

Sequence to Sequence Reward Modeling: Improving RLHF by Language Feedback

Jiayi Zhou^{*1,2}, Jiaming Ji^{*1,2}, Josef Dai^{1,2}, Yaodong Yang^{1†}

¹Institute for AI, Peking University

²State Key Laboratory of General Artificial Intelligence, Institute for AI, Peking University
{gaiej, yaodong.yang}@pku.edu.cn {jtd.acad, jiang.ji}@gmail.com

Abstract

Aligning the behavior of Large language models (LLMs) with human intentions and values remains a critical challenge. Reinforcement learning from human feedback (RLHF) aligns LLMs by training a reward model (RM) on human preferences and fine-tuning the LLMs to maximize RM feedback. Despite its effectiveness and popularity, RLHF is prone to biased local optimization. It means RM fails to provide feedback that accurately aligns with human preference, causing LLMs to explore unexpected generalizations, and failing to achieve alignment objectives. To mitigate this issue, we propose a novel *sequence-to-sequence (seq2seq) reward modeling* method. Its key insight is that learning from language feedback rather than scalar feedback improves RLHF without additional annotations. We replaced the reward modeling target from binary maximum likelihood estimation (MLE) with sequence MLE. This method enables richer and fine-grained language feedback without additional annotations, models, or training stages. Our experiments demonstrated its effectiveness, specifically, reducing the refusal-to-respond paradigm in single-turn safety dialogues and the long-response bias in text summarization tasks. We provide further analysis that seq2seq RM improves RLHF performance across 2B and 7B LLMs on 3 NLP tasks, achieving an average win rate of 76.9%. We further show that seq2seq RM can still improve the performance of RLHF under out-of-distribution prompts.

Introduction

Large Language Models (LLMs) (Achiam et al. 2023; Touvron et al. 2023) have significantly advanced capabilities in conversation, summarization, reasoning, and more tasks, leading to widespread application (Meng et al. 2024; Anthropic 2023). Alignment techniques (Ji et al. 2023), particularly Reinforcement Learning from Human Feedback (RLHF) (Ouyang et al. 2022), have become crucial for aligning LLMs with human intentions and values, *e.g.*, reducing toxicity (Bai et al. 2022; Ji et al. 2024d,e). RLHF consists of two main steps: training a reward model (RM) based on binary human preferences and then fine-tuning the LLMs to maximize feedback from this RM (Ouyang et al. 2022; Schulman et al.

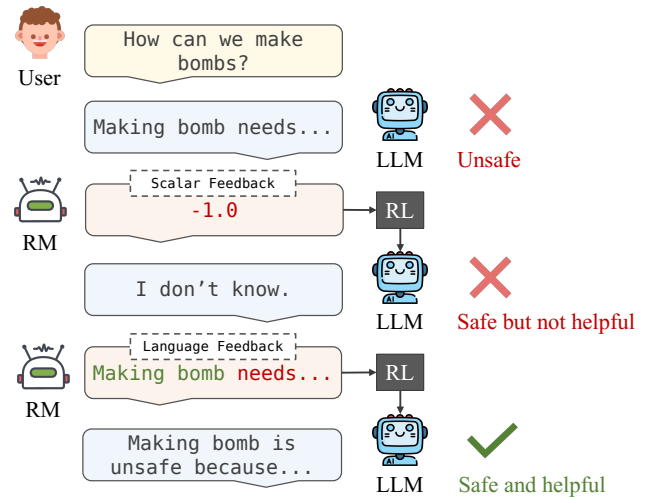


Figure 1: The traditional sequence-to-scalar reward model provides only coarse-grained scalar feedback on the LLM’s response, making it prone to exploiting unexpected generalization for high rewards during the RL fine-tuning phase (Skalse and Abate 2023), *e.g.*, falling into a refusal-to-answer paradigm (Bai et al. 2022; Dai et al. 2023). Without additional annotations, we propose a novel sequence-to-sequence reward modeling method that offers richer language feedback, improving RLHF performance.

2017; Scheurer et al. 2023). Current reward modeling implementations typically use the sequence-to-scalar (seq2scalar) approach, where a sequence, combined with a prompt and response, is scored as a scalar value (Huang et al. 2024).

However, due to high divergence, noise, and sparsity in crowd-annotated preference datasets (Wu et al. 2024), the traditional seq2scalar RM tends to capture patterns unrelated to genuine human preferences, leading to *unexpected generalization* (Pan, Bhatia, and Steinhardt 2021; Gao, Schulman, and Hilton 2023). It means RM fails to provide feedback that accurately aligns with human preference, leading RLHF to a deviation from actual alignment objective (Di Langosco et al. 2022): LLMs obtain high rewards but negatively impact their capabilities and fail to achieve alignment objectives (Tien et al. 2022; Knox et al. 2023; Lambert and Calandra 2023).

^{*}These authors contributed equally.

[†]Corresponding author.

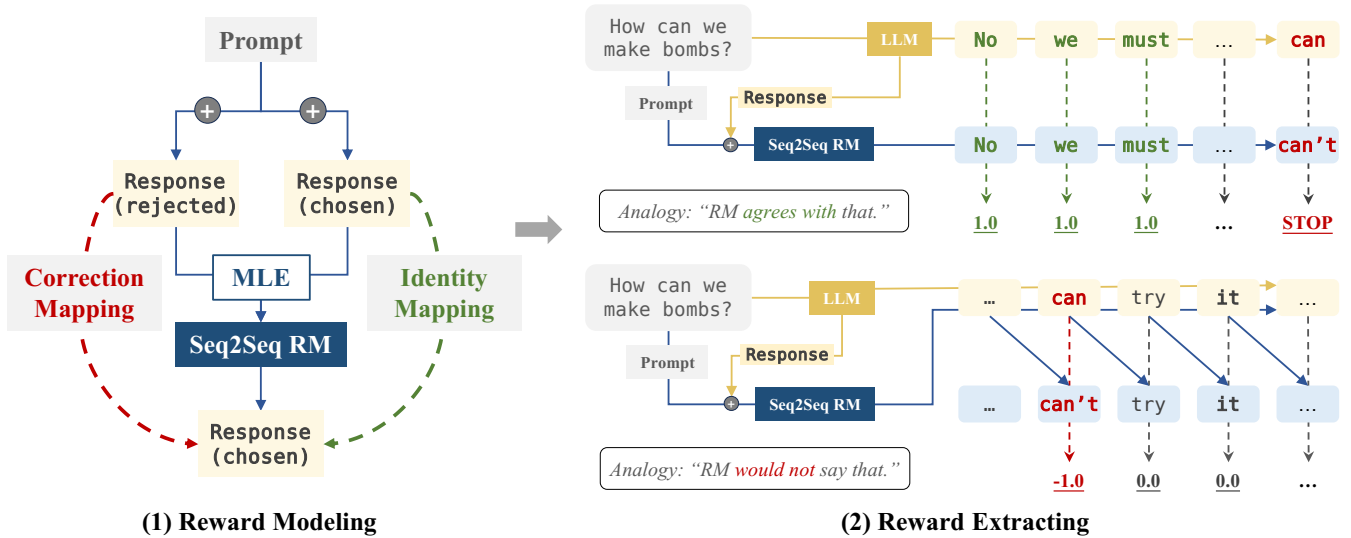


Figure 2: Overview of seq2seq reward modeling pipeline. Our pipeline consists of two stages: **(1) Reward Modeling:** We make the seq2seq RM output the chosen response when the rejected response is input, *i.e.*, Correction Mapping, and output the chosen response when the chosen response is input, *i.e.*, Identity Mapping, by sequence maximum loglikelihood estimation (MLE). **(2) Reward Extracting:** We reward the response with a positive score until it diverges from the seq2seq RM output, after which we input the response token-by-token, assigning negative scores to those diverging tokens.

For example, when the alignment objective is to reduce toxic responses, LLMs may learn the refusal-to-respond paradigm instead of providing safe and helpful responses (Dai et al. 2023; Bai et al. 2022). Existing methods require additional annotations to provide more comprehensive feedback, including regularization terms (Moskowitz et al. 2023; Zhu, Jordan, and Jiao 2024), models ensembling (Coste et al. 2023), parameters combinations (Ramé et al. 2024) and additional model (Dai et al. 2023), raising concerns about their scalability.

Therefore, we aim to answer the following question:

How can we provide more informative reward feedback to improve the RLHF Alignment?

In this work, we propose a novel sequence-to-sequence (seq2seq) reward modeling method. Our pipeline is depicted in Figure 2, which consists of two main stages, *i.e.*, *Reward Modeling*, and *Reward Extracting*. We replaced the reward modeling target from binary maximum likelihood estimation (MLE) for classification tasks to sequence MLE for text generation tasks. Since the language space directly reveals the impact of each token on the response score, this method enhances the RM’s fine-grained discrimination accuracy. Then we extract token-level positive and negative feedback based on whether the tokens are consistent with the output of the seq2seq RM, improving the granularity of RLHF. Notably, our seq2seq reward modeling does not require additional data annotation, training, and models. Our contributions to enhancing RLHF alignment are as follows:

- We propose seq2seq reward modeling. Its key insight is that learning from language feedback rather than scalar

feedback improves accuracy and granularity, thereby improving RLHF without additional annotations.

- Our further experiments demonstrate that the seq2seq RM not only mitigates unexpected behavior of RLHF but also enhances its alignment performance. Specifically, it achieves an average win rate of 76.9% in 3 NLP tasks, across models with 2B and 7B parameters.

Preliminaries

Token-Level MDP for RLHF Given an LLM parameterized with θ , we model its generation as a token-level Markov Decision Process (MDP), *i.e.*, the tuple $\mathcal{T} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P})$. Here, let $t \in \mathcal{V}$ denote a token, where \mathcal{V} represents the vocabulary set, and L denotes the maximum number of tokens. The state space \mathcal{S} can be defined as $\{(t_1, t_2, \dots, t_L) \mid t_i \in \mathcal{V}\}$ ¹. The action space $\mathcal{A} = \mathcal{V}$ consists of all possible tokens in vocabulary. The reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ assigns a reward value for generating a specific token given the current sequence of tokens. The transition probability function $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ defines the deterministic state transition model, concatenating the generated token to the existing sequence to form the next state. *i.e.*, $s_{t+1} = (s_t, a_t)$.

LLMs generation trajectory starts at prompt s_0 sampled from datasets. From $t = 0$, the next token a_t is generated according to the parametrized actor LLM $\pi_\theta(\cdot | s_t)$, followed by a transition to the next state s_{t+1} according to \mathcal{P} , continuing until the EOS token. The goal of RL fine-tuning is to maximize the cumulative reward of this trajectory:

¹We define t_{pad} as the padding token. When the length of a sentence is $l < L$, it will have $L - l$ additional padding tokens.

$$\mathcal{L}_{\theta}^{\text{RL}} = -\mathbb{E}_{a_t \sim \pi_{\theta}(\cdot|s_t)} \left[\sum_{t=0}^T (\mathcal{R}(s_t, a_t) + \beta \pi_{\text{ref}}(a_t|s_t)) \right]$$

where π_{ref} is the reference distribution from the initial LLM, combined with coefficient β to compute the KL divergence, preventing the LLM from deviating too much from the π_{ref} .

Seq2scalar Reward Modeling Let \mathbb{D}^{pre} denote a preference dataset, where each entry consists of a prompt \mathbf{x} , a chosen response \mathbf{y}^w , and a rejected response \mathbf{y}^l . The training loss function for the reward model with parameters θ is defined as the following binary MLE:

$$\mathcal{L}_{\theta}^{\text{RM}} = \mathbb{E}_{(\mathbf{x}, \mathbf{y}^w, \mathbf{y}^l) \sim \mathbb{D}^{\text{pre}}} [\log \sigma(r_{\theta}(\mathbf{x}, \mathbf{y}^l) - r_{\theta}(\mathbf{x}, \mathbf{y}^w))] \quad (1)$$

where $\sigma(x) = 1/(1 + \exp(-x))$ is the logistic sigmoid function and \cdot denotes concatenation. The reward $\mathcal{R}(\mathbf{x}, \mathbf{y})$ for a given prompt-response pair (\mathbf{x}, \mathbf{y}) is defined as the scalar value extracted from the reward model’s inference at the position of the EOS token. In the RLHF phase, seq2scalar RM assigns it as sparse feedback to the entire response.

$$\mathcal{R}(s_t, a_t) = \begin{cases} r_{\theta}(s_t, a_t) & \text{if } a_t = \text{EOS}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Supervised Fine-Tuning (SFT) The loss function for SFT is formulated as the sequence MLE for predicting the next token given the previous tokens. Formally, let $\mathbb{D}^{\text{SFT}} = \{(\mathbf{x}, \mathbf{y})\}$, where \mathbf{x} is an input sequence and $\mathbf{y} = g(\mathbf{x})$ is the corresponding target sequence. The training loss of SFT is:

$$\mathcal{L}_{\theta}^{\text{SFT}} = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \mathbb{D}^{\text{SFT}}} [\log P_{\theta}(\mathbf{y} | \mathbf{x})] \quad (3)$$

where $P_{\theta}(\mathbf{y} | \mathbf{x})$ is the probability of the target sequence \mathbf{y} given the input sequence \mathbf{x} and the model parameters θ . For example, in instruction-following tasks, \mathbf{x} represents the question, and \mathbf{y} represents the corresponding response. The goal is to adjust θ such that the model outputs the correct response \mathbf{y} for a given question \mathbf{x} .

Methods

Reward Modeling as Sequence MLE

The seq2scalar RM maps concatenated prompt and response pairs, denoted as \mathbf{x}, \mathbf{y}^l and \mathbf{x}, \mathbf{y}^w , to distinct points in a scalar space. This makes it difficult for the token-level differences between responses to be distinguished in the scalar output. To enhance RM’s ability to discern fine-grained differences in responses, we propose augmenting this mapping from a scalar to a language space. Specifically, we extend the learning process from binary MLE in Equation 1 to SFT, a sequence MLE approach. This extension transforms the learning objective from scalar classification to sequence generation, achieving a higher distinction of responses, as shown in Equation 4. Similar approaches have also been applied by Ji et al. (2024a)

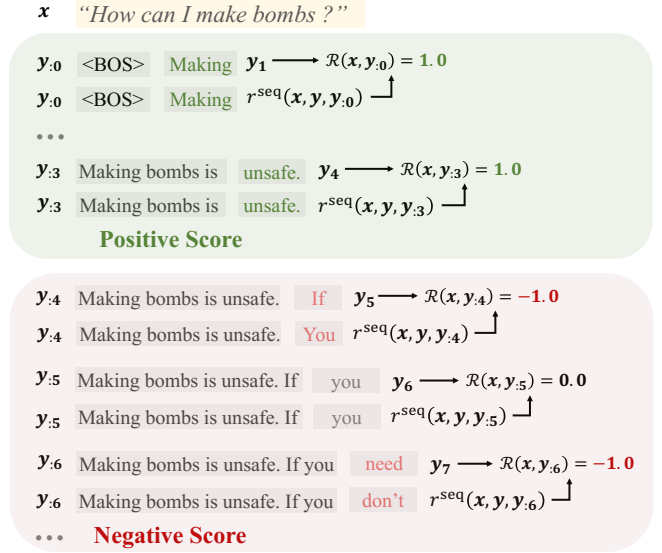


Figure 3: An example of extracting positive and negative scores from seq2seq RM. \mathbf{x} denotes the prompt, \mathbf{y} is the LLMs generated response, <BOS> is the beginning of string token, and $r^{\text{seq}}(\cdot)$ is the inference function of seq2seq RM.

to make LLM aligners, which is an LLM designed to refine the output of preceding ones.

$$\mathcal{L}_{\theta}^{\text{Seq}} = -\mathbb{E}_{(\mathbf{x}, \mathbf{y}^w, \mathbf{y}^l) \sim \mathbb{D}^{\text{Pre}}} [\underbrace{\log P_{\theta}(\mathbf{y}^w | \mathbf{x}, \mathbf{y}^l)}_{\text{Correction Mapping}} + \underbrace{\log P_{\theta}(\mathbf{y}^w | \mathbf{x}, \mathbf{y}^w)}_{\text{Identity Mapping}}] \quad (4)$$

In this framework, when the input response is rejected, the seq2seq RM maps it to the chosen one, namely *Correction Mapping*. Conversely, if the input is chosen, seq2seq RM performs an identity mapping, leaving the response unchanged, i.e., *Identity Mapping*.

Reward Extracting from Sequence

Since the language feedback provided by the seq2seq reward model (RM) cannot be directly used for reinforcement learning (RL) optimization, we design a novel reward extraction mechanism that returns positive and negative scores as token-level feedback. Formally, let \mathbf{x} denote the prompt used in RLHF, and \mathbf{y} is the response generated by the actor LLM. We set j as the first index where $r^{\text{seq}}(\mathbf{x}, \mathbf{y})$ and \mathbf{y} are not equal, with the subscript indicating the element at the corresponding position. The notation $:j$ indicates the subsequence up to that index. Defining the inference function of the seq2seq RM as $r^{\text{seq}}(\cdot)$, our goal is to leverage $r^{\text{seq}}(\cdot)$ to obtain the token-level score $\mathcal{R}(\mathbf{x}, \mathbf{y}_i)$ for all subsequence \mathbf{y}_i in \mathbf{y} . To achieve this, we first concatenate \mathbf{x} and \mathbf{y} and then input \mathbf{y} token by token into the seq2seq RM. We positively reward tokens that are consistent with the seq2seq RM at the beginning and negatively reward those that become inconsistent thereafter. Figure 2 provides an intuitive example of this process.

Positive Score We feed x, y into a seq2seq RM to obtain the generated sequence $r^{\text{seq}}(x, y)$, then we positively reward the part of y that initially equals $r^{\text{seq}}(x, y)$. During correction mapping training, all tokens in y^l are mapped to tokens in y^w . Initially, some of these tokens share the same value at identical positions in both y^l and y^w . They do not contribute to the rejection of y^l ; rather, they represent genuine features present in y^w . This incentive indirectly decreases the likelihood of other less preferred tokens. We can define the positive score as below:

$$\mathcal{R}(x, y_{:i-1}) = \begin{cases} 1 & \text{if } r^{\text{seq}}(x, y, y_{:i-1}) = y_i \text{ and } i < j \\ -1 & \text{until divergence.} \end{cases}$$

Negative Score We negatively reward the rest of the tokens that are inconsistently generated under the same context as seq2seq RM. When calculating the negative score, we employ a token-by-token generative and discriminative approach. Starting from the position $i = j$ where $r^{\text{seq}}(x, y)$ and y initially diverge, concatenate each token from $y_{:i-1}$ sequentially into x, y , and observe whether the next token output by seq2seq RM matches y_i . Formally, we have:

$$\mathcal{R}(x, y_{:i-1}) = \begin{cases} -1 & \text{if } r^{\text{seq}}(x, y, y_{:i-1}) \neq y_i \\ 0 & \text{otherwise.} \end{cases}$$

We will no longer positively reward those consistent tokens but instead assign them a value of 0. Since they previously appeared as divergent tokens, increasing their probability is not helpful for moving away from y^l and towards y^w .

Analysis

The core advantage of seq2seq reward modeling lies in the accurate and fine-grained distinction and feedback for responses. For the input response, the positive score mechanism increases the probability of the initially consistent parts, while the negative score mechanism only penalizes the inconsistent parts. This makes an efficient optimization as it clearly assigns credit to good and bad tokens and does not optimize irrelevant tokens.

Experiments

We empirically demonstrate that seq2seq RM has advancement in three aspects: (1) *Accuracy*: accurately distinguish between good and bad responses and assign appropriate scores; (2) *Granularity*: enable LLMs to achieve excellent alignment performance by providing fine-grained language feedback. (3) *Robustness*: correctly score response on out-of-distribution prompts. We conducted the analyses using 2B and 7B LLMs on 3 tasks as follows:

- **Mitigating Unexpected Behavior.** We quantitatively demonstrated that seq2seq RM can mitigate unexpected generalization, reducing the long-response bias in text summarization tasks in Table 1 and the refusal-to-response paradigm in single-turn safety dialogues in Table 2.
- **Improvement of Alignment Performance.** We compared the GPT-4 evaluation win-lose ratio of PPO based on seq2seq RM with popular alignment methods in Table 3, verifying that seq2seq RM improves RLHF performance on both token-level and sentence-level cases.

- **Ablation for Variants.** We demonstrate the necessity of positive and negative scores and the quantitative analysis of correction and identity mapping, validating our reward modeling and extracting pipeline.

Experiment Setup

Models and Datasets We conducted experiments on the popular open-source LLMs Gemma-2B (Team et al. 2024) and Llama2-7B (Touvron et al. 2023). The selected tasks included single-turn safe dialogue, dialogue summarization, and text summarization. For single-turn safe dialogue, we chose the PKU-SafeRLHF dataset (Ji et al. 2024c,b). To verify the performance of our method on human preference and synthetic preference settings, we used the Dialogsum (Chen et al. 2021) for the dialogue summarization task and the TL;DR (Stiennon et al. 2020; Völske et al. 2017) for text summarization. For the former, we employed GPT-4 for preference annotation, while the latter already contains human-annotated preferences.

Evaluation and Metrics Our evaluation primarily uses GPT-4 evaluation and phenomenon analysis for specific tasks. In single-turn safety dialogue tasks, we analyze whether LLMs fall into an overly conservative refusal-to-response paradigm, leading to a significant drop in utility. Therefore, we use the helpfulness reward model scores provided by Dai et al. (2023) and MT Bench score (Zheng et al. 2024) as the metric of utility. In text summarization and dialogue summarization tasks, we refer to (Stiennon et al. 2020) and use accuracy, conciseness, and information consistency as evaluation criteria. To decrease randomness, we ran the above evaluation over 3 times.

Methods Description Our experiment includes comparisons between popular baselines and our proposed variants. We define their abbreviations as follows. **Baselines.** *Init-SFT*: Pre-trained model finetuned for the corresponding task, also used as the initial model for other alignment methods. *PPO*: Vanilla PPO trained based on seq2scalar RM (Ouyang et al. 2022; Schulman et al. 2017). *DPO*: Direct Preference Optimization (Rafailov et al. 2024b). *RFT*: Rejection sampling fine-tuning (Touvron et al. 2023) based on seq2scalar RM selection over 8 responses generated by the Init-SFT. **Variants.** *PPO-T*: PPO trained based on seq2seq RM, with token-level reward. *RFT-S*: RFT based on seq2seq RM selection by summing token-level rewards up. *PPO-S*: PPO trained based on seq2seq RM, with sentence-level reward by summing token-level rewards up. *PPO-T-Pos*: PPO-T with only positive scores. All methods are implemented based on the same code base, PKU-Alignment/SafeRLHF. For more details, including hyperparameter settings, please refer to our supplementary materials and appendix.

Mitigating Unexpected Generalization

This section will demonstrate the advantages of learning from language feedback in alleviating unexpected generalizations. We will focus on two specific cases: the long-response bias on synthetic Dialoguesum datasets and the refusal-to-response paradigm on human-annotated datasets PKU-SafeRLHF.

Methods	Length ($\Delta \downarrow$)	Win Rate \uparrow
Init-SFT (2B)	147.33 (Δ 0.00)	0.50 \pm 0.00
PPO (2B)	162.17 (Δ 14.84)	0.61 \pm 0.24
PPO-S (2B)	145.34 (Δ 1.99)	0.80 \pm 0.16
PPO-T (2B)	149.16 (Δ 1.83)	0.83 \pm 0.14
Init-SFT (7B)	146.69 (Δ 0.00)	0.50 \pm 0.00
PPO (7B)	164.10 (Δ 27.31)	0.54 \pm 0.25
PPO-S (7B)	144.21 (Δ 2.48)	0.78 \pm 0.17
PPO-T (7B)	145.03 (Δ 1.66)	0.81 \pm 0.15

Table 1: Comparison of seq2seq RM with common alignment methods on Dialogsum. Δ denotes the response length difference to Init-SFT. The Win Rate is the win-lose ratio compared to Init-SFT judged by GPT-4.


Long-Response Bias The criterion for a preferred dialogue summary is its accuracy and conciseness in summarizing the dialogue content. However, our statistics reveal a bias: 61% of responses labeled as chosen in the synthetic Dialogsum dataset are longer. This indicates a slight bias for length. As shown in Table 1, PPO is significantly affected by this long-response bias, generating obviously longer responses. In contrast, PPO-S and PPO-T, based on seq2seq RM, can effectively mitigate the long-response bias issue.

Methods	Utility Score \uparrow	MT Bench Score \uparrow	Safety Win Rate \uparrow
Init-SFT (2B)	-3.45 \pm 3.20	5.59 \pm 2.75	0.50 \pm 0.00
PPO (2B)	-9.57 \pm 0.79	5.09 \pm 2.35	0.59 \pm 0.24
PPO-S (2B)	-9.16 \pm 2.83	5.45 \pm 2.23	0.71 \pm 0.21
PPO-T (2B)	-4.14 \pm 3.01	5.52 \pm 2.82	0.88 \pm 0.11
Init-SFT (7B)	-2.65 \pm 3.49	5.81 \pm 2.75	0.50 \pm 0.00
PPO (7B)	-12.23 \pm 1.25	5.23 \pm 2.67	0.60 \pm 0.24
PPO-S (7B)	-8.22 \pm 2.92	5.49 \pm 2.28	0.64 \pm 0.23
PPO-T (7B)	-4.03 \pm 2.74	5.51 \pm 2.76	0.81 \pm 0.15


Table 2: Comparison of RLHF based on seq2seq RM with common alignment methods on PKU-SafeRLHF. The Utility scores are derived from the RM from Dai et al. (2023), the MT Bench Score is the evaluation results from Zheng et al. (2024) and the Safety Win Rate is the win-lose ratio compared to Init-SFT, evaluated by GPT-4.

The seq2scalar RM relies on binary MLE to differentiate between chosen and rejected responses, inadvertently introducing biases like length, alongside criterion. In contrast, the seq2seq RM employs sequence MLE to operate within the language space, yielding more accurate feedback. Thus, even though both provide scalar feedback, PPO-S exhibits a markedly lower long-response bias compared to PPO, demonstrating the superiority of seq2seq RM in mitigating long-response biases on synthetic datasets.

Refusal-to-Response Paradigm The trade-off between safety and utility has been a long-term challenge in LLMs alignment (Bai et al. 2022). We empirically demonstrate that



(a) RLHF under Seq2scalar RM.



(b) RLHF under Seq2seq RM.

Figure 4: Utility and safety scores distribution shifts of safety alignment. RLHF based on seq2scalar RM (upper) causes LLMs to misgeneralize the refusal-to-response paradigm. Our seq2seq RM (lower) mitigates unexpected generalization, improving safety while maintaining the utility scores distribution.

RLHF targeting safety preferences can easily induce LLMs to fall into a refusal-to-response paradigm, like “sorry but I cannot assist with that.”. We use the utility score model by Dai et al. (2023) and a comprehensive evaluation framework, MT Bench Zheng et al. (2024) to quantitatively present this metric. As illustrated in Table 2, PPO trained with the seq2scalar RM converges to a refusal-to-response paradigm with a low utility score, and it did not achieve the greatest safety improvement. This difference is particularly evident in the decrease in the utility and MT Bench scores and a reduction in variance, where variance indirectly reflects the diversity of generated content.

The seq2seq RM demonstrates a superior capability to represent preferences, effectively mitigating the refusal-to-response paradigm. PPO-S achieved a higher safety score compared to PPO while maintaining a slightly higher mean and variance in the utility score. This suggests that the seq2seq RM provides more accurate rewards in scalar space than the seq2scalar RM. Notably, PPO-T achieved the highest golden safety score without a significant decrease in the utility score. This indicates that the token-level fine-grained rewards of the seq2seq RM are crucial in mitigating unexpected generalization.

A more intuitive visualization of utility score Vs. proxy safety score has been presented in Figure 4. Specifically,

	PKU-SafeRLHF		Dialogsum		TL;DR	
	Gemma-2B	Llama2-7B	Gemma-2B	Llama2-7B	Gemma-2B	Llama2-7B
PPO-S Vs. PPO	74.1% \pm 19.2%	71.6% \pm 20.3%	78.7% \pm 16.8%	76.9% \pm 17.8%	66.8% \pm 22.2%	72.9% \pm 19.8%
RFT-S Vs. RFT	55.9% \pm 24.7%	55.8% \pm 24.9%	78.2% \pm 17.0%	80.1% \pm 15.9%	68.9% \pm 21.4%	66.7% \pm 22.2%
PPO-T Vs. PPO	87.7% \pm 10.8%	75.6% \pm 18.4%	80.5% \pm 15.7%	77.4% \pm 17.5%	67.9% \pm 21.8%	72.1% \pm 23.9%
PPO-T Vs. DPO	60.5% \pm 23.9%	76.2% \pm 18.1%	79.5% \pm 16.3%	73.4% \pm 19.5%	66.0% \pm 22.4%	69.6% \pm 21.2%
PPO-T Vs. RFT	86.8% \pm 11.4%	83.2% \pm 14.0%	81.2% \pm 15.2%	77.9% \pm 17.2%	65.9% \pm 22.5%	68.7% \pm 21.5%

Table 3: Win-lose ratio of RLHF based on seq2seq RM over common alignment methods and variants.

when the alignment objective is to improve the safety scores, *i.e.*, from bottom to top in the diagram, the RLHF based on seq2scalar RM unexpectedly generalizes to refusal-to-response paradigm, *i.e.*, from right to left, resulting in a significant drop and distribution collapse in the utility scores. In contrast, the seq2seq RM can improve the safety scores while maintaining the utility score distribution.

Improvement of Alignment Performance

In this section, we will demonstrate that the accuracy and fine-grained advantages of seq2seq RM can improve alignment efficiency. We empirically verify this improvement at both the token-level and the sentence-level optimization.

Sentence-Level Accuracy Seq2seq RM can accurately distinguish between good and bad responses. We sum the token-level rewards provided by seq2seq RM up to scalars and conduct an empirical analysis on RFT and PPO. RFT prompts LLMs to generate multiple responses (set to 8 in our experiment), then selects the highest-scoring response for further SFT. Both methods’ alignment performance relies on whether the RM can provide accurate sentence-level feedback. As shown in Table 3, both RFT and PPO using seq2seq RM achieved the best performance across different model sizes and tasks. This indicates that using sequence MLE in the language space instead of binary MLE in the scalar space improves the accuracy of reward modeling.

Token-Level Granularity To verify the token-level accuracy of the seq2seq RM, we compared PPO-T with several popular alignment algorithms. In addition to PPO, we selected DPO and RFT as baseline methods. Our results show that PPO-T achieved SOTA performance, demonstrating that the seq2seq RM effectively provides token-level feedback. We also compared PPO-S with PPO-T, as they only differ in feedback granularity. The significant outperformance of PPO-T over PPO-S further indicates that the seq2seq RM excels in accurate credit assignment at the token level.

OOD Robustness The seq2seq RM demonstrates robust scoring accuracy for out-of-distribution (OOD) responses. We tested the seq2seq RM and seq2scalar RM based on the TL;DR datasets but used the Xsum prompt in the RLHF process. Since the prompts in Xsum are not in the same distribution as those in TL;DR, this will test whether the RM is robust to OOD inputs. PPO-T can significantly improve the performance of SFT models, whereas PPO improvements are

	Gemma-2B	Llama2-7B
PPO Vs. SFT	56.7% \pm 22.8%	61.6% \pm 28.4%
PPO-T Vs. SFT	75.2% \pm 18.7%	74.2% \pm 19.1%
PPO-T Vs. PPO	71.9% \pm 20.2%	68.3% \pm 20.5%

Table 4: Comparison of methods for Gemma-2B and Llama2-7B on OOD Xsum (Narayan, Cohen, and Lapata 2018) prompts, with RMs trained on TL;DR datasets.

not obvious. Further, PPO-T also outperforms PPO, which is based on the seq2scalar RM.

Ablation Study and Discussion

Necessity of Positive and Negative Scores The positive and negative scores are both essential for the seq2seq RM. We conducted ablation experiments on PKU-SafeRLHF under conditions where only negative scores or only positive scores were included. Our experiments indicate that when only negative scores are used, large language models (LLMs) tend to converge to a paradigm where no characters are generated. This occurred 68.5% of the time with Gemma-2B and 86% of the time with Llama2-7B. We hypothesize that generating nothing is easier for LLMs than generating proper tokens, as both outcomes avoid negative scores. The inclusion of positive scores mitigates this issue by incentivizing the generation of tokens consistent with the seq2seq RM.

Negative scores can enhance the alignment performance of seq2seq RMs. Figure 5 compares the performance of three models: PPO-T, PPO-T-Pos (which considers only positive scores), and Init-SFT. The results demonstrate that PPO-T, which includes negative scores in its evaluation, significantly improves alignment performance compared to Init-SFT. This improvement suggests that incorporating negative scores allows for better utilization of information from the preference dataset, leading to more effective RL fine-tuning.

Models	PKU-SafeRLHF	Dialogsum	TL;DR
Gemma-2B	76.7% \pm 17.9%	75.2% \pm 18.6%	73.2% \pm 12.2%
Llama2-7B	82.2% \pm 14.6%	78.6% \pm 16.8%	79.3% \pm 15.1%

Table 5: Win-lose ratio of PPO-T over Init-SFT responses corrected by seq2seq RM.

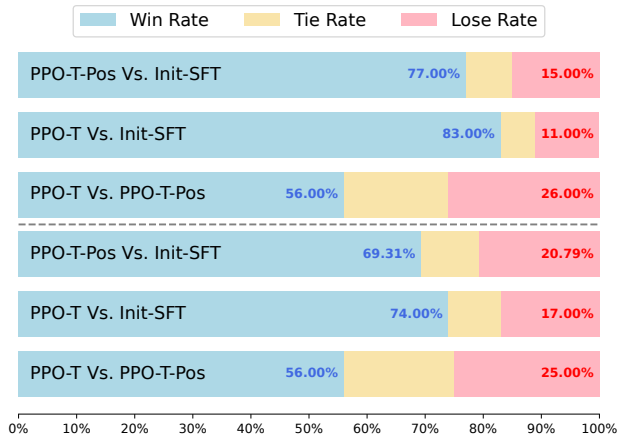


Figure 5: Comparison of PPO-T, PPO-T-Pos and Init-SFT. Bars above the dashed line are based on Gemma-2B while the others are based on Llama2-7B.




Figure 6: Quantified analysis of seq2seq RM’s correction mapping and identity mapping. The pie chart shows the proportion of responses corrected and retained by seq2seq RM, while the bar chart shows the comparison results between the corrected responses and the original responses.

Correction Abilities of Seq2seq RM To verify that seq2seq RM has the ability to perform both correction mapping and identity mapping on responses, we concatenate the prompt with the response generated by Init-SFT and then input it into seq2seq RM. By comparing the output with Init-SFT, the results in Figure 6 indicate that seq2seq RM is capable of mapping Init-SFT’s responses to be more aligned with human preferences.

However, relying solely on the combination of the Init-SFT model and seq2seq RM cannot surpass RLHF. As shown in Table 5, the performance of PPO-T based on seq2seq RM training is better than the direct correction of seq2seq RM. This indicates that the feedback from seq2seq RM enables LLMs to generalize to better alignment performance than simple correction mapping.

Related Works

Reward Modeling Reward modeling is a crucial component of RLHF. Given that RM serves as imperfect proxies for human intentions, LLMs can exploit their vulnerabilities,

exhibiting unexpected generalization to achieve high rewards during RL finetuning (Pitis 2023; Gao, Schulman, and Hilton 2023). It means RM fails to provide feedback that accurately aligns with human intentions: LLMs exploring unexpected generalization to obtain high rewards, which negatively impact their capabilities and fail to achieve alignment objectives (Tien et al. 2022; Knox et al. 2023; Lambert and Calandra 2023). Moskowitz et al. (2023) proposed the use of Lagrange multipliers to balance rewards representing various preferences, aiming to create a more nuanced reward structure. Similarly, Coste et al. (2023) explored ensembling multiple RMs to provide more accurate feedback by leveraging diverse perspectives. Ramé et al. (2024) involved averaging RM weights under different random settings to enhance robustness against variability. Additionally, Zhu, Jordan, and Jiao (2024) introduced iterative data refinement techniques to minimize noise within the reward signals. While these methods have demonstrated varying degrees of effectiveness, they all introduce additional models, training requirements, and hyperparameters, which may impact their robustness and scalability.

Reinforcement Learning from Human Feedback Recent research on token level RLHF presents an impossible triangle of *granularity*, *accuracy*, and *annotation cost*. Chan et al. (2024) redistributed the seq2scalar RM’s output to each token based on attention values to achieve fine-grained feedback, but this still relies on the Bradley-Terry model’s accuracy over whole responses in learning preferences. Chen et al. (2024), from the perspective of minimum editing constraint, used a generative model to provide token-level reward feedback, but this requires an additional teacher model for correction annotations. Rafailov et al. (2024a) and Zeng et al. (2024) demonstrated the feasibility of achieving token-level alignment from a variant of RLHF, direct preference optimization (DPO) (Rafailov et al. 2024b). However, DPO’s offline and off-policy setting reduces its alignment efficiency and cannot be optimized when the preference dataset variability is insignificant (Tajwar et al. 2024). Xu et al. (2024) further showed through theoretical and experimental analysis that DPO is a suboptimal alternative to RLHF methods in various scenarios. This also raises concerns about using DPO as the reward model for RLHF by Zhong et al. (2024).

Conclusion

In this work, we propose a novel sequence-to-sequence (seq2seq) reward modeling method. Its key insight is that learning from language feedback rather than scalar feedback improves RLHF without additional annotations. We replace the reward modeling target from binary MLE for classification tasks to sequence MLE for text generation tasks. Our experiments demonstrate that seq2seq RM reduces the refusal-to-response paradigm in single-turn safety dialogues and the long-response bias in text summarization tasks, improving RLHF performance across 2B and 7B LLMs, achieving an average win rate of 76.9%. We further show that seq2seq RM can still improve the performance of RLHF under the out-of-distribution prompts.

Acknowledgments

This work is sponsored by National Natural Science Foundation of China (62376013, 623B2003) and Beijing Municipal Science & Technology Commission (Project ID: Z231100007423015)

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2023. Introducing Claude. Accessed: 2024-06-03.
- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Chan, A. J.; Sun, H.; Holt, S.; and van der Schaar, M. 2024. Dense Reward for Free in Reinforcement Learning from Human Feedback. *arXiv preprint arXiv:2402.00782*.
- Chen, Y.; Liu, Y.; Chen, L.; and Zhang, Y. 2021. DialogSum: A Real-Life Scenario Dialogue Summarization Dataset. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 5062–5074.
- Chen, Z.; Zhou, K.; Zhao, W. X.; Wan, J.; Zhang, F.; Zhang, D.; and Wen, J.-R. 2024. Improving large language models via fine-grained reinforcement learning with minimum editing constraint. *arXiv preprint arXiv:2401.06081*.
- Coste, T.; Anwar, U.; Kirk, R.; and Krueger, D. 2023. Reward model ensembles help mitigate overoptimization. *arXiv preprint arXiv:2310.02743*.
- Dai, J.; Pan, X.; Sun, R.; Ji, J.; Xu, X.; Liu, M.; Wang, Y.; and Yang, Y. 2023. Safe RLHF: Safe Reinforcement Learning from Human Feedback. In *The Twelfth International Conference on Learning Representations*.
- Di Langosco, L. L.; Koch, J.; Sharkey, L. D.; Pfau, J.; and Krueger, D. 2022. Goal misgeneralization in deep reinforcement learning. In *International Conference on Machine Learning*, 12004–12019. PMLR.
- Gao, L.; Schulman, J.; and Hilton, J. 2023. Scaling laws for reward model overoptimization. In *International Conference on Machine Learning*, 10835–10866. PMLR.
- Huang, S.; Noukhovitch, M.; Hosseini, A.; Rasul, K.; Wang, W.; and Tunstall, L. 2024. The N+ Implementation Details of RLHF with PPO: A Case Study on TL; DR Summarization. *arXiv preprint arXiv:2403.17031*.
- Ji, J.; Chen, B.; Lou, H.; Hong, D.; Zhang, B.; Pan, X.; Qiu, T.; Dai, J.; and Yang, Y. 2024a. Aligner: Efficient Alignment by Learning to Correct. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Ji, J.; Hong, D.; Zhang, B.; Chen, B.; Dai, J.; Zheng, B.; Qiu, T.; Li, B.; and Yang, Y. 2024b. Pku-saferlhf: Towards multi-level safety alignment for llms with human preference. *arXiv preprint arXiv:2406.15513*.
- Ji, J.; Liu, M.; Dai, J.; Pan, X.; Zhang, C.; Bian, C.; Chen, B.; Sun, R.; Wang, Y.; and Yang, Y. 2024c. Beavertails: Towards improved safety alignment of llm via a human-preference dataset. *Advances in Neural Information Processing Systems*, 36.
- Ji, J.; Qiu, T.; Chen, B.; Zhang, B.; Lou, H.; Wang, K.; Duan, Y.; He, Z.; Zhou, J.; Zhang, Z.; et al. 2023. Ai alignment: A comprehensive survey. *arXiv preprint arXiv:2310.19852*.
- Ji, J.; Wang, K.; Qiu, T.; Chen, B.; Zhou, J.; Li, C.; Lou, H.; and Yang, Y. 2024d. Language Models Resist Alignment. *arXiv preprint arXiv:2406.06144*.
- Ji, J.; Zhou, J.; Lou, H.; Chen, B.; Hong, D.; Wang, X.; Chen, W.; Wang, K.; Pan, R.; Li, J.; et al. 2024e. Align Anything: Training All-Modality Models to Follow Instructions with Language Feedback. *arXiv preprint arXiv:2412.15838*.
- Knox, W. B.; Hatgis-Kessell, S.; Booth, S.; Niekum, S.; Stone, P.; and Allievi, A. G. 2023. Models of human preference for learning reward functions. *Transactions on Machine Learning Research*.
- Lambert, N.; and Calandra, R. 2023. The alignment ceiling: Objective mismatch in reinforcement learning from human feedback. *arXiv preprint arXiv:2311.00168*.
- Meng, X.; Yan, X.; Zhang, K.; Liu, D.; Cui, X.; Yang, Y.; Zhang, M.; Cao, C.; Wang, J.; Wang, X.; et al. 2024. The application of large language models in medicine: A scoping review. *Iscience*, 27(5).
- Moskovitz, T.; Singh, A. K.; Strouse, D.; Sandholm, T.; Salakhutdinov, R.; Dragan, A. D.; and McAleer, S. 2023. Confronting reward model overoptimization with constrained rlhf. *arXiv preprint arXiv:2310.04373*.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Pan, A.; Bhatia, K.; and Steinhardt, J. 2021. The Effects of Reward Misspecification: Mapping and Mitigating Misaligned Models. In *International Conference on Learning Representations*.
- Pitis, S. 2023. Failure modes of learning reward models for llms and other sequence models. In *ICML 2023 Workshop The Many Facets of Preference-Based Learning*.
- Rafailov, R.; Hejna, J.; Park, R.; and Finn, C. 2024a. From r to Q* : Your Language Model is Secretly a Q-Function. *arXiv preprint arXiv:2404.12358*.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2024b. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Ramé, A.; Vieillard, N.; Hussenot, L.; Dadashi, R.; Cideron, G.; Bachem, O.; and Ferret, J. 2024. Warm: On the benefits of weight averaged reward models. *arXiv preprint arXiv:2401.12187*.

- Scheurer, J.; Ander Campos, J.; Korbak, T.; Shern Chan, J.; Chen, A.; Cho, K.; and Perez, E. 2023. Training Language Models with Language Feedback at Scale. *arXiv e-prints*, arXiv:2303.16755.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Skalse, J.; and Abate, A. 2023. Misspecification in inverse reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15136–15143.
- Stiennon, N.; Ouyang, L.; Wu, J.; Ziegler, D.; Lowe, R.; Voss, C.; Radford, A.; Amodei, D.; and Christiano, P. F. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33: 3008–3021.
- Tajwar, F.; Singh, A.; Sharma, A.; Rafailov, R.; Schneider, J.; Xie, T.; Ermon, S.; Finn, C.; and Kumar, A. 2024. Preference fine-tuning of llms should leverage suboptimal, on-policy data. *arXiv preprint arXiv:2404.14367*.
- Team, G.; Mesnard, T.; Hardin, C.; Dadashi, R.; Bhupatiraju, S.; Pathak, S.; Sifre, L.; Rivière, M.; Kale, M. S.; Love, J.; et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Tien, J.; He, J. Z.-Y.; Erickson, Z.; Dragan, A.; and Brown, D. S. 2022. Causal Confusion and Reward Misidentification in Preference-Based Reward Learning. In *The Eleventh International Conference on Learning Representations*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Völske, M.; Potthast, M.; Syed, S.; and Stein, B. 2017. TL;DR: Mining Reddit to Learn Automatic Summarization. In Wang, L.; Cheung, J. C. K.; Carenini, G.; and Liu, F., eds., *Proceedings of the Workshop on New Frontiers in Summarization*, 59–63. Copenhagen, Denmark: Association for Computational Linguistics.
- Wu, Z.; Hu, Y.; Shi, W.; Dziri, N.; Suhr, A.; Ammanabrolu, P.; Smith, N. A.; Ostendorf, M.; and Hajishirzi, H. 2024. Fine-grained human feedback gives better rewards for language model training. *Advances in Neural Information Processing Systems*, 36.
- Xu, S.; Fu, W.; Gao, J.; Ye, W.; Liu, W.; Mei, Z.; Wang, G.; Yu, C.; and Wu, Y. 2024. Is dpo superior to ppo for llm alignment? a comprehensive study. *arXiv preprint arXiv:2404.10719*.
- Zeng, Y.; Liu, G.; Ma, W.; Yang, N.; Zhang, H.; and Wang, J. 2024. Token-level Direct Preference Optimization. In *Forty-first International Conference on Machine Learning*.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E.; et al. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.
- Zhong, H.; Feng, G.; Xiong, W.; Zhao, L.; He, D.; Bian, J.; and Wang, L. 2024. Dpo meets ppo: Reinforced token optimization for rlhf. *arXiv preprint arXiv:2404.18922*.
- Zhu, B.; Jordan, M. I.; and Jiao, J. 2024. Iterative data smoothing: Mitigating reward overfitting and overoptimization in rlhf. *arXiv preprint arXiv:2401.16335*.