. This issue is likely related to its inherent nature as a language model: the generation process of language models exhibits strong stylistic tendencies, which, in turn, leads to a significantly higher preference for responses that align with its style (as reflected in the generation probabilities and the implicit reward values




Figure 3: Win rate of responses selected by the HAF model compared to the baseline model.

of the DPO model). Consequently, when the response distribution deviates from its stylistic norms (e.g., responses that are too short or too long, or use different vocabulary), the output probabilities become highly inaccurate. This indicates that the DPO model is not suitable for use as a conventional reward model.

## 5.2 Extrinsic Evaluation on Downstream Task

In assessing the practical applicability of reward models, intrinsic performances alone provide an incomplete picture of their efficacy. To comprehensively evaluate their utility in real-world applications, it is essential to examine how these models perform in downstream tasks that simulate practical scenarios.

This section aims to investigate the robustness and effectiveness of HAF model in such scenarios. Specifically, we explore its performance in two distinct downstream tasks: best-of-N sampling as a training-free response generation strategy (Stienon et al., 2020; Gao et al., 2023; Jinnai et al., 2024), and RLHF as a training-dependent aligning methods.

### 5.2.1 Best-of-N

We demonstrate the reliability of our trained reward model through Best-of-N pick, in which the reward model should pick the best one (the response with the highest reward) from several responses sampled from the same language model. The backbone for the reward model and the sampling model are the same, 8 and 4 responses are provided to the Mistral-based reward model the Phi-2-based reward model respectively, because Phi-2 is more likely to generate the same responses. The prompts for comparisons and ranking are listed in Appendix E, which reference AlpacaEval (Li et al., 2023b).

We report two evaluation metrics. **Win rate:** We use GPT-4-turbo to directly compare the responses

	Top-1(%)		Top-2(%)	
	HAF	Baseline	HAF	Baseline
Phi-2	<b>33.77</b>	26.68	<b>58.30</b>	49.47
Phi-2 <sub>No harm</sub>	<b>37.21</b>	28.97	<b>64.33</b>	53.41
Mistral	<b>13.31</b>	11.55	<b>25.27</b>	23.49
Mistral <sub>No harm</sub>	<b>15.70</b>	13.88	<b>29.20</b>	27.67

Table 4: Top-k recall for HAF and the baseline. There are 4 candidate responses for Phi-2 and 8 for Mistral. The results are averaged over the recall values from all five datasets. The subscript “No harm” indicates that the result in that row is averaged over the AHP, CA, and harmless datasets instead of all datasets.

from HAF reward model and baseline and report the win rate (Jang et al., 2023). **Consistency with GPT:** we use GPT-3.5-turbo to rank the sampled responses and calculate the recall of the top-1 and top-2 responses.

As shown in Figure 3 and Table 4, HAF demonstrates significant advantages over the baseline reward model in selecting responses especially for Phi-2 model in terms of both evaluation metrics. It is important to note that the average performance of the baseline reward model is comparable to random selection, suggesting that it has poor sensitivity and cannot effectively distinguish between responses when the quality differences are minimal. In contrast, the reward model obtained using HAF demonstrates good discriminative ability. Considering that the model can only learn to distinguish harmful from non-harmful responses from the BS and Harmless datasets, and that the responses generated by Phi-2 and Mistral are mostly harmless, we also report the average results on the remaining three datasets. When the safety-related datasets are excluded, both HAF and baseline show an improvement in average performance. Due to space limitations, the detailed results are presented in the appendix in Table 10.

### 5.2.2 RLHF

We also test HAF in the regular RLHF process: we train two reward models with HAF and the baseline method and then use them to train the policy models with RLHF. After training, GPT-3.5-turbo is introduced to compare the generations from the two policy models.

We conduct experiments using the Mistral model along with the AHP, CA, and Helpful datasets to investigate the reward model’s capability in optimizing for comprehensive preferences. Phi-2 is not used here as it shows great instability during

	#Win	#Lose	Win rates(%)
AHP	285	215	<b>57.00</b>
CA	346	154	<b>69.20</b>
Helpful	243	256	48.70

Table 5: Win rates for the policy model trained with HAF reward model by RLHF.

training which may not exhibit any performance improvement. Setups for reward model training and PPO are listed in Appendix A.

HAF demonstrates a significant advantage on the AHP and CA datasets, while showing slightly worse performance compared to the baseline on the Helpful dataset. This indicates that the HAF reward model provides more effective guidance for the policy model. Given the widespread application of RLHF-like methods, HAF shows promising potential for active use in language model alignment in the near future. However, due to the simple experimental setup and the inherent instability of RLHF at small scales, the effectiveness of the HAF method in language model alignment still requires extensive exploration.

## 6 Related Work

Reward model was proposed to modeling human language preferences (model that outputs preference values based on questions and answers) (Christiano et al., 2017), then the explosive growth of research on reward models (McKinney et al., 2023) and large language models (Wei et al., 2022; Park et al., 2023; Zheng et al., 2023) emerged after the popularity of ChatGPT.

From training to practical applications, an increasing number of studies have also featured the presence of quantifiable preferences (usually known as “reward”). For example, RLHF (Christiano et al., 2017; Stiennon et al., 2020) uses the PPO algorithm (Schulman et al., 2017) to maximize the reward of the policy model; RAFT (Dong et al., 2023) and RRHF (Yuan et al., 2023) remove substandard data by scoring the candidate responses with reward model; LLM-as-a-judge (Zheng et al., 2023) employs GPT-4 to score the text.

Therefore, how to construct a model offering explicit preference feedback has naturally become a focal point of much research. To train a precise and robust reward model, many studies start from training with human preference data, and many

works in the data field are largely centered around this. (Touvron et al., 2023) and (Zhao et al., 2022) provided different methods for using ranking data; (Wang et al., 2024) explored ways of measuring the strength of the data; while concerning datasets themselves, (Azar et al., 2023), (Knox et al., 2022) and (Hong et al., 2022) analyzed the impact of data preference strength on training from theoretical or practical perspectives. In addition, similar to the RAG technique (Lewis et al., 2020) in large language models, many methods (Li et al., 2023a; Sun et al., 2023) using external tools or references have also emerged, injecting new vitality into the development of reward models.

Although many data-oriented methods have greatly enhanced the performance of reward models, the field of reward model optimization has been rarely explored. Currently, the training of reward models basically follows the process proposed by OpenAI (Christiano et al., 2017). It involves initializing the reward model using a fine-tuned model, then transforming the model’s predictions into probability values through the Bradley-Terry model, and optimizing these probabilities using cross-entropy loss. Considering the widespread practical applications of reward models, the attention given to their training paradigms does not match their importance.

## 7 Conclusion

In this paper, we extend and improve the training framework of the current reward model. We split the training mechanism of the reward model into two stages: aligning model preference and optimizing the reward layer. Through introducing an additional constraint of policy loss, our hybrid alignment framework supervises the internal preference model at the token level while simultaneously optimizing the mapping layer at the sequence level, significantly improving the training effectiveness. We theoretically verify the validity of our method and demonstrate its reliability through systematic experiments.

Our method allows for a consistent customization of the reward model. In the future, we will thoroughly explore the potential of the reward model and its variants across various tasks, and investigate whether the logistic distribution is the optimal prior for reward modeling.



553  
554  
555  
556  
557  
558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604

## Impact Statements

This paper presents work whose goal may benefit the training of large language models in the field of deep learning. Among the many possible consequences, we do not believe that there is a significant possibility of adverse effects on society.

## Limitations

In this paper, we discuss the potential of enhancing the alignment process of reward models by incorporating policy constraints, where the policy loss functions similarly to a regularization loss, acting as an auxiliary function to guide model training. However, since DPO can be directly used to train an implicit reward model, replacing the reward model with a DPO model for downstream tasks can also be a feasible approach, while we do not explore methods for combining the outputs of the policy layer and the reward layer, which remains a direction for our future research.

## References

Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. 2023. [A general theoretical paradigm to understand learning from human preferences](#). *ArXiv*, abs/2310.12036.

Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, T. J. Henighan, Nicholas Joseph, Saurav Kadavath, John Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Christopher Olah, Benjamin Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *ArXiv*, abs/2204.05862.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, T. J. Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *ArXiv*, abs/2005.14165.

Meng Cao, Lei Shu, Lei Yu, Yun Zhu, Nevan Wichers, Yinxiao Liu, and Lei Meng. 2024. [Drlc: Reinforce-](#)

[ment learning with dense rewards from llm critic](#). *Preprint*, arXiv:2401.07382. 605  
606

Paul Francis Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). *ArXiv*, abs/1706.03741. 607  
608  
609  
610

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [Bert: Pre-training of deep bidirectional transformers for language understanding](#). In *North American Chapter of the Association for Computational Linguistics*. 611  
612  
613  
614  
615

Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and T. Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#). *ArXiv*, abs/2304.06767. 616  
617  
618  
619  
620

Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#). *ArXiv*, abs/2305.14387. 621  
622  
623  
624  
625

Leo Gao, John Schulman, and Jacob Hilton. 2023. [Scaling laws for reward model overoptimization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10835–10866. PMLR. 626  
627  
628  
629  
630  
631

Dongyoung Go, Tomasz Korbak, Germán Kruszewski, Jos Rozen, and Marc Dymetman. 2023. [Compositional preference models for aligning lms](#). *Preprint*, arXiv:2310.13011. 632  
633  
634  
635

Joey Hong, Kush Bhatia, and Anca D. Dragan. 2022. [On the sensitivity of reward inference to misspecified human models](#). *ArXiv*, abs/2212.04717. 636  
637  
638

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*. 639  
640  
641  
642  
643

Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. [Personalized soups: Personalized large language model alignment via post-hoc parameter merging](#). *Preprint*, arXiv:2310.11564. 644  
645  
646  
647  
648  
649

Jiaming Ji, Mickel Liu, Juntao Dai, Xuehai Pan, Chi Zhang, Ce Bian, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. 2023. [Beavertails: Towards improved safety alignment of llm via a human-preference dataset](#). *ArXiv*, abs/2307.04657. 650  
651  
652  
653  
654

Yuu Jinnai, Tetsuro Morimura, Kaito Ariu, and Kenshi Abe. 2024. [Regularized best-of-n sampling to mitigate reward hacking for language model alignment](#). *Preprint*, arXiv:2404.01054. 655  
656  
657  
658

659	W. B. Knox, Stephane Hatgis-Kessell, Serena Booth, Scott Niekum, Peter Stone, and Alessandro Allievi. 2022. <a href="#">Models of human preference for learning reward functions</a> . <i>ArXiv</i> , abs/2206.02231.	715
660		716
661		717
662		
663	Andreas Kopf, Yannic Kilcher, Dimitri von Rutte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Rich'ard Nagyfi, ES Shahul, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. <a href="#">Openassistant conversations - democratizing large language model alignment</a> . <i>ArXiv</i> , abs/2304.07327.	718
664		719
665		720
666		721
667		
668		722
669		723
670		724
671		725
		726
672	Yuhang Lai, Siyuan Wang, Shujun Liu, Xuanjing Huang, and Zhongyu Wei. 2024. <a href="#">Alarm: Align language models via hierarchical rewards modeling</a> . <i>Preprint</i> , arXiv:2403.06754.	727
673		
674		
675		
676	Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. <a href="#">Rewardbench: Evaluating reward models for language modeling</a> . <i>Preprint</i> , arXiv:2403.13787.	728
677		729
678		730
679		
680		731
681		732
		733
		734
		735
682	Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. <a href="#">Retrieval-augmented generation for knowledge-intensive nlp tasks</a> . <i>ArXiv</i> , abs/2005.11401.	736
683		737
684		
685		
686		
687		
688	Lei Li, Yekun Chai, Shuohuan Wang, Yu Sun, Hao Tian, Ningyu Zhang, and Hua Wu. 2023a. <a href="#">Tool-augmented reward modeling</a> . <i>ArXiv</i> , abs/2310.01045.	738
689		739
690		740
		741
		742
		743
691	Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023b. <a href="#">AlpacaEval: An automatic evaluator of instruction-following models</a> . <a href="https://github.com/tatsu-lab/alpaca_eval">https://github.com/tatsu-lab/alpaca_eval</a> .	744
692		745
693		746
694		747
695		748
		749
696	Lev McKinney, Yawen Duan, David Krueger, and Adam Gleave. 2023. <a href="#">On the fragility of learned reward functions</a> . <i>ArXiv</i> , abs/2301.03652.	750
697		751
698		752
		753
699	Sidharth Mudgal, Jong Lee, Harish Ganapathy, YaGuang Li, Tao Wang, Yanping Huang, Zhifeng Chen, Heng-Tze Cheng, Michael Collins, Trevor Strohman, Jilin Chen, Alex Beutel, and Ahmad Beirami. 2024. <a href="#">Controlled decoding from language models</a> . <i>Preprint</i> , arXiv:2310.17022.	754
700		755
701		756
702		757
703		758
704		759
		760
705	Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. <a href="#">Generative agents: Interactive simulators of human behavior</a> . <i>Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology</i> .	761
706		762
707		763
708		764
709		765
710		766
711	Silviu Pitis. 2023. <a href="#">Failure modes of learning reward models for LLMs and other sequence models</a> . In <i>ICML 2023 Workshop The Many Facets of Preference-Based Learning</i> .	767
712		768
713		769
714		770
	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. <a href="#">Language models are unsupervised multitask learners</a> .	715
		716
		717
	Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. 2024. <a href="#">From r to q*: Your language model is secretly a q-function</a> . <i>arXiv preprint arXiv:2404.12358</i> .	718
		719
		720
		721
	Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. <a href="#">Direct preference optimization: Your language model is secretly a reward model</a> . <i>ArXiv</i> , abs/2305.18290.	722
		723
		724
		725
		726
	Stuart Russell. 2014. <a href="#">Of myths and moonshine</a> . online.	727
	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. <a href="#">Proximal policy optimization algorithms</a> . <i>ArXiv</i> , abs/1707.06347.	728
		729
		730
	Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. <a href="#">Learning to summarize with human feedback</a> . In <i>Advances in Neural Information Processing Systems</i> , volume 33, pages 3008–3021. Curran Associates, Inc.	731
		732
		733
		734
		735
		736
		737
	Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liangyan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. 2023. <a href="#">Aligning large multimodal models with factually augmented rlhf</a> . <i>ArXiv</i> , abs/2309.14525.	738
		739
		740
		741
		742
		743
	Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. <a href="#">Llama 2: Open foundation and fine-tuned chat models</a> . <i>ArXiv</i> , abs/2307.09288.	744
		745
		746
		747
		748
		749
		750
		751
		752
		753
		754
		755
		756
		757
		758
		759
		760
		761
		762
		763
		764
		765
		766
	Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. <a href="#">Attention is all you need</a> . In <i>Neural Information Processing Systems</i> .	767
		768
		769
		770

771	Bing Wang, Rui Zheng, Luyao Chen, Yan Liu, Shihan Dou, Caishuang Huang, Wei Shen, Senjie Jin, Enyu Zhou, Chenyu Shi, Songyang Gao, Nuo Xu, Yuhao Zhou, Xiaoran Fan, Zhiheng Xi, Jun Zhao, Xiao Wang, Tao Ji, Hang Yan, Lixing Shen, Zhan Chen, Tao Gui, Qi Zhang, Xipeng Qiu, Xuanjing Huang, Zuxuan Wu, and Yuanyuan Jiang. 2024. <a href="#">Secrets of rlhf in large language models part ii: Reward modeling</a> . <i>ArXiv</i> , abs/2401.06080.	
780	Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. <a href="#">Large language models are not fair evaluators</a> . <i>Preprint</i> , arXiv:2305.17926.	
784	Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. <a href="#">Self-instruct: Aligning language models with self-generated instructions</a> . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	
790	Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021. <a href="#">Finetuned language models are zero-shot learners</a> . <i>ArXiv</i> , abs/2109.01652.	
794	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. <a href="#">Chain of thought prompting elicits reasoning in large language models</a> . <i>ArXiv</i> , abs/2201.11903.	
799	Zejiu Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. <a href="#">Fine-grained human feedback gives better rewards for language model training</a> . In <i>Advances in Neural Information Processing Systems</i> , volume 36, pages 59008–59033. Curran Associates, Inc.	
806	Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Feiran Huang. 2023. <a href="#">Rrhf: Rank responses to align language models with human feedback without tears</a> . <i>ArXiv</i> , abs/2304.05302.	
810	Shengbin Yue, Wei Chen, Siyuan Wang, Bingxuan Li, Chenchen Shen, Shujun Liu, Yuxuan Zhou, Yao Xiao, Song Yun, Wei Lin, et al. 2023. <a href="#">Disc-lawllm: Fine-tuning large language models for intelligent legal services</a> . <i>arXiv preprint arXiv:2309.11325</i> .	
815	Yao Zhao, Misha Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J. Liu. 2022. <a href="#">Calibrating sequence likelihood improves conditional language generation</a> . <i>ArXiv</i> , abs/2210.00045.	
819	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Haoteng Zhang, Joseph Gonzalez, and Ion Stoica. 2023. <a href="#">Judging llm-as-a-judge with mt-bench and chatbot arena</a> . <i>ArXiv</i> , abs/2306.05685.	
825	Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, and Jiantao Jiao. 2023. <a href="#">Starling-7b: Improving llm helpfulness &amp; harmlessness with rlaiif</a> .	
	<b>A Experiments Setup</b>	828
	Our default setup is shown in Table 6.	829
	To train the reward model, we use DPO Loss as the policy loss in HAF and set policy ratio $\alpha = 0.2$ . The learning rate is $1.0 \times 10^{-5}$ for Phi-2 and Mistral-lora-baseline, $3.0 \times 10^{-5}$ for Mistral-lora-HAF. A single RTX A6000 with 48GB memory is used for training the reward model. The model used for testing is the checkpoint that achieves the highest reward on the validation set.	830
	For PPO training in Section 5.2.2, we utilize two RTX A6000 GPUs for parallel training with a total batch size of 4. The maximum number of new tokens generated is set to 128, and the learning rate is 1e-6. The training is conducted over a maximum of 20,000 episodes. We employ score scaling and score normalization and clip the scores between -3 and 3. All other settings follow the implementation in the TRL library. The model used for testing is the checkpoint that achieves the highest reward on the validation set. The generation config includes $top_p = 0.8$ , $temperature = 0.5$ , $length_penalty = 1.3$ , $repetition_penalty = 1.2$ , $do_sample = True$	831
		832
		833
		834
		835
		836
		837
		838
		839
		840
		841
		842
		843
		844
		845
		846
		847
		848
		849
		850
		851
	<b>B Discussions for Policy Loss Ratio</b>	852
	Figure 4 reveals that incorporating even a mere 0.1x of policy loss can significantly impact the results. Using reward loss alone leads to slow training; to achieve the same loss value, the model with policy loss requires only a fraction of the time. However, this rapid training characteristic also accelerates overfitting, necessitating the use of early stopping strategies to halt training in time. When the policy loss ratio is negative, model performance deteriorates, and the variations in various metrics resemble those of the baseline. This indicates a correlation between the policy model and the reward model.	853
		854
		855
		856
		857
		858
		859
		860
		861
		862
		863
		864
	<b>C Loss Functions</b>	865
	<b>C.1 Deriving the Reward Loss Functions</b>	866
	In practice, there is no access to the ground truth reward of a response, so it is not applicable to solve the reward regression problem by directly optimizing the discrepancy between every predicted reward and the true reward. The Bradley-Terry model is introduced here to construct a solvable classification problem with one additional assumption – if one response is better than the other, then it wins with the probability of 100%. For a query	867
		868
		869
		870
		871
		872
		873
		874
		875

setup	value	setup	value	setup	value
<b>lora rank</b>	64	<b>optimizer</b>	AdamW	<b>precision</b>	bf16
<b>lora alpha</b>	16	<b>adam_beta1</b>	0.9	<b>max gradient norm</b>	1.0
<b>training steps</b>	3200	<b>adam_beta2</b>	0.999	<b>max sequence length</b>	512
<b>evaluation steps</b>	0.025	<b>weight_decay</b>	0.0	<b>global random seed</b>	0
<b>batch size</b>	16	<b>adam_epsilon</b>	1e-5	<b>framework</b>	PyTorch

Table 6: Default setup




Figure 4: Results for different policy ratios. “margin” is the average difference between a better and worse response’s rewards. A policy ratio of 0 equals to Baseline method.

876  $x$ , a preferred response  $y$  and a dispreferred re- 888  
877 sponse  $y'$ , the predicted winning probability is 889  
878  $P(y \succ y') = \sigma(\mathbf{r}(x, y) - \mathbf{r}(x, y'))$ , and the ground 890  
879 truth  $P^*(y \succ y') = \sigma(\mathbf{r}^*(x, y) - \mathbf{r}^*(x, y')) = 1$ , 891  
880 so the standard reward loss is essentially a cross- 892  
881 entropy loss of the predicted winning probability.

$$\begin{aligned}
\mathcal{L} &= -P^*(y \succ y') \log P(y \succ y') \\
&\quad - P^*(y' \succ y) \log P(y' \succ y) \\
&= -\log \sigma(\mathbf{r}(x, y) - \mathbf{r}(x, y'))
\end{aligned}$$

882  
883 The optimal model  $F^*$  and  $\phi^*$  are secretly hidden 894  
884 in the coefficient “1”.

## 885 C.2 DPO as the Policy Loss

886 The derivation for policy loss is the same as re- 898  
887 ward loss in their essence. The policy model can 899

888 be treated as a reward model with sequence proba- 889  
889 bilities reflecting the rewards (Rafailov et al., 2023, 890  
890 2024).  $reward(x, y) = \pi(x, y) / \pi_{ref}(x, y)$ . With 891  
891 the Bradley-Terry model and the assumption of 892  
892  $P(y \succ y') = 1$ , DPO is also a legal loss function.

893 From this perspective, the DPO loss and reward 894  
894 loss share the same assumption of  $P(y \succ y') = 1$ . 895  
895 The reward model and the DPO-trained policy 896  
896 model are essentially doing the same task despite 897  
897 some formal differences (Rafailov et al., 2023, 898  
898 2024). This may provide insight into why DPO 899  
899 is the most suitable among all policy losses.

## D Mathematical Derivations

### D.1 Inequality Scaling

$$\begin{aligned}
& \min_{F, \phi, K} \mathbb{E}_{d \sim \mathcal{P}} [D_1(F \circ \phi(d), F^* \circ \phi^*(d)) \\
& \quad + \alpha \cdot D_2(K \circ \phi(d), K^* \circ \phi^*(d))] \\
& \leq \min_{\substack{F=F_s \\ \phi=\phi_s \\ K}} \mathbb{E}_{d \sim \mathcal{P}} [D_1(F \circ \phi(d), F^* \circ \phi^*(d)) \\
& \quad + \alpha \cdot \mathcal{L}_2(K \circ \phi(d), K^* \circ \phi^*(d))] \\
& = \min_K \mathbb{E}_{d \sim \mathcal{P}} [\alpha \cdot D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))] \\
& \quad + \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_s \circ \phi_s(d), F^* \circ \phi^*(d))]
\end{aligned}$$

With the definition of  $\phi_H, K_H, F_H$ , we have:

$$\begin{aligned}
& \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_H \circ \phi_H(d), F^* \circ \phi^*(d)) \\
& \quad + \alpha \cdot D_2(K_H \circ \phi_H(d), K^* \circ \phi^*(d))] \\
& \leq \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_s \circ \phi_s(d), F^* \circ \phi^*(d))] \\
& \quad + \min_K \mathbb{E}_{d \sim \mathcal{P}} [\alpha \cdot D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))] \\
& \leq \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_H \circ \phi_H(d), F^* \circ \phi^*(d))] \\
& \quad + \min_K \mathbb{E}_{d \sim \mathcal{P}} [\alpha \cdot D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))]
\end{aligned}$$

In practical settings, “ $\leq$ ”s do not hold at the same time (simultaneously optimizing two objectives is preferable to optimizing them sequentially). With the premise that the model is fully optimized with the hybrid alignment loss for any  $\alpha \in (0.1, 2)$ , which means both of the objectives have an impact on the final optimization result, namely  $\phi_H \neq \phi_s$ , there exists a little gap  $\epsilon > 0$  such that

$$\begin{aligned}
& \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_H \circ \phi_H(d), F^* \circ \phi^*(d)) \\
& \quad + \alpha \cdot D_2(K_H \circ \phi_H(d), K^* \circ \phi^*(d))] \\
& \leq \mathbb{E}_{d \sim \mathcal{P}} [D_1(F_H \circ \phi_H(d), F^* \circ \phi^*(d))] \\
& \quad + \min_K \mathbb{E}_{d \sim \mathcal{P}} [\alpha \cdot D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))] - \epsilon
\end{aligned}$$

Then, there goes

$$\begin{aligned}
& \mathbb{E}_{d \sim \mathcal{P}} [D_2(K_H \circ \phi_H(d), K^* \circ \phi^*(d))] \\
& \leq \min_K \mathbb{E}_{d \sim \mathcal{P}} [D_2(K \circ \phi_s(d), K^* \circ \phi^*(d))] - \frac{\epsilon}{\alpha}
\end{aligned}$$

Here we get Prop. 1.

### D.2 Derive the Final Inequality with the 3 Properties

#### Convergence:

Since the trained model  $K \circ \phi$  is close to  $K^* \circ \phi^*$ , we can therefore linearize  $D_2$  with a certain positive number  $k$ :

$$\begin{aligned}
& \mathbb{E}_{d \sim \mathcal{P}} [D_2(K \circ \phi(d), K^* \circ \phi^*(d))] \\
& = \mathbb{E}_{d \sim \mathcal{P}} k |K \circ \phi(d) - K^* \circ \phi^*(d)|
\end{aligned} \tag{7}$$

Separating little disturbance:

$$\mathbb{E}_{d \sim \mathcal{P}} |N \circ \phi(d)| < \delta \tag{8}$$

holds for any fully-optimized model  $K \circ \phi$  with  $N := K - K^*$ . Given that the trained model and its preferences closely approximate those of the true model and preferences, we are able to scale down the error terms by a small margin.

**Gradient scaling:**

Intuitively, the optimal model is unique, so  $\mathbb{E}_{d \sim \mathcal{P}} |K^* \circ \phi(d) - K^* \circ \phi^*(d)| > 0$ . Here we make a slightly stronger assumption that  $K^*$  is locally  $g_{max}$ -Lipschitz continuous and has the lower bound  $g_{min}$ , which means for any  $\phi$  that is close to  $\phi^*$ , there exists

$$\begin{aligned}
& g_{min} \mathbb{E}_{d \sim \mathcal{P}} \|\phi(d) - \phi^*(d)\| \\
& < \mathbb{E}_{d \sim \mathcal{P}} |K^* \circ \phi(d) - K^* \circ \phi^*(d)| \\
& < g_{max} \mathbb{E}_{d \sim \mathcal{P}} \|\phi(d) - \phi^*(d)\|
\end{aligned} \tag{9}$$

Based on these three properties, we can derive the result from Prop. 1.

**Prop. 1**

$$\begin{aligned}
& \xrightarrow{\text{Eq. 7}} \mathbb{E}_{d \sim \mathcal{P}} |K_H \circ \phi_H(d) - K^* \circ \phi^*(d)| \\
& \leq \min_K \mathbb{E}_{d \sim \mathcal{P}} |K \circ \phi_s(d) - K^* \circ \phi^*(d)| - \frac{\epsilon}{\alpha \cdot k} \\
& \xrightarrow{\text{Ineq. 8}} \mathbb{E}_{d \sim \mathcal{P}} |K^* \circ \phi_H(d) - K^* \circ \phi^*(d)| - \delta \\
& < \mathbb{E}_{d \sim \mathcal{P}} |K^* \circ \phi_s(d) - K^* \circ \phi^*(d)| + \delta - \frac{\epsilon}{\alpha \cdot k} \\
& \xrightarrow{\text{Ineq. 9}} \\
& g_{min} \mathbb{E}_{d \sim \mathcal{P}} [|\|\phi_H(d) - \phi^*(d)\| - \|\phi_s(d) - \phi^*(d)\||] \\
& < (g_{max} - g_{min}) \mathbb{E}_{d \sim \mathcal{P}} \|\phi_s(d) - \phi^*(d)\| \\
& \quad + 2\delta - \frac{\epsilon}{\alpha \cdot k}
\end{aligned}$$

which is Proposition 2.

## E GPT Judgement

**Comparing two responses** The prompt we used for judgement is listed in Table 8. The sentence between “<SYSTEM PROMPT>” is the system prompt, and the others are the user prompt. “{question}”, “{response 1}”, “{response 2}” will be replaced with the actual query or responses respectively. As GPT does not exhibit a strong “positional bias” (Wang et al., 2023), so we just randomly interchange the order of the two responses rather than prompting twice with the responses swapped.

**Ranking responses** Table 7 shows the consumption approximation for getting top-1, top-2 responses and the complete order out of 4/8 responses. We consider that performing a single sorting operation on eight responses with the model may result in a loss of precision. Besides, while binary comparisons exhibit high accuracy, repeated binary comparisons inevitably lead to cumulative errors and erroneous outcomes. Therefore, whether from a cost or accuracy standpoint, it is not a favorable option. In practice, we obtain the top 2 responses by ranking 4 responses with GPT-3.5-turbo at once. For 8 candidate responses, we first evenly divide them into two groups and use GPT to rank the responses of each group, then we rank the two sets of the top 2 responses to get the top 2 responses among 8 candidates.




Figure 5: Three times of interactions with GPT to get top-2 responses

The prompt for ranking four responses is shown in Table 9. GPT’s answer will be parsed in JSON format.

# responses	Top-1		Top-2		Complete sort	
	4	8	4	8	4	8
binary comparison	<b>6</b> <sub>3×2</sub>	<b>14</b> <sub>7×2</sub>	<b>8</b> <sub>4×2</sub>	<b>20</b> <sub>10×2</sub>	<b>10</b> <sub>5×2</sub>	<b>32</b> <sub>16×2</sub>
rank 4 responses	<b>4</b> <sub>1×4</sub>	<b>12</b> <sub>3×4</sub>	<b>4</b> <sub>1×4</sub>	<b>12</b> <sub>3×4</sub>	<b>4</b> <sub>1×4</sub>	<b>20</b> <sub>5×4</sub>
rank 8 responses	<b>4</b> <sub>1×4</sub>	<b>8</b> <sub>1×8</sub>	<b>4</b> <sub>1×4</sub>	<b>8</b> <sub>1×8</sub>	<b>4</b> <sub>1×4</sub>	<b>8</b> <sub>1×8</sub>

Table 7: Approximation for resources consumption. The first column is three different ways of interacting with GPT. The first row is the target response(s) and the second row is the number of candidate responses. “ $a \times b$ ” means we should engage with GPT-3.5 a total of  $a$  times, with each interaction requiring an input of  $b$  responses. For example, “**6**<sub>3×2</sub>” means when using binary comparison, to get the top-1 response among 4 candidate responses, we need 3 turns of interactions with each turn requiring an input of 2 responses, hence our expenditure amounts to approximately 6 units

Prompt for comparing two responses.

<SYSTEM PROMPT>You are a helpful instruction-following assistant that prints the best model by selecting the best outputs for a given instruction.<SYSTEM PROMPT>  
 Select the output (a) or (b) that best matches the given instruction. Choose your preferred output, which can be subjective. Your answer should ONLY contain: Output (a) or Output (b).  
 Here’s an example:

# Example:

## Instruction:

Give a description of the following job: "ophthalmologist"

## Output (a):

An ophthalmologist is a medical doctor who pokes and prods at your eyes while asking you to read letters from a chart.

## Output (b):

An ophthalmologist is a medical doctor who specializes in the diagnosis and treatment of eye diseases and conditions.

## Which is best, Output (a) or Output (b)?

Output (b)

Here the answer is Output (b) because it provides a comprehensive and accurate description of the job of an ophthalmologist. In contrast, output (a) is more of a joke.

# Task:

Now is the real task, do not explain your answer, just say Output (a) or Output (b).

## Instruction:

{question}

## Output (a):

{response 1}

## Output (b):

{response 2}

## Which is best, Output (a) or Output (b)?

Table 8: We use 1-shot for response comparison.

Prompt for ranking four responses.

<SYSTEM PROMPT>You are a helpful assistant, that ranks models by the quality of their answers<SYSTEM PROMPT>

I want you to create a leaderboard of different models. To do so, I will give you the instructions (prompts) given to the models, and the responses of four models. Please rank the models based on which responses would be preferred by humans. All inputs and outputs should be python dictionaries.

Here is the prompt:

```
{
  "instruction": {question},
}
```

Here are the outputs of the models:

```
[
  {
    "model": "model_1",
    "answer": {output_1}
  },
  {
    "model": "model_2",
    "answer": {output_2}
  },
  {
    "model": "model_3",
    "answer": {output_3}
  },
  {
    "model": "model_4",
    "answer": {output_4}
  }
]
```

Now please rank the models by the quality of their answers, so that the model with rank 1 has the best output. Then return a list of the model names and ranks, i.e., produce the following output:

```
[
  {"model": "model_1", "rank": <model-rank>},
  {"model": "model_2", "rank": <model-rank>},
  {"model": "model_3", "rank": <model-rank>},
  {"model": "model_4", "rank": <model-rank>}
]
```

Your response must be a valid Python dictionary and should contain nothing else because we will directly execute it in Python. Please provide the ranking that the majority of humans would give.

Table 9: We rank four responses in order of quality in a single interaction.

	AHP		BS		CA		Helpful		Harmless	
	Top-1	Top-2	Top-1	Top-2	Top-1	Top-2	Top-1	Top-2	Top-1	Top-2
Phi-2 <sub>HAF</sub>	<b>28.68</b>	<b>52.51</b>	<b>32.69</b>	<b>53.35</b>	<b>37.52</b>	<b>66.21</b>	<b>45.44</b>	<b>74.26</b>	<b>24.52</b>	<b>45.15</b>
Phi-2 <sub>baseline</sub>	15.46	34.64	29.28	49.72	27.83	51.68	43.62	73.92	17.22	37.29
Mistral <sub>HAF</sub>	<b>17.42</b>	<b>31.22</b>	<b>9.94</b>	<b>17.70</b>	16.00	28.81	<b>13.68</b>	<b>27.57</b>	<b>9.50</b>	<b>21.07</b>
Mistral <sub>baseline</sub>	10.97	23.87	7.45	17.08	<b>17.99</b>	<b>32.78</b>	12.68	26.36	8.68	17.36

Table 10: Top-k recall for best-of-N sampling on each dataset. The results are presented as the percentage of the chosen responses included in top-k responses.