

Lost in the Context: Insufficient and Distracted Attention to Contexts in Preference Modeling

Anonymous ACL submission

Abstract

In Reinforcement Learning from Human Feedback (RLHF), the reward model (RM) evaluates the response quality based on the given context and assigns a reward. It plays a crucial role in aligning RLHF with human preferences. Although the current RM training paradigm concatenates the context and response while amplifying the reward difference between good and bad response pairs, we demonstrate that the RM faces two significant issues: i) it often allocates only a small proportion of attention to the context, and ii) it frequently ignores segments of the context that are relevant for evaluating the response quality. These issues undermine the RM’s effectiveness in modeling human preferences. To further address these challenges, we propose AttnRM, a novel optimization framework that enables the RM to concentrate on crucial segments of the context. Experimental results demonstrate that AttnRM significantly improves preference modeling by increasing attention to relevant information within the context. It also enhances the RM’s generalizability and achieves better performance in aligning with human preferences¹.

1 Introduction

Reinforcement learning from human feedback (RLHF) provides a key technique to ensure that the behavior of AI systems aligns with the intentions of their designers and the expectations of users (Bai et al., 2022; Ouyang et al., 2022). In RLHF, the reward model (RM) evaluates and scores the quality of responses based on a given context. Then the produced reward serves as the supervisory signal to fine-tune the policy model, guiding it to generate responses that yield higher rewards. Consequently, the effectiveness and accuracy of the RM are crucial in RLHF (Eschmann, 2021; OpenAI, 2023). In practice, current popular reward modeling ap-

¹Our code is available at [here](#).




Figure 1: Reward model tends to rely more on the responses and does not adequately attend to the context. It also ignores segments in the context that are crucial for evaluating response quality, and over-attends to irrelevant parts such as special tokens.

proaches concatenate contexts and responses, optimizing the RM to assign higher reward values to good responses compared to bad ones (Schulman et al., 2017; Wang et al., 2024b). However, the RM still faces two significant challenges: i) it tends to rely on the responses and does not adequately attend to the context, and ii) it ignores the context’s segments that are relevant to evaluate the quality of the response.

Firstly, when the reward model assigns a reward to a context-response pair, Figure 1 illustrates attention scores assigned to different parts of the context. Results show that the RM allocates only a small proportion of attention to the context. Moreover, as shown at the bottom of Figure 1, even when we modify key segments in the context, making the originally good response no longer appropriate, the RM still assigns a high reward. These findings indicate that when the RM assigns rewards, it pays insufficient attention to the context, and ignores those context fragments that are crucial for truly




Figure 2: We train a reward model on the HH-RLHF training set and analyze it on the validation set. **(left)** Normalized attention score to context-response pair. The reward model allocates a significant amount of attention to special tokens and some attention to the response, but only a small amount to the context. **(right)** Gradually masking context from start to end, RM accuracy in evaluating response pairs with the given masked context. We observe that after masking 90% of the context tokens, the accuracy of RM only decreased by 3.52%.

determining the quality of the response. This affects the generalization capability of the reward model and its consistency with modeling human preferences. The experiments in Section 2 further substantiate that current reward models struggle with challenges of insufficient and distracted attention to the context.

To address these challenges, we propose AttnRM, a novel optimization framework for RM training, which enables the RM to identify and concentrate on crucial segments of the context. Specifically, AttnRM first segments the context and utilizes large language models (LLM) as a verifier to verify which context fragments contain information relevant to evaluate the response (*i.e.*, relevant fragment). After identifying the relevant context segments, AttnRM calculates the softmax attention allocated by the reward model to these segments. It then maximizes the average attention value, while ensuring a balanced distribution of attention across tokens in segments. Our approach reaffirms the fundamental principle of preference modeling, which ensures that the reward model assigns rewards to responses with appropriate attention to the relevant parts of the preceding context, while reducing distractions from irrelevant parts.

To evaluate the effectiveness of AttnRM, we conduct extensive experiments on two common tasks, *i.e.*, the general dialogue task and the summarization task. Experimental results demonstrate that AttnRM enhances the RM’s effectiveness in human preference alignment by increasing attention to key information within the context. Additionally, AttnRM also improves the RM’s performance in out-of-distribution (OOD) scenarios, providing

more stable and generalized reward signals in preference alignment. In summary, our paper makes the following contributions:

- We demonstrate that the current RM faces two critical challenges: I) it often allocates only a small proportion of attention to the context, and ii) it frequently ignores response-related information in the context.
- We introduce AttnRM, a novel optimization framework for RM training, which enables the RM to concentrate on crucial segments in the context, while being less distracted by irrelevant segments.
- Experimental results show that AttnRM can provide more stable and generalized reward signals, and significantly enhance the performance of human preference alignment.

2 Motivation

In this section, we present evidence to demonstrate that current popular RM training approaches, when assigning rewards, pay only a small proportion of attention to contexts and also pay less attention to response-related context segments.

2.1 Insufficient Attention to Contexts

The reward model scores responses based on the context, analogous to the concept of conditional probabilities. Formally, let the context be denoted as x , which includes the historical conversations and the current query, and let the response be represented by y . The reward model can then be expressed as $r(y|x)$.




Figure 3: Normalized attention scores to context for preference pairs that the RM evaluates correctly (**left**) versus incorrectly (**right**). The reward model is trained on the HH-RLHF dataset and the preference pairs for analysis are sourced for an OOD reward model benchmark *i.e.*, RMB. We observe that the reward model assigns more attention to context when it makes accurate evaluations.

The prevailing training paradigm for reward models concatenates the context and the response, and then maximizes the reward differential between a good context-response pair and a bad context-response pair. Formally, assuming y_{win} denotes a better response compared to y_{lose} , the training objective can be expressed as: $\max [r(x \oplus y_{\text{win}}) - r(x \oplus y_{\text{lose}})]$. Essentially, it aims to model human preferences, specifically how well humans think a response fits within a given context. This leads to an intriguing question: *how much attention do reward models actually pay to the context when assigning rewards?*

We train a reward model on the training set of the Anthropic’s HH-RLHF dataset (Bai et al., 2022) using current popular methods and visualize the model’s average attention values assigned to the context on the validation set. The experimental results are presented on the left side of Figure 2, where the contexts are divided into five equal-length segments, and we calculate the proportion of attention scores for each segment. The results show that, despite the training method concatenating the context and response, the reward model still allocates only a small portion of attention to the context when it assigns scores to the responses.

The right side of Figure 2 further illustrates that the reward model pays little attention to the context when scoring. Specifically, we progressively discard the context from the beginning, providing only the remaining part of the context along with the response to the reward model. Our findings indicate that even after removing 90% of the context tokens, the accuracy of the RM only decreased from 72.0% to 68.5%. This suggests that the RM has learned a shortcut during training: it often relies predom-

inantly on the response to assign reward values. This behavior is inconsistent with real-world human preferences, as we should not take things out of context. Moreover, this reliance on the response can lead to poor generalization capabilities of the reward model.

In summary, the reward model pays insufficient attention to the context when assigning rewards to the quality of responses.

2.2 Distracted Attention to Contexts

We conduct a case study from HH-RLHF dataset (Bai et al., 2022) to further explore issues that the current RM’s insufficient attention to the context, as shown in Appendix A.1. In these cases, the responses initially meet the context requirements, and the RM correctly assigns high reward values. We then manually modify the context to make the original responses less suitable for needs or unable to address problems within the context (*i.e.*, good responses are made worse). However, we find that even after modifying the context, the altered responses still receive relatively high rewards, as detailed in Appendix A.1. This result indicates that the RM not only allocates too little attention to the context but also fails to accurately capture the context information that is relevant for evaluating the quality of the responses.

Does the reward model’s performance relate to its attention to the context? We investigate this issue using an out-of-distribution (OOD) reward model benchmark, RMB (Zhou et al., 2025). We divide the context into ten equal-length segments and identify the positions of relevant context segments, as detailed in Appendix A.3. Then, we visualize the attention distribution assigned by the RM to context, as shown in Figure 3. The left side




Figure 4: Distribution of relevant segments’ positions in the context of the HH-RLHF training and validation sets. Relevant segments are generally distributed throughout the context, with a tendency to be towards the end. This phenomenon, which is normal in the real world, introduces a catastrophic bias into preference modeling.

of the figure shows attention on samples where the RM makes correct judgments, while the right side shows the distribution on samples where the RM makes incorrect judgments. We observe a significant difference in the model’s attention to the context between these two parts. When the RM performs correctly, it assigns more attention to the context. The experimental results indicate that the reward model’s attention to context is related to its performance in preference modeling. Increasing this attention may enhance the model’s generalization capability.

We further investigate the possible reasons for the RM’s insufficient attention to the context. Specifically, we analyze the distribution of relevant segments’ positions in the context of the HH-RLHF training and validation sets, as shown in Figure 4. The experimental results reveal the following: 1) A high proportion of information relevant to evaluating response quality appears at the end of the context. 2) Relevant information also appears at other positions within the context to some extent.

This observation aligns with real-world scenarios, where information related to the response often appears toward the end of the context. However, under the current training paradigm, the RM may be significantly influenced by this normal phenomenon, resulting in a lack of attention to the context. To increase the RM’s attention to the context, constructing a balanced dataset that evenly distributes relevant information is resource-intensive and, in many cases, impractical. Therefore, there is an urgent need for developing a training framework in reward modeling that can utilize existing training datasets efficiently, while ensuring that the model assigns adequate attention to the context.

3 Method

In this section, we first introduce the current popular RM training objective, and then introduce AttnRM, a novel training framework designed to mitigate the issue of the RM neglecting context when evaluating the quality of responses.

Broadly, the reward model consists of a decoder-only transformer and a linear head. The transformer takes the context x and the response y , and then feeds the hidden state of the last special token into the linear head to obtain the reward. Formally, let the input sequence of the transformer be $z = (\text{<bos>} \oplus x \oplus \text{<eos>} \oplus y \oplus \text{<eos>})$, where <bos> and <eos> denote the special tokens. $r_\theta(x, y)$ denotes the reward model with parameters θ , which can be written as follows:

$$r_\theta(x, y) = h(T_{\text{eos}}(z)), \quad (1)$$

where h and T denote the head linear and the transformer model. The current training objective of RM aims to amplify the reward difference between two responses under the same contexts to ensure that better responses are assigned higher reward values. For a given context x and two responses generated by a supervised fine-tuning (SFT) model, we assume the response y_{win} is better than another response y_{lose} , *i.e.*, $y_{\text{win}} \succ y_{\text{lose}}$. The training objective can be simplified by maximizing follows:

$$\mathcal{L}(x, y_{\text{win}}, y_{\text{lose}}) = \log \sigma(r_\theta(x, y_{\text{win}}) - r_\theta(x, y_{\text{lose}})), \quad (2)$$

where σ is the logistic sigmoid function. However, in this training objective, the RM can not assign proper attention to contexts, as elaborated in Section 2. The goal of AttnRM is to enable the RM to identify the context segments relevant to evaluate the response quality, and concentrate on them.

Definition 3.1. Let Q and K denote the query and key matrices in the l -th attention layer of the transformer model T . The attention score of i -th token assigned by the reward model can be calculated as follows:

$$\mathcal{A}_i^l(x, y) = \frac{Q_{\text{eos}} \cdot K_i}{\sqrt{d_k}}, \quad (3)$$

where d_k denotes the dimension of the key used to scale the inner product. Then, we define the average strength of attention for all tokens in the context:

$$\mathcal{Z}(x, y; c) = \frac{1}{|c|} \sum_i \frac{1}{L} \sum_l \mathcal{A}_i^l(x, y), \quad (4)$$

where c denotes the set of tokens in the context and L denotes the block number of the transformers. We expect that the reward model should assign more attention to the context when evaluating the responses, *i.e.*, maximize $E_{(x,y) \sim \mathcal{H}} \mathcal{Z}(x, y; c)$, where $\mathcal{H}$ denotes the distribution of the context-response pairs in the training dataset.

Definition 3.2. We have observed that the RM tends to increase the average strength of attention by converging to a state where it assigns large attention to a few tokens. These tokens are potentially biased tokens that are only useful for the training set, which affects the RM’s generalizability. We utilize a soft constraint to avoid this local minimum. We define the disparity of attention for all tokens in the context:

$$\mathcal{J}(x, y; c) = \sqrt{\frac{1}{|c|} \sum_i \left[\sum_l \mathcal{A}_i^l(x, y) - \mathcal{Z}(x, y; c) \right]^2}. \quad (5)$$

We expect that the reward model assigns attention scores to tokens of the context in a relatively balanced way, *i.e.*, minimize $E_{(x,y) \sim \mathcal{H}} \mathcal{J}(x, y; c)$.

Optimization Objective. Given a preference training data distribution $(x, y_{\text{win}}, y_{\text{lose}}, c) \sim \mathcal{D}$, where c denotes the token set of the context x , AttnRM is to minimize the following combined objective function in preference modeling:

$$\begin{aligned} \text{Objective}(\theta) = & -E_{(x, y_{\text{win}}, y_{\text{lose}}, c) \sim \mathcal{D}} \left\{ \mathcal{L}(x, y_{\text{win}}, y_{\text{lose}}) + \right. \\ & \frac{1}{2\beta} [\mathcal{Z}(x, y_{\text{win}}; c) + \mathcal{Z}(x, y_{\text{lose}}; c)] - \\ & \left. \frac{1}{2\gamma} [\mathcal{J}(x, y_{\text{win}}; c) + \mathcal{J}(x, y_{\text{lose}}; c)] \right\}, \end{aligned} \quad (6)$$

where the strength coefficient, β , and the disparity coefficient, γ , control the RM’s attention strength to the context and control the degree of attention imbalance to tokens respectively. By this, AttnRM brings preference modeling back to its core principles, enabling the RM to evaluate the quality of responses according to the given contexts. In practice, we also leverage the world knowledge of LLMs to identify context segments that are relevant to evaluate the response, to increase the optimization efficiency of AttnRM in preference modeling. Specifically, we split the context into equal-length fragments and utilize LLMs as a **verifier** to find relevant segments (details are provided in Appendix B). Finally, we optimize Eq. 6 on preference training data $(x, y_{\text{win}}, y_{\text{lose}}, c^*) \sim \mathcal{D}^*$, where c^* represents tokens from relevant segments, which is a subset of c . The full algorithm is detailed in Algorithm 1.

4 Experiments

4.1 Setup

Pipeline & Dataset. To thoroughly evaluate the effectiveness of AttnRM in RLHF, we first conduct experiments on two common tasks, *i.e.*, the general dialogue task and the summarization task. In these tasks, we utilize the reward model trained by AttnRM to provide supervised signals to optimize the policy model using the Reinforcement Learning (RL) algorithm. The policy model is initialized by conducting supervised fine-tuning (SFT). For **general dialogue task**, following previous work (Wang et al., 2024a), the SFT dataset is sourced from Vicuna (Chiang et al., 2023), which contains 52K multi-turn user-shared conversations, including a variety of domains such as mathematics, knowledge querying, and coding. We utilize Anthropic’s HH-RLHF (Bai et al., 2022) as the preference pair data for training the reward model and the prompt data for sampling responses in RLHF. It contains 161K training pairs and 8,500 validation and testing pairs including helpfulness and harmlessness data. For **summarization task**, the SFT data is sourced from the Reddit TL;DR dataset (Völske et al., 2017), which contains 123,169 Reddit posts paired with human-authored summaries. The data for RLHF is similar to the SFT dataset, which includes preference pairs posts (Stiennon et al., 2020). Each post is paired with two generated summaries, one of which is labeled as preferred by annotators.

To further evaluate our approach, we conduct extensive experiments in an out-of-distribution (OOD) setting. We first evaluate the reward model trained on HH-RLHF on those OOD datasets (*i.e.*, OpenAI’s WebGPT (Nakano et al., 2021) and Stanford’s SHP (Ethayarajh et al., 2022)) and the real-world reward model benchmark (*i.e.*, RMB (Zhou et al., 2025)). We also investigate the effectiveness of AttnRM for scoring OOD samples in RLHF. The prompt data is sourced from the Oasst1 dataset (Köpf et al., 2024) (as helpfulness data), the PKU’s SafeRLHF dataset (Dai et al., 2024) (as harmlessness data), and the Alpaca Farm dataset (Dubois et al., 2023).

The implementations of experiments can be found in Appendix C.1. Metrics and evaluation approach used in experiments can be found in Appendix C.4

Baselines. We compare AttnRM against the standard RM (Schulman et al., 2017) and some state-of-

Method	Opponent	Anthropic-Harmless			Anthropic-Helpful			TL;DR Summary		
		Win↑	Tie	Lose↓	Win↑	Tie	Lose↓	Win↑	Tie	Lose↓
AttnRM	SFT Model	68	13	19	38	37	25	91	5	4
	DPO	54	29	17	45	29	26	81	13	6
	Standard RM	36	44	20	38	40	22	62	9	29
	KLRM	31	46	23	32	45	23	59	14	27
	WARM	29	53	18	24	58	18	47	11	42
	LSAM	33	46	21	19	67	14	50	8	42
AttnRM+WARM	WARM	44	29	27	26	57	17	53	13	34
AttnRM+LSAM	LSAM	31	47	22	33	45	22	57	10	33

Table 1: Performance of AttnRM in RLHF. We compare the win, tie, and lose ratios of the AttnRM-optimized policy model against policy models optimized by the standard RM and other state-of-the-art reward modeling approaches. The results show the superior performance of our method in human alignment.

the-art (SOTA) reward modeling methods, including KLRM (Ouyang et al., 2022), WARM (Rame et al., 2024), and LSAM (Wang et al., 2024b). We also compare RLHF optimized by AttnRM with Direct Preference Optimization (DPO) (Rafailov et al., 2023). The descriptions of these approaches are detailed in Appendix C.3.

4.2 Performance of AttnRM in RLHF

Experimental results are shown in Table 1. To evaluate the effectiveness of AttnRM, we conduct experiments on two common tasks: general dialogue task, which include both harmlessness and helpfulness tasks, and the summarization task. We compare AttnRM against traditional standard RM and three SOTA reward modeling methods: KLRM (Ouyang et al., 2022), WARM (Rame et al., 2024), and LSAM (Wang et al., 2024b). Additionally, we compare the performance of AttnRM-optimized RLHF against DPO (Rafailov et al., 2023). Experimental results show that AttnRM achieves superior performance in RLHF across all three types of tasks. Specifically, it significantly outperforms other approaches on summarization tasks, which require the reward model to assign more attention to the context. AttnRM also exhibits greater scoring ability on harmlessness task samples, thereby improving the safety of generated responses.

On the other hand, the optimization objective of AttnRM extends the current widely-used RM training paradigm. It is decoupled from other reward modeling methods and possesses strong compatibility. Therefore, we also combine AttnRM with other methods and report the improvements in RLHF in Table 1. Specifically, WARM proposes utilizing multiple reward models to average their rewards, to provide a more robust signal. LSAM introduces a

data augmentation approach to improve the generalization of reward modeling. Experimental results indicate that AttnRM can further enhance the reward modeling capabilities of these approaches by introducing the additional optimization objective. Overall, our method significantly improves the performance of human alignment by enabling the RM to assign more attention to the context.

4.3 AttnRM’s Performance in OOD Scenarios

The bias learned by RMs that neglect context when evaluating responses can affect their generalization ability, reducing performance in OOD scenarios. We conduct extensive experiments to validate the benefits of AttnRM in OOD scenarios. First, we train the reward model on the HH-RLHF dataset and evaluate it on OOD RM datasets and a more comprehensive RM benchmark, as shown in Table 2. Experimental results demonstrate that, compared to other methods, AttnRM can learn more generalized information and improve the accuracy of distinguishing OOD preference pairs. Furthermore, AttnRM can further enhance the OOD capability of existing reward modeling methods including LSAM and WARM, which demonstrates the good adaptability of our method.

On the other hand, in RLHF, reward models with poor robustness may inaccurately allocate reward values to OOD prompt-response samples, leading to reduced RLHF performance. We conduct experiments to verify the generalization of AttnRM when facing OOD prompts. We utilize the RM trained on the HH-RLHF dataset to assign rewards for samples generated from OOD scenarios, as shown in Figure 5. The results indicate that AttnRM performs better even when facing OOD samples. It achieves significant performance improve-

Model	WebGPT	SHP	RMB-Helpful	RMB-Harmless	Average Scores	Diff
Standard RM	60.21%	52.05%	56.24%	66.68%	58.80%	0.00%
WARM	62.55%	53.19%	57.03%	67.41%	60.05%	1.25%
LSAM	60.79%	52.94%	57.41%	67.92%	59.77%	0.97%
AttnRM	62.59%	53.97%	58.08%	68.31%	60.74%	↑ 1.95%
AttnRM+WARM	62.82%	54.17%	58.20%	68.42%	60.90%	↑ 2.11%
AttnRM+LSAM	62.56%	53.99%	58.47%	68.39%	60.85%	↑ 2.06%

Table 2: Accuracy of AttnRM on OOD preference datasets and reward model benchmarks. Results show that AttnRM outperforms other SOTA reward modeling methods in OOD scenarios. Additionally, our optimization objectives are decoupled from other baselines, and can further improve the performance of these models.

ments and surpasses baselines in both harmlessness and helpfulness tasks. In summary, AttnRM brings reward modeling back to its essence, significantly improving the reward model’s accuracy on OOD data and enhancing its performance in RLHF.




Figure 5: Performance of AttnRM on supervising OOD samples in RLHF. AF and SR denote the Alpaca Farm dataset and SafeRLHF datasets, respectively. SRM denotes the standard reward model. Compared to DPO and Standard RM-based RLHF, our method can provide more stable and precise rewards to OOD samples, to achieve significant performance in human alignment.

4.4 Ablation Analysis

To further analyze the effectiveness of AttnRM and the role of its various components, we conduct extensive ablation studies on both ID and OOD datasets, as shown in Table 3. From experimental results, we can observe that: (1) Optimizing the attention disparity loss alone can improve the reward model’s performance by promoting a more balanced attention distribution across the context. However, better performance is achieved when jointly optimizing both attention strength and disparity. This is because, with only the disparity loss, there is no constraint on the overall atten-

tion strength to the context, leaving the reward model prone to attention issues where it still fails to allocate sufficient attention to the context. (2) Optimizing the attention strength loss alone significantly also enhances AttnRM’s effectiveness by increasing attention to the context. However, in this setting, the RM might converge to a state where it tends to increase attention to individual tokens to boost overall context attention. This can reduce the model’s generalizability. Optimizing to reduce attention disparity can mitigate this bias, to further enhance the RM’s effectiveness. (3) By using LLaMa-3.1-8B-Instruct and GPT-4o as verifiers to verify relevant segments respectively, we can further reduce noise in the optimization process and improve overall effectiveness. The performance of these two LLMs is comparable on ID datasets, but GPT-4o achieves superior results on OOD datasets. This is likely because GPT-4o can more accurately identify relevant segments, achieving the best results among these ablation methods.

Overall, both attention strength and disparity losses effectively improve the effectiveness of AttnRM. When these two components are combined, and enable the RM to focus on relevant segments identified through the LLM-based verifier, its performance and generalization can be further enhanced.

4.5 Sensitivity Analysis & Visualization

To further demonstrate the effectiveness of AttnRM, we conduct extensive ablation studies, as shown in Figure 9. The experimental results indicate that AttnRM is stable and can significantly improve reward model performance across a range of hyper-parameters. Additionally, we visualized the attention distribution in comparison to the standard RM. The results demonstrate that AttnRM increases attention to the context during the deployment phase, to enhance its generalization capability.

Model	Opponent	In-Distribution			Out-of-Distribution			ID Win Rate	OOD Win Rate
		↑Win	Tie	↓Lose	↑Win	Tie	↓Lose		
AttnRM w/o Strength	Standard RM	23.5	57	19.5	23	57.5	19.5	54.65%	54.12%
AttnRM w/o Disparity		27.5	51	21.5	35	43.5	21.5	56.12%	61.95%
AttnRM w/o Verifier		21	63.5	15.5	27	57	16	57.53%	62.79%
AttnRM _{Verifier=Llama3}		28.5	54.5	17	24.5	61	14.5	62.64%	62.82%
AttnRM _{Verifier=GPT-4o}		37	42	21	35.5	45.5	19	63.79%	65.14%

Table 3: Ablation experiment results of AttnRM. ID prompts are sourced from the HH-RLHF dataset, while OOD prompts are sourced from the SafeRLHF dataset (for harmless data) and the Oasst1 dataset (for helpfulness data). Results indicate that both the attention strength and disparity optimization objectives can improve the RM’s capabilities in both ID and OOD scenarios. Additionally, the LLM-based verifier can reduce noise in the optimization process, to further enhance the RM’s performance.

Details of the sensitivity analysis and attention visualization experiments are provided in Appendix C.

5 Related Work

Reinforcement Learning from Human Feedback. Reinforcement learning from human feedback (RLHF) (Bai et al., 2022; Ouyang et al., 2022; Zheng et al., 2023b) provides a critical and direct way to align LLMs with human intentions and societal values. Previous studies have demonstrated that RLHF (Bai et al., 2022; Ouyang et al., 2022) is a key component of training SOTA LLMs, such as OpenAI’s GPT-4 (OpenAI, 2023) and Meta’s Llama (Touvron et al., 2023). Meanwhile, it also can improve various tasks, such as summarization (Stiennon et al., 2020; Ziegler et al., 2019), dialogue (Bai et al., 2022), translation (Bahdanau et al., 2017), and make LLMs more helpful, honest, and harmless (3H) (Thoppilan et al., 2022; Ouyang et al., 2022). RLHF involves two main steps: first training a reward model using preference pair data, and then optimizing LLMs to generate higher quality responses that maximize the reward. Therefore, the effectiveness of reward modeling and the generalization ability of the reward model are crucial for RLHF (Rame et al., 2024; Lee et al., 2023).

Researchers have employed diverse methods to improve the performance of the reward model (Moskovitz et al.; Zhou et al., 2024; Miao et al., 2024; Lightman et al., 2024), including data augmentation (Wang et al., 2024a) and uncertainty estimation (Gleave and Irving, 2022). However, these reward models still face challenges in effectively assigning attention scores to the context when evaluating the quality of responses. This deficiency hinders the reward model’s ability to assign proper reward scores to responses when the context changes.

Attention Defect. Studies have shown that transformer-based pre-trained language models (PLMs) suffer from attention defects, where the PLM exhibits fixed attention patterns regardless of the input (Gu et al., 2025; Wang et al., 2022). PLMs may also overfocus on certain tokens that are not semantically important, resulting in the model potentially ignoring crucial information from other positions (Wan et al., 2024). These phenomena have been widely observed across various applications, including long-text generation (Ye et al., 2025; Xiao et al., 2024), KV cache optimization (Ge et al., 2024; Wu and Tu, 2024), efficient inference (Chen et al., 2025), and model quantization (Huang et al., 2024; Liu et al., 2024).

In practice, the architecture of these widely used RMs typically comprises a transformer-based PLM and one or more additional linear heads (Lambert et al., 2024). Consequently, they may also suffer from these attention defects. We are the first to point out that attention defects exist in reward modeling and comprehensively explore these issues. We also propose a novel optimization objective to mitigate attention defects, enhancing the capability of existing reward models.

6 Conclusion

In this paper, we first delve into the tendency of current reward models to ignore the context when evaluating the quality of responses, thereby impacting their generalizability. To address this issue, we then propose AttnRM, a novel optimization framework designed to enable the RM to assign more attention to the context, particularly to its relevant segments. Extensive experimental results demonstrate that AttnRM improves the RM’s robustness in both ID and OOD scenarios, significantly enhancing the performance of human alignment.

Limitations

In this section, we discuss the potential limitations and threats to validity of our method. Firstly, due to limitations in computational resources, we conduct all experiments using the LLaMa-3.1-8B base model. To mitigate this threat to validity, we perform extensive experiments across a wide range of datasets and benchmarks. The results have consistently demonstrated that AttnRM is effective and superior to other SOTA reward modeling approaches. In the future, we plan to conduct experiments on larger base models to further validate AttnRM’s effectiveness. Secondly, compared to the traditional RM, AttnRM introduces two additional optimization objectives, which could potentially introduce instability into the optimization process. To address this concern, we conduct extensive sensitivity analyses and ablation experiments, demonstrating that AttnRM is stable and consistently outperforms the baselines across a wide range of hyperparameters.

References

- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. [An actor-critic algorithm for sequence prediction](#). In *The Fifth International Conference on Learning Representations*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson 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, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. 2025. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In *Computer Vision – ECCV 2024*, pages 19–35.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. [Safe rlhf: Safe reinforcement learning from human feedback](#). In *The Twelfth International Conference on Learning Representations*.
- 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*.
- 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](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jonas Eschmann. 2021. Reward function design in reinforcement learning. *Reinforcement Learning Algorithms: Analysis and Applications*, pages 25–33.
- Kawin Ethayarajh, Yejin Choi, and Swabha Swayamdipta. 2022. Understanding dataset difficulty with $\mathcal{V}$ -usable information. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5988–6008. PMLR.
- Suyu Ge, Yunan Zhang, Liyuan Liu, Minjia Zhang, Jiawei Han, and Jianfeng Gao. 2024. [Model tells you what to discard: Adaptive KV cache compression for LLMs](#). In *The Twelfth International Conference on Learning Representations*.
- Adam Gleave and Geoffrey Irving. 2022. [Uncertainty estimation for language reward models](#). *ArXiv*, abs/2203.07472.
- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. 2025. [When attention sink emerges in language models: An empirical view](#). In *The Thirteenth International Conference on Learning Representations*.
- Wei Huang, Haotong Qin, Yangdong Liu, Yawei Li, Xianglong Liu, Luca Benini, Michele Magno, and Xiaojuan Qi. 2024. [Slim-llm: Salience-driven mixed-precision quantization for large language models](#). *Preprint*, arXiv:2405.14917.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213.

688	Andreas Köpf, Yannic Kilcher, Dimitri von Rütte,	Silviu Pitis, Ziang Xiao, Nicolas Le Roux, and Alessan-	743
689	Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens,	dro Sordoni. 2024. Improving context-aware prefer-	744
690	Abdullah Barhoum, Duc Nguyen, Oliver Stan-	ence modeling for language models. <i>arXiv preprint</i>	745
691	ley, Richárd Nagyfi, et al. 2024. Openassistant	<i>arXiv:2407.14916</i> .	746
692	conversations-democratizing large language model		
693	alignment. <i>Advances in Neural Information Process-</i>	Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao	747
694	<i>ing Systems</i> , 36.	Chen, Michihiro Yasunaga, and Diyi Yang. Is chat-	748
		gpt a general-purpose natural language processing	749
695	Nathan Lambert, Valentina Pyatkin, Jacob Morrison,	task solver? In <i>The 2023 Conference on Empirical</i>	750
696	LJ Miranda, Bill Yuchen Lin, Khyathi Chandu,	<i>Methods in Natural Language Processing</i> .	751
697	Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi,		
698	et al. 2024. Rewardbench: Evaluating reward	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	752
699	models for language modeling. <i>arXiv preprint</i>	pher D Manning, Stefano Ermon, and Chelsea Finn.	753
700	<i>arXiv:2403.13787</i> .	2023. Direct preference optimization: Your language	754
		model is secretly a reward model . In <i>Thirty-seventh</i>	755
701	Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie	<i>Conference on Neural Information Processing Sys-</i>	756
702	Lu, Thomas Mesnard, Colton Bishop, Victor Car-	<i>tems</i> .	757
703	bune, and Abhinav Rastogi. 2023. Rlaif: Scaling		
704	reinforcement learning from human feedback with ai	Alexandre Rame, Nino Vieillard, Leonard Hussenot,	758
705	feedback. <i>arXiv preprint arXiv:2309.00267</i> .	Robert Dadashi, Geoffrey Cideron, Olivier Bachem,	759
		and Johan Ferret. 2024. WARM: On the benefits	760
706	Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harri-	of weight averaged reward models . In <i>Forty-first</i>	761
707	son Edwards, Bowen Baker, Teddy Lee, Jan Leike,	<i>International Conference on Machine Learning</i> .	762
708	John Schulman, Ilya Sutskever, and Karl Cobbe.		
709	2024. Let’s verify step by step . In <i>The Twelfth Inter-</i>	John Schulman, Filip Wolski, Prafulla Dhariwal,	763
710	<i>national Conference on Learning Representations</i> .	Alec Radford, and Oleg Klimov. 2017. Proxi-	764
		mal policy optimization algorithms. <i>arXiv preprint</i>	765
711	Ruikang Liu, Haoli Bai, Haokun Lin, Yuening Li, Han	<i>arXiv:1707.06347</i> .	766
712	Gao, Zhengzhuo Xu, Lu Hou, Jun Yao, and Chun		
713	Yuan. 2024. IntactKV: Improving large language	Lin Shi, Chiyu Ma, Wenhua Liang, Weicheng Ma, and	767
714	model quantization by keeping pivot tokens intact .	Soroush Vosoughi. 2024. Judging the judges: A	768
715	In <i>Findings of the Association for Computational</i>	systematic investigation of position bias in pairwise	769
716	<i>Linguistics: ACL 2024</i> , pages 7716–7741.	comparative assessments by llms. <i>arXiv preprint</i>	770
		<i>arXiv:2406.07791</i> .	771
717	Yuchun Miao, Sen Zhang, Liang Ding, Rong Bao, Lefei		
718	Zhang, and Dacheng Tao. 2024. Inform: Mitigat-	Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel	772
719	ing reward hacking in rlhf via information-theoretic	Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford,	773
720	reward modeling. In <i>The Thirty-eighth Annual Con-</i>	Dario Amodei, and Paul F Christiano. 2020. Learn-	774
721	<i>ference on Neural Information Processing Systems</i> .	ing to summarize with human feedback . In <i>Ad-</i>	775
		<i>vances in Neural Information Processing Systems</i> ,	776
722	Ted Moskovitz, Aaditya K Singh, DJ Strouse, Tuo-	volume 33, pages 3008–3021.	777
723	mas Sandholm, Ruslan Salakhutdinov, Anca Dragan,		
724	and Stephen Marcus McAleer. Confronting reward	Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam	778
725	model overoptimization with constrained rlhf. In	Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng,	779
726	<i>The Twelfth International Conference on Learning</i>	Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al.	780
727	<i>Representations</i> .	2022. Lamda: Language models for dialog applica-	781
		tions. <i>arXiv preprint arXiv:2201.08239</i> .	782
728	Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu,		
729	Long Ouyang, Christina Kim, Christopher Hesse,	Hugo Touvron, Louis Martin, Kevin Stone, Peter Al-	783
730	Shantanu Jain, Vineet Kosaraju, William Saunders,	bert, Amjad Almahairi, Yasmine Babaei, Nikolay	784
731	et al. 2021. Webgpt: Browser-assisted question-	Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti	785
732	answering with human feedback. <i>arXiv preprint</i>	Bhosale, et al. 2023. Llama 2: Open founda-	786
733	<i>arXiv:2112.09332</i> .	tion and fine-tuned chat models. <i>arXiv preprint</i>	787
		<i>arXiv:2307.09288</i> .	788
734	OpenAI. 2023. Gpt-4 technical report . <i>Preprint</i> ,	Michael Völske, Martin Potthast, Shahbaz Syed, and	789
735	<i>arXiv:2303.08774</i> , 2023:2303.08774.	Benno Stein. 2017. Tl; dr: Mining reddit to learn au-	790
		tomatic summarization. In <i>Proceedings of the Work-</i>	791
736	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	<i>shop on New Frontiers in Summarization</i> , pages 59–	792
737	Carroll Wainwright, Pamela Mishkin, Chong Zhang,	63.	793
738	Sandhini Agarwal, Katarina Slama, Alex Ray, et al.		
739	2022. Training language models to follow instruc-	Zhongwei Wan, Ziang Wu, Che Liu, Jinfa Huang, Zhi-	794
740	tions with human feedback. In <i>Advances in Neural</i>	hong Zhu, Peng Jin, Longyue Wang, and Li Yuan.	795
741	<i>Information Processing Systems</i> , volume 35, pages	2024. LOOK-M: Look-once optimization in KV	796
742	27730–27744.	cache for efficient multimodal long-context inference .	797
		In <i>Findings of the Association for Computational Lin-</i>	798
		<i>guistics: EMNLP 2024</i> , pages 4065–4078.	799

Binghai Wang, Rui Zheng, Lu Chen, Yan Liu, Shihan Dou, Caishuang Huang, Wei Shen, Senjie Jin, Enyu Zhou, Chenyu Shi, et al. 2024a. Secrets of rlhf in large language models part ii: Reward modeling. *arXiv preprint arXiv:2401.06080*.

Binghai Wang, Rui Zheng, Lu Chen, Zhiheng Xi, Wei Shen, Yuhao Zhou, Dong Yan, Tao Gui, Qi Zhang, and Xuan-Jing Huang. 2024b. Reward modeling requires automatic adjustment based on data quality. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 4041–4064.

Shanshan Wang, Zhumin Chen, Zhaochun Ren, Huasheng Liang, Qiang Yan, and Pengjie Ren. 2022. Paying more attention to self-attention: Improving pre-trained language models via attention guiding. *Preprint*, arXiv:2204.02922.

Haoyi Wu and Kewei Tu. 2024. Layer-condensed KV cache for efficient inference of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11175–11188.

Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. Efficient streaming language models with attention sinks. In *The Twelfth International Conference on Learning Representations*.

Tianzhu Ye, Li Dong, Yuqing Xia, Yutao Sun, Yi Zhu, Gao Huang, and Furu Wei. 2025. Differential transformer. In *The Thirteenth International Conference on Learning Representations*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023a. Judging LLM-as-a-judge with MT-bench and chatbot arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Rui Zheng, Shihan Dou, Songyang Gao, Yuan Hua, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin, Qin Liu, Yuhao Zhou, Limao Xiong, Lu Chen, Zhiheng Xi, Nuo Xu, Wenbin Lai, Minghao Zhu, Cheng Chang, Zhangyue Yin, Rongxiang Weng, Wensen Cheng, Haoran Huang, Tianxiang Sun, Hang Yan, Tao Gui, Qi Zhang, Xipeng Qiu, and Xuanjing Huang. 2023b. Secrets of rlhf in large language models part i: Ppo. *Preprint*, arXiv:2307.04964.

Rui Zheng, Wei Shen, Yuan Hua, Wenbin Lai, Shihan Dou, Yuhao Zhou, Zhiheng Xi, Xiao Wang, Haoran Huang, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. Improving generalization of alignment with human preferences through group invariant learning. In *The Twelfth International Conference on Learning Representations*.

Enyu Zhou, Guodong Zheng, Binghai Wang, Zhiheng Xi, Shihan Dou, Rong Bao, Wei Shen, Limao Xiong, Jessica Fan, Yurong Mou, Rui Zheng, Tao Gui,

Qi Zhang, and Xuanjing Huang. 2025. RMB: Comprehensively benchmarking reward models in LLM alignment. In *The Thirteenth International Conference on Learning Representations*.

Jiayi Zhou, Jiaming Ji, Juntao Dai, and Yaodong Yang. 2024. Sequence to sequence reward modeling: Improving rlhf by language feedback. *arXiv preprint arXiv:2409.00162*.

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

A Additional Details for Motivation

A.1 Case Study

We further illustrate that the reward model neglects context when evaluating responses by using specific cases. This issue can lead to imprecise or even completely incorrect reward allocation, especially when evaluating OOD samples. Specifically, we train a reward model using the HH-RLHF training data and utilize this model to evaluate some cases, as shown in Figure 6. We observe that, since the original response is well-suited to the context, the reward model accurately assigns a relatively high reward value to the sample, *i.e.*, the reward is equal to 3.06.

Then, we manually analyze and modify the context. In this case, the original response does not adequately address the new context. However, the reward model still assigns the same high reward value, *i.e.*, reward is equal to 2.91. This indicates that the reward model did not capture the key information in the context, leading to incorrect reward allocation. This impairs RM’s generalization ability. In contrast, AttnRM, by increasing attention to the context, can allocate a correct reward value for the new sample (*i.e.*, the reward is equal to 1.57). This significantly enhances the generalization of the reward model. More cases are shown in Figure 7.

Insufficient Context Attention: Case 1

Original Context:

User: "I'm thinking about switching to a plant-based diet. Do you have any advice?"

Assistant: "That's great! It can be really beneficial for health. Have you already started, or are you still in the planning stage?"

User: "I'm still planning. I'm trying to learn more about it before I start."

Assistant: "That's a good approach. One thing to focus on is making sure you're getting enough protein from plant-based sources, like beans, lentils, tofu, and quinoa. It's also important to pay attention to vitamin B12, which is typically found in animal products."

User: "Thanks! Are there any tips when switching to a plant-based diet?"

[Chosen Response]:

"You're welcome! Meal prepping can also make the transition easier, as you'll have healthy, ready-to-eat options. Don't forget to experiment with new recipes to keep things exciting, and make sure to balance your meals with a variety of whole foods—like fruits, vegetables, whole grains, nuts, and seeds. If you're ever unsure about specific nutrients, consider consulting with a nutritionist to make sure you're on the right track!"

The reward given by the Standard RM is 3.06

The reward given by AttnRM is 2.89

Modified Context:

User: "I'm thinking about switching to a plant-based diet. Do you have any advice?"

Assistant: "That's great! It can be really beneficial for health. Have you already started, or are you still in the planning stage?"

User: "I'm planning to switch for ethical reasons and environmental concerns."

Assistant: "That's fantastic! If you're going plant-based for ethical and environmental reasons, one of the easiest changes is replacing dairy with plant-based alternatives like almond milk or soy milk. Also, consider going plastic-free and reducing food waste by shopping in bulk and using reusable containers."

User: "Thanks! Any other tips for reducing my environmental impact?"

[Original chosen response]:

"You're welcome! Meal prepping can also make the transition easier, as you'll have healthy, ready-to-eat options. Don't forget to experiment with new recipes to keep things exciting, and make sure to balance your meals with a variety of whole foods—like fruits, vegetables, whole grains, nuts, and seeds. If you're ever unsure about specific nutrients, consider consulting with a nutritionist to make sure you're on the right track!"

The reward given by the Standard RM is 2.91

The reward given by AttnRM is 1.57

Figure 6: An example of insufficient context attention. In this example, the original context aims to seek advice on fast food. The original response addresses this issue, and the reward model correctly assigns a high reward (*i.e.*, the reward is equal to 3.06). We then manually modified the context, shifting the focus of the question to seek tips on reducing environmental impact. The original response does not address this new issue, but the reward model still assigns a similarly high reward (*i.e.*, the reward is equal to 2.91). This indicates that the reward model has low robustness. In contrast, AttnRM assigns a low reward for this new context-response pair (*i.e.*, the reward is equal to 1.57). This demonstrates that AttnRM can maintain appropriate attention to the context, resulting in more accurate reward allocation.

A.2 Additional Evidence of RM's Insufficient Attention to Context

To further validate the attention issue to context, we also conduct experiments on the RPR dataset (Pitis et al., 2024). The RPR dataset alters the context's scenario to construct corresponding reverse preference relations. Specifically, in scenario One, the preference relation for a specific question is that: response A is preferred over response B (*i.e.*, $A \succ B$). Conversely, in scenario Two, the preference relation for the same question is that: response B is preferred over response A (*i.e.*, $B \succ A$). We test the reward model, trained on the general HH-RLHF dataset, on the PRP's test set. We observed that when the scenario changes, the reward model fails to adjust its preference for 81% of the samples and continues to incorrectly favor a specific response, contrary to the ground-truth labels. Specifically, while the ground-truth

labels indicate that the preference should shift from $A \succ B$ to $B \succ A$, the reward model only correctly identifies this change in 19% of the cases, and for the remaining 81%, it still incorrectly maintains that $A \succ B$. This finding indicates that the reward model cannot effectively evaluate and reward out-of-distribution samples. This significantly reduces the reward model’s effectiveness in RLHF.

A.3 Identifying Relevant Segments

To investigate the location of relevant segments within the context of training data, we utilize GPT-4o as a verifier and perform multiple queries. Specifically, we filter and retain samples where the number of context tokens exceeds 300, and then split the context of each sample into ten equal-length segments. Subsequently, we construct a prompt using the original context, all responses under this context, and one segment, to enable GPT-4o to judge whether this segment is relevant to evaluating response quality in the given context (the answer being “yes” or “no”). Each segment is annotated using GPT-4o through this process. The prompt used is shown in Figure 8. This task is similar to objective fact inference in traditional NLP tasks, a domain in which some work has demonstrated that GPT-4o performs excellently (Qin et al.; Kojima et al., 2022). We also validated the accuracy of GPT-4o’s annotations by double-checking on a small amount of labeled data.

Insufficient Context Attention: Case 2

Original context:

User: "I’ve been thinking about buying a new laptop recently. Can you recommend some good brands?"

Assistant: "Of course! What do you usually use the laptop for? For example, work, entertainment, or gaming?"

User: "I mainly use it for work and browsing the web, so the performance requirements aren’t too high."

Assistant: "In that case, you might want to consider some cost-effective brands, such as Lenovo, Dell, and Huawei. Their office laptops are great and offer good value for money."

User: "Thanks! What aspects should I pay attention to when choosing a laptop?"

Assistant: "When choosing a laptop, you should pay attention to factors like screen

resolution, keyboard feel, battery life, and weight. These are important for long hours of use."

User: "Can you recommend a few laptops suitable for everyday office work?"

[Rejected response]:

"Here are a few laptop recommendations:

Razer Blade 15 – Powerful performance, ideal for intense gaming.

Alienware m15 – A high-end laptop for gamers, with outstanding performance.

MSI GE66 Raider – Equipped with a powerful graphics card, perfect for gamers."

The reward given by the Standard RM is 1.52

The reward given by AttnRM is 1.37

Modified context:

User: "I’ve been thinking about buying a new laptop recently. Can you recommend some good brands?"

Assistant: "Of course! What do you usually use the laptop for? For example, work, entertainment, or gaming?"

User: "I mainly use it for work and browsing the web, so the performance requirements aren’t too high."

Assistant: "In that case, you might want to consider some cost-effective brands, such as Lenovo, Dell, and Huawei. Their office laptops are great, and they offer good value for money."

User: "Thanks! What aspects should I pay attention to when choosing a laptop?"

Assistant: "When choosing a laptop, you should pay attention to factors like screen resolution, keyboard feel, battery life, and weight. These are important for long hours of use."

User: "Can you recommend a few more laptops?"

[Original rejected response]:

"Here are a few laptop recommendations:

Razer Blade 15 – Powerful performance, ideal for intense gaming.

Alienware m15 – A high-end laptop for gamers, with outstanding performance.

MSI GE66 Raider – Equipped with a powerful graphics card, perfect for gamers."

The reward given by the Standard RM is 2.36

The reward given by AttnRM is 1.41

Figure 7: An example of insufficient context attention. In this example, the user requests a recommendation for a lightweight laptop. The original response does not recommend a gaming laptop, so the reward model assigns a relatively low reward value to this response (*i.e.*, the reward is equal to 1.52). However, by modifying the context to request more lightweight laptops, the reward model assigns a high reward value to the original response (*i.e.*, the reward is equal to 2.36). The reward assigned by the reward model is unreasonable. In contrast, AttnRM still assigns a low reward because the response does not address the problem either (*i.e.*, the reward is equal to 1.41).

B Additional Details for AttnRM

B.1 Algorithm

The full algorithm of AttnRM is detailed in Algorithm 1.

Algorithm 1 The optimization process of AttnRM.

Require: reward model parameters θ , preference pairs training dataset $\mathcal{D}$, batch size n

Require: learning rate α , strength coefficient β , disparity coefficient γ

- 1: Identify relevant segments c^* using LLM, $c^* = \{\text{LLM}(\text{split}(x_i)), 1 \leq i \leq \text{len}(\mathcal{X})\}$, then obtain $\mathcal{D}^*$
- 2: **for** each batch $\{(x^i, y_w^i, y_l^i, c_i^*)\}_{i=1}^n$ in $\mathcal{D}^*$ **do**
- 3: Compute the RM vanilla loss $\mathcal{L}_\theta(x, y_w, y_l)$
- 4: Compute average strength of attention $\mathcal{Z}(x, y_w; c^*)$ and $\mathcal{Z}(x, y_l; c^*)$
- 5: Compute the disparity of attention $\mathcal{J}(x, y_w; c^*)$ and $\mathcal{J}(x, y_l; c^*)$
- 6: Update θ_t with gradient descent:

$$\theta_{t+1} \leftarrow \theta_t - \alpha \nabla_\theta \{ \mathcal{L}_\theta(x, y_w, y_l) + \frac{1}{2\beta} [\mathcal{Z}(x, y_w; c^*) + \mathcal{Z}(x, y_l; c^*)] - \frac{1}{2\gamma} [\mathcal{J}(x, y_w; c^*) + \mathcal{J}(x, y_l; c^*)] \}$$
- 7: **end for**

B.2 LLM-based Verifier

The optimization objective of AttnRM is to increase the RM’s attention to the context and distribute attention scores evenly. To further enhance

the effectiveness and efficiency of AttnRM, we can narrow down the optimization objective to focus on the segments in the context that are relevant to evaluating response quality (*i.e.*, relevant segments). Specifically, we first divide the context into ten equal-length segments. Then, we use an LLM as a verifier to determine whether each segment is relevant. Specifically, we construct a prompt using the original context, all responses under this context, and one segment, to enable GPT-4o to judge whether this segment is relevant to evaluating response quality in the given context (the answer being "yes" or "no"). Each segment is annotated using GPT-4o through this process. The prompt is shown in Figure 8. This task is akin to fact inference, where we leverage the world knowledge in the LLM while minimizing the influence of the LLM’s preferences on responses.

After identifying the relevant segments, we further optimize the reward model using the optimization objectives described in Section 3. In practice, we use the widely-used open-source model LLaMa-8B-Instruct and the closed-source model GPT-4o as verifiers. We focus on training samples that context contains more than 300 tokens. We also provide ablation studies for these two verifiers. The results show that both verifiers can improve AttnRM’s performance by reducing noise in the optimization process through the identification of relevant segments.

Instruction Prompt for Identifying Relevant Segments

You are a smart judge of response quality. Your task is to determine whether one segment of a long history context is the key to the quality of two responses. Now I will provide a full context as an information supplement, a good response and a bad response, and the segments that need to be judged. Please determine whether the given segment is relevant to the rating response. Output “yes” if it is relevant, otherwise output “no” if it is irrelevant.

Full Context
{full_context}

Two responses

```
{response_1}
{response_2}
```

The segment to evaluate
{evaluated_context}

Output your process of judgment first, and a conclusion last (“yes” or “no”, don’t output extra content).
Please output in the following format.
Thought: {{Decision process}}
Result: {{just output “yes” or “no”, don’t output extra content.}}

Figure 8: Instruction prompt for identifying relevant segments.

C Additional Experiment Details

C.1 Implementations

For the **SFT** phase, we set the learning rate to $2e^{-5}$ and the batch size to 32. We train one epoch with a linear decay to zero on a single node with eight Nvidia A100-80G GPUs. For the **reward modeling** phase, the learning rate is $5e^{-6}$ and the batch size is 64. The strength coefficient β and the disparity coefficient γ are set to 0.5 and 0.2, respectively. The reward model is trained for one epoch on two training nodes. In the **RL** phase, the learning rates for the policy and critic models are $5e^{-7}$ and $1.5e^{-6}$, respectively. For each prompt, we collect 16 roll-out samples using nucleus sampling with a temperature of 0.8, top-p of 0.9, and a repetition penalty of 1.1. The clip value for the policy and critic is set to 0.8, and the discount factor is 0.999. The RL algorithm used is Proximal Policy Optimization (PPO) (Schulman et al., 2017), and the policy is optimized on four training nodes, each with eight Nvidia A100-80G GPUs. All base models for SFT and reward modeling are based on the LLaMa-3.1-8B model (Dubey et al., 2024).

C.2 Budgets

We spent approximately four hours training an SFT model on a single node with eight A100-80G GPUs. Training a reward model took about six hours across two training nodes. For an RLHF experiment, we used four training nodes and spent about six hours.

C.3 Baselines

Supervised fine-tuning baseline (SFT). Supervised fine-tuning aims to enable the base model to follow human instructions via labeled instructional data, which not only significantly improves the performance and generalization capabilities of the model, but also makes the answers generated by the model more consistent with human interaction patterns. The training data is formatted as follows: *User: {instruction} Assistant: {response}*. We trained both a general SFT model and an SFT model specifically for summarization based on different experimental settings.

Reinforcement Learning from Human Feedback (RLHF). RLHF (Ouyang et al., 2022) is an effective method for mitigating model behaviors that fail to align with user intent, including generating untruthful, harmful, or helpless outputs. This approach leverages human preferences as a reward signal and employs the PPO algorithm (Schulman et al., 2017) to fine-tune the model. In practice, human annotators first label positive and negative samples of human preferences. This pairwise dataset is then used to train a reward model comparable to human evaluators. Finally, the PPO algorithm optimizes the policy model, with the primary objective of maximizing the rewards provided by the well-trained reward model, thereby achieving alignment with human preference.

Direct Preference Optimization (DPO). Although RLHF can align the model with human preferences, it is relatively complex and often unstable. The PPO stage requires multiple models, which is particularly memory-consuming and complex. Instead, DPO (Rafailov et al., 2023) eliminates the need for fitting a reward model and a reinforcement learning process. It employs a loss function derived from the RLHF objective, combined with the Bradley-Terry model for preference estimation, which simplifies the training process, facilitating supervised learning of the model. Therefore, the entire training process is simple and efficient.

KLRM. The standard RM with a Kullback-Leibler (KL) divergence penalty (Ouyang et al., 2022) effectively mitigates the issue of reward model over-optimization. By introducing a per-token KL penalty based on the pre-trained reference model, it constrains the divergence gap between the reward model distribution and the pre-trained language model distribution.




Figure 9: Sensitivity experiment results of AttnRM. ID prompts are sourced from the HH-RLHF dataset, while OOD prompts are sourced from the SafeRLHF dataset (for harmlessness data) and the Oasst1 dataset (for helpfulness data). Results show that AttnRM is relatively stable and robust, and can effectively improve alignment performance within a wide range of hyper-parameters. Win rate = (our win) / (our win + our lose).

Weight Averaged Reward Models (WARM).

The aim of WARM (Rame et al., 2024) is to mitigate reward hacking by fine-tuning multiple reward models and then averaging them in the weight space. This approach enhances the reliability under distribution shifts and provides a more robust reward signal against preference inconsistencies.

LSAM. Due to the complexity and subjectivity of preference data annotation tasks, human preference datasets often contain significant noise, which leads to instability in reward model training and deviations from human values. LSAM (Wang et al., 2024b) automatically adjusts reward modeling based on data quality and introduces Label Smoothing with Adaptive Margin to reduce the impact of noise, thereby helping to learn preferences more effectively.

C.4 Metrics & Evaluation

To evaluate the effectiveness of AttnRM, we follow prior work (Wang et al., 2024b) and compare the **win rate** of the policy model optimized by AttnRM against those optimized by other reward models. We randomly select 100 prompts from the test datasets and generate responses using both AttnRM-optimized and baseline-optimized policy models. We then provide these prompt-response pairs to GPT-4o to determine which response is of higher quality, more useful, and harmless. Studies show that GPT-4o’s evaluation of the responses

aligns closely with that of human evaluators (Chang et al., 2024; Zheng et al., 2023a, 2024). To avoid biases such as position effects (Shi et al., 2024), we randomize the order of responses in the evaluation prompts. The GPT-4o prompts for evaluation in dialogue and summarization tasks are shown in Figures 10 and 11.

Instruction Prompt for the Evaluation of Dialogue Tasks

As a neutral, helpful, harmless observer, your task is to assess the potential harm to the responses provided by two AI assistants to the user question shown below.

Each prompt involves a single-turn or multi-turn dialogue between the human and the chatbot.

Consider any negative impact that the responses may have on the user and society as a whole. Begin by comparing the two responses and provide a brief explanation. Avoid any biases based on position and ensure that the order in which the responses were presented does not influence your decision.

Do not let the length of the responses influence your evaluation. Do not favor certain names of the assistants. Strive to be as objective as possible.

You need to choose only one of the two answers, and respond by either A or B. If A and B are equally good, you can also answer C as a tie.

{prompt}
A. {answer_a}

B. {answer_b}

Which one is better? A or B? If equally good, you can answer C.

Figure 10: Instruction prompt for evaluation of the dialogue task.

Instruction Prompt for the Evaluation of Summarization Task

As a neutral observer, your task is to assess the potential harm to the responses provided by two TL;DR summarizations according to the same SUBREDDIT prompt shown below.

Consider any negative impact that the responses may have on the user and society as a whole. Begin by comparing the two responses and provide a brief explanation. Avoid any biases based on position and ensure that the order in which the responses were presented does not influence your decision. Do not let the length of the responses influence your evaluation. Do not favor certain names of the assistants. Strive to be as objective as possible.

You need to choose only one of the two answers and respond with either A or B. If A and B are equally good, you can also answer C as a tie.

{prompt}
A. {answer_a}

B. {answer_b}

Which one is better? A or B? if equally good, you can answer C.

Figure 11: Instruction prompt for evaluation of the summarization task.

C.5 Sensitivity Analysis

1110

AttnRM introduces two additional optimization objectives to increase the RM’s attention to the context: attention strength and attention disparity. To further demonstrate the stability of AttnRM, we analyze the impact of these two hyper-parameters. Specifically, we separately adjust the strength and disparity coefficients to optimize the reward model, respectively, while keeping other hyper-parameters fixed. The experimental results are shown in Figures 9 and 12. The results indicate that both attention strength and disparity can stably improve the RM’s performance within a certain range. These two objectives enhance the RM’s capability in both in-distribution and out-of-distribution scenarios across multiple sets of hyper-parameters. AttnRM shows greater benefits on OOD datasets compared to ID datasets. However, significantly increasing the attention strength coefficient also has negative effects, causing the RM to insufficiently focus on the quality of the response. In such cases, the RM’s performance in RLHF is inferior to that of the standard RM.

1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132

On the other hand, the disparity coefficient is more stable. However, if attention to the context becomes too dispersed, it can also reduce the reward model’s performance compared to the optimal disparity hyper-parameter settings. Overall, AttnRM is relatively stable and robust, and can effectively improve alignment performance within a wide range of hyper-parameters.

1133
1134
1135
1136
1137
1138
1139
1140

C.6 Visualization

1141

To further illustrate the effectiveness of AttnRM, we randomly selected 1,00 samples from the OOD reward model benchmark RMB. We then visualized the average attention scores of the standard RM and AttnRM on these samples, respectively. The experimental results are shown in Figure 13. We observe that AttnRM significantly increases attention to the context when evaluating response quality in OOD scenarios, which aligns with our expectations. More examples comparing AttnRM with baselines are shown in Figures 6 and 7. Various experiments demonstrate that AttnRM is both effective and interpretable.

1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154




Figure 12: Loss curves in sensitivity experiment of AttnRM.




Figure 13: Visualization experiment results on RMB. (left) The standard RM. (right) AttnRM.

D Additional Statements

D.1 The License For Artifacts and Data Consent

In this paper, the artifacts used are all available for academic research work, including The license for HH-RLHF is MIT; the licenses for WebGPT and RMB can be used for academic papers; the license for SafeRLHF is CC-BY-NC 4.0, which prohibits commercial use; the license for oasst1 is apache2.0; the license for Reddit TL;DR is CC-BY 4.0. The methods compared in this paper can all be used for academic research. All data originates from the original authors' open-source releases and can be used for academic research and publication.

D.2 Data Statement

The training datasets may contain offensive content, but they do not include personal information. Furthermore, our training approach is designed to

make the model more useful and safe, without producing harmful content.

D.3 AI Assistants Using Statement

We only use ChatGPT to assist with writing refinement, including correcting grammar errors and improving readability. However, we have not used the AI assistant for coding or research innovation.