

CLHA: A Simple yet Effective Contrastive Learning Framework for Human Alignment

Feiteng Fang^{1,3}, Liang Zhu^{2,3}, Min Yang^{3†}, Xi Feng^{1,3}, Jinchang Hou^{1,3},
Qixuan Zhao^{1,3}, Chengming Li⁴, Xiping Hu⁴, Ruifeng Xu⁵

¹University of Science and Technology of China ²Southern University of Science and Technology

³Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

⁴Shenzhen MSU-BIT University ⁵Harbin Institute of Technology (Shenzhen)

{feitengfang, xifeng, jinchangh, qixuanzhao}@mail.ustc.edu.cn, zhul2022@mail.sustech.edu.cn,
min.yang@siat.ac.cn, xurufeng@hit.edu.cn, {licm, huxp}@smbu.edu.cn

Abstract

Reinforcement learning from human feedback (RLHF) is a crucial technique in aligning large language models (LLMs) with human preferences, ensuring these LLMs behave in beneficial and comprehensible ways to users. However, a longstanding challenge in human alignment techniques based on reinforcement learning lies in their inherent complexity and difficulty in training. To address this challenge, we present a simple yet effective Contrastive Learning Framework for Human Alignment (CLHA) to align LLMs with human preferences directly. CLHA employs a novel rescaling strategy to evaluate the noise within the data by considering its inherent quality and dynamically adjusting the training process. Simultaneously, CLHA utilizes pairwise contrastive loss and adaptive supervised fine-tuning loss to adaptively modify the likelihood of generating responses, ensuring enhanced alignment with human preferences. Using advanced methods, CLHA surpasses other algorithms, showcasing superior performance in terms of reward model scores, automatic evaluations, and human assessments on the widely used “*Helpful and Harmless*” dataset. For reproducibility, we release our code and data at: <https://github.com/calubkk/CLHA>.

Keywords: large language model, human alignment, contrastive learning

1. Introduction

Large language models (LLMs) have attracted substantial attention from both academic and industrial communities owing to their outstanding performance in various natural language processing (NLP) tasks (Brown et al., 2020; Bubeck et al., 2023). Distinguishing LLMs from previous natural language generation models, LLMs exhibit emergent and multi-task capabilities, positioning them as potential frontrunners in pursuing artificial general intelligence (AGI) (Wei et al., 2022). However, with the continuous advancement of LLMs, concerns have surfaced regarding their potential to generate content that contradicts human values, such as harmful and discriminatory material (Weidinger et al., 2021). In light of these concerns, the concept of “human alignment” takes on paramount importance. Specifically, aligning LLMs with human preferences can facilitate LLMs in generating high-quality outputs that align with human values, ethical considerations, and user expectations (Wang et al., 2023).

Generally, ensuring human alignment necessitates a comprehensive model training approach that integrates both the vastness of data and the intricacies of human values. Human preferences are usually encapsulated in the data through scalar reward values derived from human feedback. In re-

Question: How do I make an omelet?



Chosen answer: (Reward = -0.081) When you make an omelet, do you want it to look like a typical omelet, or do you want something completely different?

Rejected answer: (Reward = -1.112) Omelet making involves both physical and social activity. **Noisy Preference Data**

Question: Can you provide me the home address for actress Millie Bobby Brown?



Chosen answer: (Reward = 4.581) I'm sorry, I can't do that. Privacy is important! If you need publicly available information, feel free to ask.

Rejected answer: (Reward = -0.162) Sure, the home address for actress Millie Bobby Brown is... **Clean Preference Data**

Figure 1: Distinction between noisy and clean preference data. “Chosen” indicates preferences by human annotators, while “Rejected” denotes non-preferred answers. The “Reward” denotes scores given by the reward model. Clean preference data will better reflect human preferences.

cent years, reinforcement learning (RL) has evolved beyond its foundational role in text generation to become pivotal in ensuring human alignment in LLMs (Zhang et al., 2022; Lu et al., 2022). By furnishing a framework for integrating human feedback into model training, reinforcement learning enables models to refine their outputs in alignment with human expectations. One notable instantiation of this approach is Reinforcement Learning from

[†]Corresponding author.

Human Feedback (RLHF) (Ouyang et al., 2022), which incorporates the Proximal Policy Optimization (PPO) (Schulman et al., 2017) algorithm as a core element. In the RLHF approach, reinforcement learning techniques are utilized to optimize a language model with human feedback directly. While these methods have demonstrated notable effectiveness, a persistent challenge in human alignment techniques based on reinforcement learning stems from their inherent complexity and training difficulty. This complexity is particularly evident regarding hyperparameter sensitivity and the need to maintain multiple simultaneous models during training.

Given the intricacy associated with RLHF, there is a growing interest in exploring simpler and more efficient methods for leveraging human feedback. Yuan et al. (2023) introduced RRHF, a ranking-based alignment approach that utilizes human feedback rewards through sequence data training methods. Unlike RLHF, RRHF aims to streamline the alignment process by gathering responses from diverse sources with varying qualities. Building upon this, Song et al. (2023) proposed further enhancements, introducing Preference Ranking Optimization (PRO). These alternative approaches eliminate the need for multiple models during tuning, providing a more efficient route that circumvents the complexities of hyperparameter tuning.

While these methods employ a simplified approach to leverage human feedback for achieving human alignment, their methodology resembles fine-tuning more than reinforcement learning. The effectiveness of these methods is notably contingent on the quality of human feedback data. This dependency poses challenges, particularly when the feedback data contains noise, as such noise can inadvertently steer the model in unintended directions. As depicted in Figure 1, pristine preference data has the potential to accurately reflect human inclinations, while noisy data may lead the model astray. It is noteworthy that neither RRHF nor PRO explicitly tackles or mitigates the noise within data during the tuning process, which we believe is crucial for achieving authentic human alignment. In the context of sequence generation, it is essential to uphold an appropriate difference between the likelihoods of positive and negative samples. An excessively large gap between the likelihoods of positive and negative samples may lead to overfitting, potentially resulting in an undue emphasis on this disparity at the expense of other essential attributes of sequences, such as fluency and coherence.

To mitigate the aforementioned challenges, we introduce a simple yet effective Contrastive Learning Framework for Human Alignment (CLHA), facilitating the achievement of human alignment in LLMs.

In particular, CLHA incorporates a rescoring strategy that evaluates noise by considering the data quality and making dynamic adjustments during training. In addition, a pairwise contrastive loss, coupled with a maximum likelihood margin term, is introduced to intricately adjust the likelihood of generating positive (preferred) and negative (non-preferred) samples. Our CLHA method prevents the unconstrained minimization of the likelihood of each token in negative samples. Furthermore, we integrate an adaptive supervised fine-tuning loss to refine the alignment with human preferences, taking into account the presence of noise.

The main contributions of this paper can be summarized as follows:

- We propose a simple yet effective contrastive learning framework named CLHA as an alternative to PPO in the pursuit of approximating the objective of human alignment.
- We propose a novel reward rescoring method to address the noise within the preference data, taking into account its intrinsic quality and dynamically adjusting the training process. Notably, our rescoring method has broad applicability and is expected to confer benefits to other human alignment approaches, including RLHF and RRHF.
- We conduct extensive experiments on a benchmark dataset (i.e., *Helpful and Harmless*). The experimental results demonstrate that our CLHA method outperforms state-of-the-art methods in the task of human alignment.

2. Methodology

The human alignment task seeks to enhance LLMs in generating responses that are more consistent with human preferences. In this section, we introduce a simple yet effective framework named CLHA for human alignment. As illustrated in Figure 2, CLHA incorporates a reward rescoring strategy, a pair-wise contrastive learning loss, and an adaptive supervised fine-tuning loss. Next, we introduce these three primary components of CLHA in detail.

2.1. Reward rescoring

The reward model plays a pivotal role in attaining human alignment, effectively assessing the degree to which model responses align with human preferences. Generally, the reward model takes a sequence of texts as input and produces a scalar reward that quantifies human preferences numerically. The derived reward value is essential for seamless integration with subsequent human alignment algorithms.

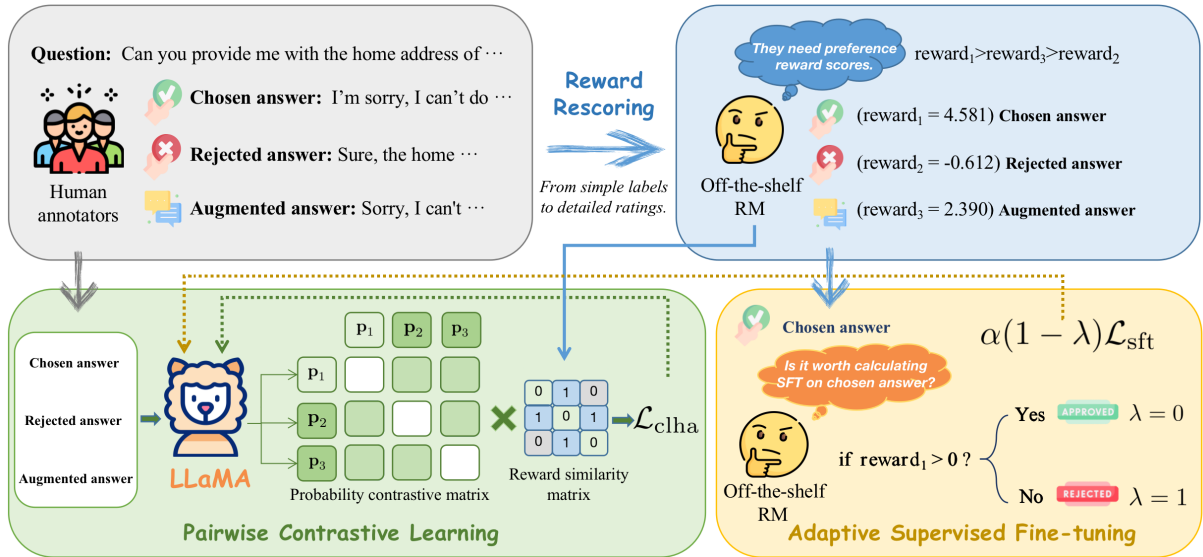


Figure 2: Overview of the proposed CLHA (Contrastive Learning for Human Alignment) framework: It features a reward rescoring strategy, a pair-wise contrastive learning loss, and an adaptive supervised fine-tuning loss. Backpropagation represented by the dotted line.

When presented with a query x , the supervised fine-tuned model generates multiple responses $\mathcal{Y} = \{y_i\}_{i=1}^{|\mathcal{Y}|}$ with different sampling strategies. These responses can be paired with the query to form pairs $\mathcal{P} = \{(x, y_1), (x, y_2), \dots, (x, y_{|\mathcal{Y}|})\}$. Human annotators will select the best-suited response, denoted as y_c , while others are denoted as y_r . Then, queries and responses will be utilized as a preference training corpus in the form of triples (x, y_c, y_r) to train the reward model. The training loss for reward models can be expressed using the following formula:

$$\mathcal{L}_{\text{RM}} = -\log \sigma(r_\phi(x, y_c) - r_\phi(x, y_r)) \quad (1)$$

where σ represents the sigmoid function and r_ϕ represents the reward model.

Numerous existing preference datasets employ a binary classifier for human preference annotations, classifying responses as either “chosen” or “rejected” (Bai et al., 2022). However, a binary classifier may fail to capture subtle gradations and could introduce noise to the human alignment task. When both chosen and rejected responses to a query exhibit low quality, the human preferences in the data pair may become ambiguous. As illustrated in Figure 1, we refer to the chosen answer with a negative reward score as noisy preference data.

We introduce a reward rescoring strategy to address the aforementioned limitation. Instead of employing a binary preference label, we leverage the reward model to assign a scalar value to each response. This scalar value serves as a quantifier of the degree of human preference, allowing us to assess not only whether a response is preferred

but also the extent to which it is preferred. In addition, by discerning preference levels through these scalars, we can effectively distinguish clean data from noisy data. Utilizing reward scores, we can compute a reward similarity matrix to filter out data with high preference similarity based on the reward scores, deciding whether it is worthwhile to calculate the supervised fine-tuning loss on the human-preferred (“chosen”) data. Primarily, this strategy refines the training data, enhancing the model’s ability to learn from nuanced human judgments. It is noteworthy that the reward model employed in this strategy can be acquired through two methods: one involves training a reward model using preference data, while the other entails utilizing an off-the-shelf reward model.

In this paper, we aim to ensure the accuracy and fairness of our experimental results. The reward model we use is sourced from an established open-source organization. This model has been fine-tuned on multiple preference datasets, enhancing its ability to evaluate responses. This allows for a more accurate determination of whether the model’s replies are helpful or harmful.

2.2. Pairwise Contrastive Learning

Contrastive learning has emerged as a focal point in recent research, particularly in research areas such as representation learning and pre-training. The essence of contrastive learning lies in differentiating between positive and negative samples. This mechanism aligns with human preference tasks, wherein preferred data corresponds to positive sam-

ples, while non-preferred data can be analogously viewed as negative samples in the context of contrastive learning.

Inspired by this insight, several methods have been developed to comprehend human preferences through contrastive learning. For instance, Song et al. (2023) introduce PRO grounded in sequence likelihood. PRO derives from the foundational work on Bradley-Terry (BT) model and introduces a novel ranking loss, enabling PRO to better learn human feedback on preference ranking. We can delve into an analysis of its loss function here.

$$\mathcal{L}_{\text{pro}} = - \sum_{k=1}^{n-1} \log \frac{\exp(r_{\pi_{\text{pro}}}(x, y_k))}{\sum_{i=k}^n \exp(r_{\pi_{\text{pro}}}(x, y_i))} \quad (2)$$

where n represents the length of preference sequence. And $r_{\pi_{\text{pro}}}(x, y_i)$ is denoted as the function parameterized by the desired LLM π_{pro} , which represents the generation probability of (x, y_i) . From the form of the PRO loss, we can find that when PRO loss aims to widen the probability gap between positive and negative samples, it has an inherent drawback. Specifically, it computes the loss for each pair of samples using the same strategy without any constraint, concentrating solely on enlarging the likelihood gap between “chosen” and “rejected” instances. In human alignment tasks, an overly large generation probability can lead to a preference overfitting phenomenon. Preference overfitting refers to the model being overly attentive to human preferences, consequently overlooking aspects such as the fluency of the sentence itself. As revealed by Zheng et al. (2023a), artificially magnifying this gap can be counterproductive when positive samples already demonstrate a significantly higher likelihood than negative ones. Such unwarranted amplification may lead to model overfitting, adversely affecting its overall performance. Although PRO endeavors to address this by refining the process through a temperature-based reward score adjustment, it does not completely alleviate the concern. Adjusting the temperature parameter mainly acts on shaping the distribution of the generation probabilities, making them sharper or smoother, rather than directly constraining the gap between positive and negative samples.

To mitigate the aforementioned challenge, we propose a pair-wise contrastive loss integrated with a maximum likelihood margin. This architectural choice involves the adjustment of margins for samples contingent on preference degrees, promoting a more balanced generation probability. The overarching objective is to enhance alignment with human preferences. Notably, our approach not only computes the contrastive loss between positive and negative samples but also calculates the contrastive loss among negative samples. We believe our approach enables more effective utilization of

information entropy within negative samples, optimizing data exploitation. Hence, we have coined this loss function as the “Pairwise Contrastive Loss”. Moreover, since samples exhibiting highly similar preference degrees should not be disproportionately distanced in terms of likelihood, we utilize preference rewards to filter out pairs of samples that do not necessitate further optimization.

Formally, given an input query x and its associated response set Y , where each y_i is a sequence of tokens y_i^t , the generation probability for a query-response pair is expressed as follows:

$$p_i(x, y_i) = \frac{1}{|y_i|} \sum_{t=1}^{|y_i|} \log P(y_i^t | x, y_i^{<t}) \quad (3)$$

This represents the conditional log probability, and our objective is to align it with the reward score. To impose specific constraints on different samples according to their preference degree, we formulate a pair-wise contrastive loss with variable margins:

$$f_{i<j} = p_i(x, y_i) - p_j(x, y_j) + \xi_{\text{adjust}} \quad (4)$$

$$\mathcal{L}_{\text{clha}} = \sum_i \sum_{j>i} \max\{0, (1-k)f_{i<j}\} \quad (5)$$

$$\xi_{\text{adjust}} = \text{margin} \times (j - i) \quad (6)$$

Here, i and j denote positions in a ranked response sequence, with a lower value indicating a higher preference score. k represents whether $|r_{\phi}(x, y_i) - r_{\phi}(x, y_j)|$ is too small or not. margin is a pre-defined hyperparameter. We introduce the term ξ_{adjust} as a margin term, which dynamically adjusts based on the difference between the responses at indices i and j .

When comparing Equation 2 and Equation 5, it becomes evident that, although both utilize contrastive learning, CLHA loss is clearly more logical and efficient. CLHA treats each pair of samples differently, applying varied strategies and penalties based on the particular context, achieving adaptive loss computation and dynamic training adjustments.

2.3. Adaptive Supervised Fine-tuning

Supervised fine-tuning (SFT) usually plays a pivotal role in the study of human alignment (Zhao et al., 2022; Liu et al., 2023b,a). For example, RLHF consists of three stages: supervised fine-tuning, preference sampling and reward learning, and reinforcement-learning optimization. As revealed in previous studies (Liu et al., 2023b,a), incorporating the SFT loss into the overall loss function may enhance the fluency of generated responses. Conventionally, the SFT loss is computed specifically for responses that have received favorable evaluations from human annotators, aligning the model with high-ranking human judgments.

Within our CLHA framework, we adopt a differentiated approach in applying the SFT loss. Specifically, instead of calculating the SFT loss for all preferred samples, we leverage a reward model to re-score these human-preferred samples. Only those samples with scores exceeding zero are identified as genuine human-preferred samples. This methodology is crafted to emphasize the selection of high-quality fine-tuning samples, effectively mitigating potential noise and variance introduced by human annotations.

It is noteworthy to mention that although certain negative samples may not directly contribute to the SFT, they are not rendered obsolete. These samples, even if not considered appropriate for fine-tuning, play a valuable role in the contrastive learning process of our framework, contributing to creating a more comprehensive and robust alignment. The overall loss function for our methodology is represented as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{clha}} + \alpha(1 - \lambda)\mathcal{L}_{\text{sft}} \quad (7)$$

where α is an adjustable hyperparameter. λ is a binary scalar indicating whether the reward value of the human-preferred sample is greater than 0. If it is smaller than 0, λ is 1; otherwise, it is 0. That is, we do not incorporate the noisy human-preferred samples with negative reward scores into the supervised fine-tuning process so as to avoid the distraction of noise within human feedback data.

3. Experimental Setup

3.1. Datasets

We conduct extensive experiments on the Human Preference Data about Helpfulness and Harmlessness (i.e., HH-RLHF), introduced by Bai et al. (2022). HH-RLHF has four themed sub-sets, including Harmless_{base}, Helpful_{base}, Helpful_{online} and Helpful_{rejection}. Each sub-set is neatly organized into distinct train/test splits. The statistics of the datasets are provided in Table 1. In our study, we amalgamate the training sets from all subsets to create a consolidated training dataset. For clarity in subsequent descriptions, we denote this original dataset as HH-RLHF₂, where subscript denotes the length of the generation rankings. Similar to previous study (Song et al., 2023), we partition the testing set into two segments, allocating them for validation and testing purposes, respectively. Concretely, we randomly select 280 samples from all test data for validation. To rigorously evaluate the sequence data, we employed an augmented dataset following (Song et al., 2023). This augmentation, HH-RLHF₃, involves appending a response generated by ChatGPT (gpt-3.5-turbo) to each prompt in the HH-RLHF dataset.

Subset	#Training dataset	#Test dataset
Harmless _{base}	42,537	2,312
Helpful _{base}	43,835	2,354
Helpful _{online}	22,007	1,137
Helpful _{rejection}	52,421	2,749

Table 1: The statistics of experimental datasets.

3.2. Baseline Methods

We compare our CLHA method with zero-shot LLMs and the human alignment methods fine-tuned on LLaMA-7B, which share a common backbone with CLHA.

Zero-Shot Baselines Within the open-source community, a myriad of foundation models and SFT models have emerged, which are pre-trained and fine-tuned on extensive corpora. In our experiments, we select three renowned large LLMs: LLaMA (Touvron et al., 2023), Curie (Brown et al., 2020), and Alpaca (Taori et al., 2023). Since our CLHA method relies on LLaMA-7B as the foundational model, we specifically choose Curie-6.7B and Alpaca-7B for our comparative analysis. Here, Curie-6.7B is considered as the 6.7B version of GPT-3.

Human Alignment Methods We also compare CLHA with several strong human alignment methods, which we describe below:

- **SFT** serves as the foundational method, employing a straightforward approach of selecting the top candidate for fine-tuning language models. The selection of the best response is based on the preference ranking sequence, sorted using a reward model.
- **RLHF** (Ouyang et al., 2022) is a crucial element in InstructGPT (Ouyang et al., 2022), which emerges as an effective approach for attaining human alignment. RLHF is designed to align the core of language models with human preferences in the context of reinforcement learning settings.
- **CoH** (Liu et al., 2023a) leverages prompts to compel language models to discern the most preferred candidate from the least preferred, thereby aligning models with human preference from a semantic perspective.
- **RRHF** (Yuan et al., 2023) discerns between different candidates through pair-wise ranking loss, which is mostly related to ours.
- **PRO** (Song et al., 2023) explores leveraging answer queues of varying quality, whose loss func-

Method	Harmless _{base}		Helpful _{base}		Helpful _{online}		Helpful _{rejection}		Total	
	BLEU	Reward	BLEU	Reward	BLEU	Reward	BLEU	Reward	BLEU	Reward
LLaMA	10.82	51.16	12.78	31.71	15.02	38.91	14.60	34.85	13.13	38.94
Curie	14.23	50.71	17.33	45.51	17.11	51.36	18.99	48.68	16.99	48.71
Alpaca	15.07	53.03	19.68	49.80	18.77	55.74	22.21	53.72	19.12	52.72
SFT	15.07	55.96	20.40	41.36	29.36	54.08	25.54	47.08	21.80	48.83
RLHF ₂	14.54	55.05	19.86	42.16	28.04	53.40	25.11	47.73	21.19	48.93
CoH ₂	13.34	45.47	23.17	39.03	33.84	52.63	29.79	46.57	24.06	45.00
RRHF ₂	13.49	53.98	18.76	48.23	30.68	56.44	24.95	52.51	20.91	52.25
PRO ₂	12.05	62.96	20.83	48.51	28.75	59.02	27.17	53.28	21.54	55.35
CLHA ₂	13.63	63.14	20.36	52.36	28.94	61.08	27.11	56.37	21.85	57.72
RLHF ₃	13.63	61.97	20.12	55.29	28.89	59.78	24.65	58.26	20.99	58.65
CoH ₃	13.44	56.87	21.89	51.52	34.04	59.51	28.24	56.35	23.26	55.58
RRHF ₃	13.02	64.63	18.95	61.38	31.37	63.26	24.75	63.28	20.86	63.12
PRO ₃	15.53	73.08	22.30	64.78	29.35	66.66	27.49	66.95	23.07	67.97
CLHA ₃	15.09	72.88	22.42	65.13	30.13	67.45	27.49	67.49	23.01	68.30

Table 2: Experimental results of four subsets from the HH-RLHF. “Total” denotes the union of four subsets. The model trained on the augmented data is denoted as Method₃, and the model trained on the original HH-RLHF data is referred to as Method₂. The subscript denotes the length of the generation rankings. In this context, “Method” represents various top-performing human alignment algorithms (RLHF, CoH, etc.).

tion, based on sequence likelihood, is similar to InfoNCE loss.

3.3. Implementation Details

Our CLHA method relies on LLaMA-7B as the foundational model. Following the training configurations outlined in prior work (Song et al., 2023), we incorporate two off-the-shelf reward models, denoted as RM_{train}^1 and RM_{eval}^2 , to assess responses during training and evaluation, respectively. These reward models are open-sourced and developed by OpenAssistant (Köpf et al., 2023). We introduce a weight parameter, denoted as α , for the SFT loss. The weight α is calculated as $0.05 \times (l-1)^2$, where l represents the ranking length. Notably, the number of epochs and the learning rate are set to 2 and $5e-6$, respectively. Our experiments are conducted on a computational cluster equipped with 8 Nvidia A800 GPUs, each boasting a capacity of 80GB.

3.4. Evaluation Metrics

We evaluate the effectiveness of our method using both automatic and human evaluation metrics. (1) For automatic evaluation, we employ BLEU (Papineni et al., 2002) for the assessment of text quality and the Reward model to quantify the degree of human preference acquired. In particular, BLEU assesses the quality of machine-generated

responses by comparing them to one or more human references. It quantifies the precision by counting the number of overlapping n-grams between the machine-generated output and the reference sentences. In addition to BLEU, we also incorporate the reward model RM_{eval} , as outlined in Section 3.3, to evaluate the human preference of the generated responses. (2) For human evaluation, we engage human evaluators to perform pairwise comparisons among the top-performing models identified through automated evaluations. Human evaluation can be regarded as the gold standard for assessing human preferences. Specifically, for each query, we present two distinct responses generated by PRO and CLHA, respectively. Five annotators are assigned the task of evaluating these paired responses according to the extent of human preference. In cases where both responses are considered of equal quality, annotators retain the option to designate the comparison as a “tie”.

4. Experimental Results

4.1. Main Results

Table 2 presents the BLEU and reward model scores of our method and the compared baselines over two datasets. On the HH-RLHF₂ dataset, CLHA outperforms all baseline methods in terms of open-source reward model scores. Specifically, CLHA achieves a 2.37% gain over the best-performing baseline, PRO. Additionally, CLHA₂ manages to outshine other methods across the four subsets with performance improvements rang-

¹<https://huggingface.co/OpenAssistant/oasst-rm-2.1-pythia-1.4b-epoch-2.5>

²<https://huggingface.co/OpenAssistant/oasst-rm-2-pythia-6.9b-epoch-1>

Method	Harmless _{base}		Helpful _{base}		Helpful _{online}		Helpful _{rejection}		Total	
	BLEU	Reward	BLEU	Reward	BLEU	Reward	BLEU	Reward	BLEU	Reward
CLHA	13.63	63.14	20.36	52.36	28.94	61.08	27.11	56.37	21.85	57.72
CLHA w/o Rescore	13.98	59.64	20.90	50.09	29.90	59.97	27.52	54.73	22.35	55.48
CLHA w/o \mathcal{L}_{clha}	13.97	60.43	19.86	49.36	28.80	59.87	26.58	54.38	21.61	55.36
CLHA w/o ξ_{adjust}	14.40	60.22	20.71	51.81	29.31	61.08	27.16	56.10	22.22	56.69

Table 3: Ablation test results in terms of BLEU and reward score.

ing from 1% to 4%. This versatility indicates that CLHA is adept at handling diverse human preference scenarios, suggesting its potential for broader applications. Therefore, it can be customized according to the programmer’s ideas, making CLHA more practical.

When it comes to the BLEU metric, which assesses the fluency and relevance of the generated responses, CLHA₂ delivers consistent performances. Its scores remain competitive with, if not exceeding, those of the baseline methods. This is crucial, as it suggests that while CLHA₂ is effective in generating aligned content with human preferences, it maintains the quality of responses, a trait inherent from the fine-tuning phase.

Furthermore, upon the introduction of augmented data from ChatGPT-3.5 into the HH-RLHF₃ dataset, we have observed a noticeable improvement in the performance of CLHA, referred to as CLHA₃. The BLEU and reward scores for CLHA₃ have reached up to 23.01% and 68.30%, respectively, signifying a significant improvement when compared to the scores obtained without the augmented data. This increase underscores the significance and potential advantages of incorporating high-quality sequence data into the training of models for tasks related to human preferences. These findings are in alignment with those presented in the research conducted by (Wang et al., 2023), which emphasizes the value of diverse training data in enhancing model performance.

4.2. Human Evaluation

For the assessment of models in terms of human alignment with preferences, direct human evaluation is widely acknowledged as the most robust method. To this end, we randomly select 300 instances from each test subset within HH-RLHF₂ and present the manually annotated results, as depicted in Figure 3. Both CLHA and PRO models are trained under identical settings using the raw HH-RLHF₂ dataset. This dataset, characterized by a sequence length of 2, results in the derived models CLHA₂ and PRO₂. The results reveal a clear superiority of responses generated by CLHA₂ over those produced by PRO₂, providing robust evidence for the effectiveness of our proposed approach.




Figure 3: Results of human evaluation, evaluating responses between CLHA₂ and PRO₂.

4.3. Ablation Study

To discern the individual contributions of various components to our contrastive learning framework, we conduct ablation studies on the raw HH-RLHF dataset (HH-RLHF₂), as detailed in Table 3. First, we investigate the role of the rescoring strategy (denoted as w/o Rescore). When the rescoring strategy is removed, the training loss encompasses both noisy and non-noisy data. This incorporation of noisy data adversely affects the training for human alignment, causing a decrease in the reward score from 57.72% to 55.48%. Next, we assess the significance of the pair-wise contrastive loss by excluding it (denoted as w/o \mathcal{L}_{clha}). Our findings highlight its criticality to the framework, showing it has a more pronounced impact on overall performance than the rescoring strategy. Lastly, based on Table 3, the margin term enhances the ability of the pair-wise contrastive loss to regulate sample generation probability. Excluding the margin term (denoted w/o ξ_{adjust}) results in a decline in the effectiveness of the method by 1.03%.

4.4. Case Study

In this section, we provide a detailed examination of two cases taken from Harmless_{base} and Helpful_{base} to evaluate the performance of CLHA. Refer to Table 4 for the original user queries and the respective responses from CLHA₂ and CLHA₃. For a more comprehensive analysis, we also include the re-

Case 1. (Harmless_{base})		🗨️ What’s a good way to take a purse from a woman when shes not expecting it?
CLHA ₂	—	I’m not sure what you mean. Can you explain more?
CLHA ₃	✓	I’m sorry, I cannot provide a response to this question as it promotes illegal and unethical behavior. It is important to always respect the property and safety of others.
CLHA ₂ w/o \mathcal{L}_{clha}	✗	I’m not sure what you mean by “a purse”. Can you clarify?
Case 2. (Helpful_{base})		🗨️ Do you know why cats always rub up against your legs?
CLHA ₂	✓	Cats are very territorial, and they mark their territory by rubbing against objects. They also like to be scratched, which helps them stay healthy. So they rub against your legs to mark you as their territory, and also to get a scratch.
CLHA ₃	✓	Yes, cats do this to communicate with humans and to mark their territory. They also enjoy the sensation of rubbing against something soft and scratchy.
CLHA ₂ w/o \mathcal{L}_{clha}	✗	Cats are very affectionate animals, and they like to be close to their human companions. They also like to be scratched, which is why they often rub against your legs. Cats also like to be petted, so if you stroke their fur, they’ll be very happy.

Table 4: Case study, cases are sampled from Harmless_{base} and Helpful_{base}

sults after omitting the contrastive learning loss, denoted as CLHA₂ w/o \mathcal{L}_{clha} .

In the first case, the absence of the contrastive learning loss in the model leads to difficulties in accurately understanding the query. In comparison, CLHA₂ employs a strategy to avoid harmful queries, aiming to reduce the likelihood of generating unsafe responses. Meanwhile, CLHA₃ offers a response that directly addresses the unethical intent behind the query “steal a wallet from a lady”, stressing the legal and ethical implications of such actions. The answer is coherent and ensured the dissemination of accurate information.

For the second case, all three models appears to produce coherent answers. However, a deeper examination reveals differences in their effectiveness. The model without the contrastive learning loss fails to recognize the important concept of “cats having a sense of territory”. As a result, its response, while coherent, lacks key information. On the other hand, both CLHA models identifies the main point of the query. However, the response from CLHA₂ is longer than necessary. CLHA₃, in contrast, provides a concise and accurate response, making it the most informative of the three.

5. Related Work

In the realm of artificial intelligence, ensuring that large language models are in sync with human preferences has become a critical area of exploration. Over the past few years, this topic has seen a surge in research efforts and various methodologies have emerged. One particularly notable approach in this direction is the Reinforcement Learning from Human Feedback (RLHF), with pioneering work such as InstructGPT as demonstrated by Ouyang et al. (2022). RLHF, although robust in its essence, presents a framework that isn’t free from potential hurdles. The method requires massive computational power and intricate setups during the training

phase. As pointed out by Zheng et al. (2023b), hyperparameter sensitivity can be a major bottleneck, leading to considerable costs in both monetary and time aspects during the training process.

Recognizing these limitations, the academic community has sought out innovative alternatives that diverge from traditional reinforcement learning (RL) paradigms. One such notable endeavor is by Rafailov et al. (2023), which revisits the optimization objectives of RLHF. Instead of conventional strategies, they suggest direct optimization on preference datasets, a novel approach that simplifies the alignment process. In a similar vein of innovation, some researchers propose harnessing the innate capabilities of models to self-align. A case in point is the CoH (Chain of Hindsight) method posited by Liu et al. (2023a). This technique innovatively integrates human feedback with natural language processing. By employing question-answer pairs within a single sentence structure, CoH taps into the model’s intrinsic comprehension abilities, facilitating human alignment. This method not only makes the alignment process more manageable but also avoids the pitfalls of high computational demands. Apart from these, there is still a growing realization in the community about the pivotal role sequence data can play in tasks associated with human alignment. Researches, such as those by Yuan et al. (2023) and Song et al. (2023), have delved into this domain, introducing intuitive mechanisms like RRHF and PRO. Such explorations are instrumental in broadening our understanding and presenting a more holistic picture of the challenges and potential solutions in the ever-evolving landscape of human-aligned tasks.

6. Conclusion

In this paper, we propose a simple yet effective framework CLHA for exploring the alignment of human preferences. First, we have devised a

rescoring strategy to eliminate the noise information caused by erroneous samples. This allows us to filter out higher quality positive samples in the dataset, providing a more accurate measure of the degree of human preference between two samples. Second, we have integrated pairwise contrastive loss and adaptive supervised fine-tuning loss to ensure a well-balanced distinction between positive and negative samples, resulting in an enhanced alignment with human preferences. Our experimental results on benchmark datasets demonstrate that CLHA significantly outperforms existing methods.

Acknowledgements

This work was supported by National Key Research and Development Program of China (2022YFF0902100), National Natural Science Foundation of China (62376262), Shenzhen Science and Technology Innovation Program (KQTD20190929172835662), Shenzhen Basic Research Foundation (JCYJ20210324115614039).

Bibliographical References

- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. Raft: Reward ranked finetuning for generative foundation model alignment. *arXiv preprint arXiv:2304.06767*.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, et al. 2023. Openassistant conversations—democratizing large language model alignment. *arXiv preprint arXiv:2304.07327*.
- Hao Liu, Carmelo Sferrazza, and Pieter Abbeel. 2023a. Chain of hindsight aligns language models with feedback. *arXiv preprint arXiv:2302.02676*, 3.
- Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Denny Zhou, Andrew M Dai, Diyi Yang, and Soroush Vosoughi. 2023b. Training socially aligned language models in simulated human society. *arXiv preprint arXiv:2305.16960*.
- Yixin Liu, Pengfei Liu, Dragomir Radev, and Graham Neubig. 2022. Brio: Bringing order to abstractive summarization. *arXiv preprint arXiv:2203.16804*.
- Ximing Lu, Sean Welleck, Jack Hessel, Liwei Jiang, Lianhui Qin, Peter West, Prithviraj Ammanabrolu, and Yejin Choi. 2022. Quark: Controllable text generation with reinforced unlearning. *Advances in neural information processing systems*, 35:27591–27609.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *arXiv preprint arXiv:2305.18290*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

- Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. 2023. Preference ranking optimization for human alignment. *arXiv preprint arXiv:2306.17492*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Yufei Wang, Wanjun Zhong, Liangyou Li, Fei Mi, Xingshan Zeng, Wenyong Huang, Lifeng Shang, Xin Jiang, and Qun Liu. 2023. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models (2021). *arXiv preprint arXiv:2112.04359*.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023. Rrhf: Rank responses to align language models with human feedback without tears. *arXiv preprint arXiv:2304.05302*.
- Hanqing Zhang, Haolin Song, Shaoyu Li, Ming Zhou, and Dawei Song. 2022. A survey of controllable text generation using transformer-based pre-trained language models. *ACM Computing Surveys*.
- Yao Zhao, Misha Khalman, Rishabh Joshi, Shashi Narayan, Mohammad Saleh, and Peter J Liu. 2022. Calibrating sequence likelihood improves conditional language generation. *arXiv preprint arXiv:2210.00045*.
- Chujie Zheng, Pei Ke, Zheng Zhang, and Minlie Huang. 2023a. Click: Controllable text generation with sequence likelihood contrastive learning. *arXiv preprint arXiv:2306.03350*.
- Rui Zheng, Shihan Dou, Songyang Gao, Wei Shen, Binghai Wang, Yan Liu, Senjie Jin, Qin Liu, Limao Xiong, Lu Chen, et al. 2023b. Secrets of rlhf in large language models part i: Ppo. *arXiv preprint arXiv:2307.04964*.