

# VickreyFeedback: Cost-efficient Data Construction for Reinforcement Learning from Human Feedback

Guoxi Zhang<sup>1\*</sup> and Jiuding Duan<sup>2\*</sup>

<sup>1</sup> State Key Laboratory of General Artificial Intelligence, BIGAI, Beijing, China

<sup>2</sup> Allianz Global Investors Japan Co., Ltd., Tokyo 106-0032, Japan  
{altriaex86,jiuding.duan}@gmail.com

**Abstract.** This paper addresses the cost-efficiency aspect of Reinforcement Learning from Human Feedback (RLHF). RLHF leverages datasets of human preferences over outputs of large language models (LLM)s to instill human expectations into LLMs. While preference annotation comes with a monetized cost, the economic utility of a preference dataset has not been considered by far. What exacerbates this situation is that, given complex intransitive or cyclic relationships in preference datasets, existing algorithms for fine-tuning LLMs are still far from capturing comprehensive preferences. This raises severe cost-efficiency concerns in production environments, where preference data accumulate over time. In this paper, we see the fine-tuning of LLMs as a monetized economy and introduce an auction mechanism to improve the efficiency of the preference data collection in dollar terms. We show that introducing an auction mechanism can play an essential role in enhancing the cost-efficiency of RLHF, while maintaining satisfactory model performance. Experimental results demonstrate that our proposed auction-based protocol is cost-efficient for fine-tuning LLMs by concentrating on high-quality feedback.

## 1 Introduction

Large language models (LLMs) have revolutionized natural language processing and artificial intelligence, enabling unprecedented capabilities in generating, summarizing, and understanding text at a high cognitive level [29,36,19,10]. The delivery of these models in domain-specific applications relies on fine-tuning, i.e., aligning the LLMs closely with human preferences. A key to successful fine-tuning is the construction of a high-quality dataset, which is an emerging but challenging problem [27,23,35]. Conventional fine-tuning methods rely heavily on human-annotated preference datasets to guide the model toward generating qualified content for domain usage. The success of these methods is backed by available datasets and has reinforced the urgent demand for large-scale high-quality datasets [24,11]. Despite the proven effectiveness in the conventional

---

\* Both authors contributed equally.

Reinforcement Learning from Human Feedback (RLHF) pipelines, these methods often overlook the economic aspects when creating responses and gathering preference feedback [15]. This can be a significant oversight in real-world applications, because model owners who have only a limited budget will compromise the model’s performance due to a budgeted situation in data quantity and quality.

Specifically, there are two issues associated with an economic consideration. First, gathering high-quality preference annotation is inherently expensive and economically inefficient due to a lack of quality assurance in the process of data collection [4,38,11]. Existing data quality assurance protocols either leverage human annotations or LLM models to simulate the preference labeling [1,16,32,40]. The second issue is the computational redundancy in the phase of reward modeling, where existing algorithms, such as those based on Bradley-Terry (BT) models for reward prediction may fail to capture the incorrect and ambiguous preference pairs due to these preferences’ complex, often intransitive or cyclic nature [5,31,25]. Moreover, it is observed that the misalignment between model design and dataset quality is structural and can lead to large inference errors as well as redundant use of low-quality data and computing resources, diminishing the marginal economic value of the training process [8,10]. These economic issues are critical in a production environment, because service owners confront evolving user demands and expect an organic lifelong evolution of the dataset and the models [26].

In this paper, we outline the practical limitations of existing RLHF methodologies in the context of economic efficiency and present a novel auction-based mechanism designed to optimize the cost-effectiveness of fine-tuning processes. The proposed method addresses three missing aspects in recent efforts to construct human preference datasets for RLHF:

1. Regarding the quality assurance aspect of LLM agents’ responses, we argue that by introducing an auction mechanism as a communication tool, LLM agents can be incentivized to provide pricing information truthfully, therefore, avoidance of myopic pricing/bidding of LLM agents, and a discovery of a fair price of LLM responses can be made possible. We propose an auction mechanism for procuring datasets on the fly, where judicious exploitation of resources by the dataset owner, paving the way for *fair valuation* of contents provided by agents and a sustainable construction of high-quality dataset.
2. Regarding the cost efficiency concerns by dataset owners, we argue that the dataset owner can leverage the proposed mechanism to achieve a balanced cost efficiency during data accumulation. This can be resolved as a cost minimization problem in a procurement setting by incorporating the dataset owner’s budget into quantity allocation and payment decisions [3].
3. Given the mechanism, we further devise **Vickrey-QA**, an algorithm that can leverage the bidding information provided by LLM agents in the fine-tuning process to enhance the performance based on small but selective data, while maintaining the overall quality of fine-tuned LLM. Experimental results verify our arguments in the cost-efficiency of dataset construction and downstream fine-tuning of the base model.

## 2 Related Work

### 2.1 Datasets for RLHF

Tremendous efforts have been made in dataset construction to facilitate RLHF, enabling diverse applications and contexts for RLHF. Such datasets combine human feedback, training data, interaction logs, reward signals, and evaluation metrics to train and fine-tune language models according to human preferences [27,23,11,13].

Alignment of different perspectives, e.g., helpfulness and harmlessness [1] can improve performance on almost all NLP evaluations, and is fully compatible with training for specialized skills such as coding and text summarization [12]. As helpfulness and harmlessness often stand in opposition to each other, preference models trained to evaluate one of these qualities primarily have unfavorable performance (much worse than chance) on the other. Assembled with data that contains multiple aspects, the preference model can nevertheless learn the right lessons and behave helpfully when appropriate, while encouraging the polite refusal of harmful requests.

Moreover, many other qualitative aspects of human preferences, e.g., instruction following, honesty, and truthfulness have been explored [4]. Given the diverging nature of multiple aspects in preference annotation, primitive quality assurance can be adopted [11]. For example, instruction following ratings ensure the comprehension of the intention of human instructions without deviating from the requirements. Honesty reflects what they (don't) know and express uncertainty when they are in waver towards the given problem. Truthfulness rating reveals the alignment between the instructions and real-world knowledge, not fabricating any facts without references or introducing any self-contradiction.

### 2.2 Mechanism Design in RLHF

Mechanism design is an area in economics and game theory that focuses on creating systems or protocols (mechanisms) that lead to desirable outcomes for individual agents and social good. In the context of RLHF, mechanism design can have presence in 2 prospective tasks. The first task is aggregating inputs from multiple LLM agents. This task requires designing auction mechanisms that can incentivize the agents to reveal their preferences truthfully. Token Auction Model is representative and operates on a token-by-token basis, allowing LLM agents to influence content generation through single-dimensional or informed multi-dimensional bids [6,28]. The second prospective task is cost-efficient data construction, which is a prerequisite for achieving desirable agent performance. Collecting instruction responses and human preference requires cost in dollar term and has been a bottleneck in budgeting and model operation in production environment. This bottleneck is alternatively termed as alignment tax, due to its connection with degenerated model performances [7,17]. Other related efforts include adopting game-theoretic perspectives for fine-tuning, e.g., Nash Direct Optimization[25], nonetheless limited attention is paid on data construction.

Mechanism design research in economic theory provides succinct tools for optimal mechanism design. Among others, Vickrey auction or sealed-bid second-price auction (SBSPA) is a type of well-studied as a primitive but useful mechanism to incentivize truthful bidding of each bidders or agents [21].

In this paper, we argue that designing effective mechanisms for dataset construction in RLHF can be seen as a procurement problem that has been extensively studied in economic theory and operations research [3]. The common objective is to ensure the service/content provider has an incentive for truthful bidding and the buyer of data has control over his payments at the expense of introducing uncertainty in the quantity acquired in the process.

### 3 Preliminaries

#### 3.1 Vanilla Preferences

*Notations* In RLHF, human preferences are encoded as pairwise comparisons between model responses. Specifically, a preference sample is a triplet  $(x, y_a, y_r)$ , where  $x$  is the instruction that describes the desired response,  $y_a$  is the accepted (preferred) response, and  $y_r$  is the rejected (not preferred) response. For example,  $x$  can be “Write a code snippet to compute the Fibonacci sequence.”,  $y_a$  can be a program that runs successfully, and  $y_r$  can be a program with syntactic errors.

To construct a preference dataset  $D$ , we need to sample instructions, generate responses, and annotate responses with preferences. For a specific application, the instructions can be sampled from history logs and thus often free of charge. For each instruction, we need at least two responses, which are typically generated by commercial LLMs. Finally, preference annotation means to determine the accepted and rejected responses, and it also done by invoking commercial LLMs. Take the UltraFeedback [4] dataset as an example. Its instructions are collected from public NLP datasets such as FLAN [18] and Evol-Instruct[34]. For each instruction, four responses are generated by either commercial or open-source LLMs and rated by GPT-4. The accepted responses are the ones with the highest rating, while the rejected response is sampled from the remaining ones [2]. This paper refers to preference samples generated with the above-mentioned pipeline as the *vanilla* preferences.

*Cost of Vanilla Preferences* When instructions are sampled from historical data, the cost of a vanilla preference dataset comprises of the cost of model responses and preference annotation. In this paper, we focus on the cost of model responses. Currently, commercial LLMs charge for the amount of input and output texts, which are measured in tokens. So the cost of a model response is proportionate to the length of its instruction and itself.

#### 3.2 Learning from Preference Datasets

Technically, an LLM is a generative model for texts. Denote by  $\pi_\theta(y|x)$  the likelihood of response  $y$  conditioned on instruction  $x$ , where  $\theta$  refers to the parameters

to be optimized. The aim of RLHF is to instill the knowledge of human preferences into LLMs. That is, the model should learn to assign higher likelihood to the accepted responses of preference samples.

The direct preference optimization (DPO) [24] algorithm is based on the fact that, before learning from preferences, the parameters of a LLM  $\theta$  have been pre-trained on a large collection of datasets. Therefore, the likelihood of a response under pre-trained model can be consider as a legitimate baseline for its goodness. Denote by  $\pi_{\text{ref}}$  the likelihood function parameterized by a fixed copy of  $\theta$  constructed before learning. The objective function to be minimized is given by

$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(x, y_a, y_r) \sim D} \left[ \log \sigma \left( \beta \log \frac{\pi_{\theta}(y_a|x)}{\pi_{\text{ref}}(y_r|x)} - \beta \log \frac{\pi_{\theta}(y_r|x)}{\pi_{\text{ref}}(y_r|x)} \right) \right], \quad (1)$$

where  $\sigma(\cdot)$  is the sigmoid function and  $\beta$  is a hyper-parameter. Intuitively, this objective enlarges the difference between the “improved goodness” measured by  $\log \frac{\pi_{\theta}(y|x)}{\pi_{\text{ref}}(y|x)}$  of accepted and rejected responses, which means the model is encouraged to generate accepted responses.

Fine-tuning entire language models is prohibitively expansive for specific applications. Low-rank adaptation [9] (LoRA) is a recent parameter efficiency fine-tuning approach that effectively reduces the number of parameters to be optimized. Given a pre-trained LLM, with LoRA we only learn “adaptors” of its large parameter matrices and keep the parameter matrices themselves unchanged. After being fine-tuned on a dataset, these adaptors can encode the knowledge in the dataset. Thus, we can gauge the improvement resulting from fine-tuning by comparing the performance of the original LLM (called the base model) and the combination of the LoRA adaptor and the base model.

### 3.3 Vickrey Auction

In Vickrey auction, bidders submit written bids without knowing the other people’s bids in the auction. The highest bidder wins but the price paid is the second-highest bid. Generalized variants of the Vickrey auction for multiunit auctions exist, such as the generalized second-price auction used in online advertisement programs and the Vickrey–Clarke–Groves (VCG) auction [30]. In VCG auction, truthful bidding of each bidder is theoretically justifiable and incentivized through the mechanism that the buyer will adopt the highest bid but only pay for the second-highest bid price. The property of VCG is well studied in economic theory and stretched under reasonable assumptions [21].

## 4 VickreyFeedback

In this paper, we propose a data collection pipeline based on Vickrey auction to achieve quality assurance under a limited budget. Inspired by the multi-unit VCG auction mechanism in procurement [3], we consider the sampling and

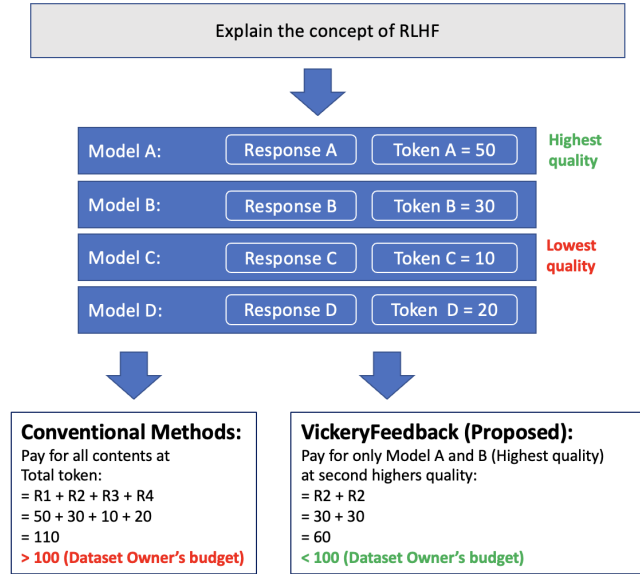


Fig. 1: Illustration of VickreyFeedback for RLHF. Instruction is given to four LLM agents simultaneously and each of the agents responds with a  $(x, y_a, y_r)$ , where  $x$  represents the model’s response,  $y_a$  is acceptance of the response by the mechanism,  $y_r$  is the model’s true bid price/valuation of the respective  $x$ . Assuming that the quality of model responses is proportional to the length of the respective response, accepting a longer response is a proxy for selecting a higher-quality response when no additional quality control protocol is available.

dataset construction from subsidiary LLM agents as a procurement problem. This pipeline is illustrated in Fig. 1.

#### 4.1 Assumptions

Recent discussions on data diversity, quality, and quantity argue that scaling up diversity and output quality has measurable positive effects for achieving better alignment while scaling up quantity alone might not [39]. We follow the intuition and assume that the quality of model responses is proportional to the length of the respective response. Therefore, accepting a longer response is a proxy for selecting a higher-quality response when no additional quality control protocol is available [14].

#### 4.2 Protocol

The data collection protocol of VickreyFeedback is illustrated in Algorithm 1. The protocol requires multiple LLM agents as response suppliers given an in-

---

**Algorithm 1 VickreyFeedback:** A mechanism for procuring LLM responses

---

**Require:** An instruction  $I$ , a set of LLMs  $\{LLM_1, LLM_2, LLM_3, LLM_4\}$ **Ensure:** Selected response and payment to the winning LLMEach  $LLM_i$  submits a response  $r_i$  and quality  $q_i$ , where  $q_i = \text{length}(r_i)$ Evaluate each response  $r_i$  by agreeing at the valuation  $v_i$ , where  $v_i = q_i$ Identify the winning LLM  $i^*$  with the highest data quality:  $i^* = \arg \max_i q_i$ Identify the second-highest quality:  $j = \arg \max_{i \neq i^*} q_i$ Pay equally to both agent  $i^*$  and  $j$  at  $q_j$  in dollar termCalculate total payment  $C_{total}$  in the collection of responses for instruction  $I$  by

$$C_{total} = q_j + q_j = 2 \times q_j$$

**return** A selective response pair  $(r_{i^*}, r_j)$  for instruction  $I$ 

---

struction  $I$ . Each supplier is required to submit a response  $r_i$  and a self-evaluated quality  $q_i$ . In our experiment,  $q_i = \text{length}(r_i)$  holds by assumption, where  $\text{length}(r)$  denotes the token length of a response  $r$ . The mechanism selects the two agents that provided the longest responses and pays only the second longest responses to each winning agent. Consequentially, the responses that have a lower quality or shorter token length will not receive rewards and the information will not be fed to the fine-tuning of LLMs.

Considering the cost budget of the dataset owner in the dataset construction for RLHF, truth-telling from all the suppliers is important because it is the prerequisite for allocating limited monetary resources from the dataset owner to the data suppliers. The rationale for data suppliers doing so relies on whether bidding truthfully is a dominant strategy for each of the suppliers [22].

It is trivial to prove that under mild conditions, the total cost  $C_{total}$  by the proposed VickreyFeedback protocol is smaller than the total construction cost of the conventional methods in RLHF. Besides, our proposed protocol is organic to RLHF by enabling a control over the dataset construction cost, while not degenerating the performance. This is the first effort that incorporates a mechanism into dataset construction in RLHF, supported by the truthful bidding property of the auction mechanism. Moreover, by proposing an associated algorithm, we compensate the information loss that might happen when dropping diversified answers in the data collection phase.

### 4.3 Quality-Adjusting DPO

A downside of VickreyFeedback is that it sacrifices data diversity for data quality to attain better cost efficiency. Recall that for each instruction, we do not collect the responses whose declared quality is below the second-to-best quality. The responses included in a Vickrey preference dataset are thus mostly *good* responses. This data skewness could be problematic for RLHF as models cannot learn a wide spectrum of human preferences but overfit to a narrow part of them. Put differently, a model might capture the nuances of high-quality responses but is not trained to prevent generating low-quality responses.

We tackle this drawback by enhancing models’ ability to differentiate responses. The intuition is that, if a model is capable of differentiating good responses from bad ones, then it should have encoded typical patterns of bad ones and will not generate them. Our key observation is that the quality differences between the best responses and the second-to-best responses can be different for different instructions. The samples whose responses have large quality differences are more valuable than the rest as they are more helpful for differentiating responses. Therefore, we propose a direct extension of the DPO algorithm that weights samples according to their quality difference, which is called quality-adjusting DPO (QA-DPO). In specific, the loss function of QA-DPO is given by

$$\mathcal{L}_{\text{QA-DPO}}(\theta) = -\mathbb{E}_{(x, y_a, y_r) \sim D} \left[ w(b_a, b_r) \log \sigma \left( \beta \log \frac{\pi_\theta(y_a|x)}{\pi_{\text{ref}}(y_r|x)} - \beta \log \frac{\pi_\theta(y_r|x)}{\pi_{\text{ref}}(y_r|x)} \right) \right], \quad (2)$$

where  $w(b_a, b_r)$  is a function of the declared quality of the two responses. The weighting function  $w(b_a, b_r)$  should be larger for samples with diverse response. We use  $w(b_a, b_r) = 0.5 + \sigma(b_a - b_r)$ , which is monotonically increasing with  $b_a - b_r$  and satisfies  $w(b_a, b_r) = 1$  if  $b_a = b_r$  and  $w(b_a, b_r) \rightarrow 1.5$  if  $b_a - b_r \rightarrow +\infty$ . This weighting mechanism emphasizes the influence of samples containing diverse responses and thus mitigates the downside of Vickrey preferences.

## 5 Experiments

In this section, we aim to demonstrate that the cost for preference collection can be reduced with VickreyFeedback if the assumptions in Section 4 are satisfied. Further, we shall show that our proposed QA-DPO can address the downside of Vickrey preferences.

### 5.1 Setup

*Datasets* We use the binarized version [2] of the UltraFeedback [4] as vanilla preference samples. The original UltraFeedback dataset contains 63,967 instructions and four responses for every instruction. The responses are annotated by GPT-4 with four scores (from one to five) to evaluate them from four aspects: instruction-following, truthfulness, honesty, and helpfulness. In the binarized version [2], the four scores of response are averaged into an overall score. The accepted response of an instruction is the one with the highest overall score, and the rejected response is sampled uniformly at random from the remaining three responses. This dataset is referred to as vanilla preferences in our results.

We simulate the proposed VickreyFeedback pipeline and generate synthetic Vickrey preferences. Based on the assumption that response providers report response quality truthfully, we use the overall scores as declared quality. For each instruction, the accepted response is still the one with the highest overall score, but the rejected response is the one with the second-highest overall score.

To investigate the influence of data size, we experiment with 25%, 50%, and 100% of samples.

*Models* To showcase the effect of RLHF, we use the `Dolphin-7B` model<sup>3</sup> as the base model, since it is trained on datasets without any alignment or bias. This model is called the Base in our results. We compare three fine-tuned models in our experiments. The model fine-tuned on vanilla preferences is referred to as Vanilla-DPO. The model fine-tuned on Vickrey preferences using the DPO algorithm is called Vickrey-DPO, and the model fine-tuned on Vickrey preferences with our QA-DPO algorithm is denoted as **Vickrey-QA**.

*Metrics* We use the pairwise model evaluation protocol of the MT-bench [37] benchmark to evaluate model performance. The first step is to generate model responses for the instructions in this benchmark using the models of interest. Then pairs of responses, accompanied with the instructions and reference answers, are presented to another more powerful LLM to determine the better one in each pair. This LLM-based evaluation is of low cost and highly correlated with human evaluation [37]. In our experiments, we use the `GPT-4o` to compare model responses. In specific, we report the win rate of a model against another model as the metric for performance. For two models A and B, the win rate of A against B is the number of times that A’s responses are considered to be better than B’s responses divided by the total number of pairwise comparisons between their responses. A tie counts for half a win. This metric reflects the extent A is better than B. The higher, the better. A win rate of 0.5 indicates that the two models have similar performance.

As for cost, we report the number of tokens of the responses included in vanilla preferences and Vickrey preferences. The instructions are excluded when computing cost because they are the same for both types of preferences. For simplicity, we use the tiktoken library<sup>4</sup> when computing tokens.

*Technical Details* We use the implementation of LoRA from the PEFT package [20] and report the values of hyperparameters in Table 1. All models are fine-tuned for two epochs using eight A800 GPUs. It takes about six hours to fine-tuned models on 100% of the preference samples.


## 5.2 Results

We start with presenting results for model performance and the trade-off between data quality and diversity. Then, we compare models from the perspective of cost and discuss the pros and cons of VickreyFeedback.

Figure 2a shows the win rate against Base when using different subsampling ratios. When using 25% of the samples, both Vickrey-DPO and Vickrey-QA outperform Vanilla-DPO. As we use more samples, Vanilla-DPO begins to outperform Vickrey-DPO, but Vickrey-QA is still comparable with Vanilla. Figure 2b provides an in-depth analysis of the performance of Base, Vanilla-DPO,

<sup>3</sup> <https://huggingface.co/cognitivecomputations/dolphin-2.9.3-mistral-7B-32k>

<sup>4</sup> <https://github.com/openai/tiktoken>



(a) **Vickrey preferences lead to better performance on small datasets.** The win rate is with respect to Base. Both Vickrey-DPO and Vickrey-QA outperform Vanilla-DPO when the subsampling ratio is 25%, and Vickrey-QA remains competitive against Vanilla for larger subsampling ratios.

(b) **Vickrey-QA matches Vanilla-DPO on large datasets.**

The win rate of a model (row) against another model (column) when using 100% of the preferences. Vanilla-DPO, Vickrey-DPO, and Vickrey-QA all show improvement against Base, and Vickrey-QA matches the performance of Vanilla-DPO and Vickrey-DPO.

Fig. 2: Results for model performance.

Vickrey-DPO, and Vickrey-QA on 100% of the samples. It shows that Vanilla-DPO, Vickrey-DPO, and Vickrey-QA all outperform Base. Moreover, Vickrey-QA matches the performance of Vickrey-DPO and Vanilla-DPO, as its win rate against the two models is 0.49 and 0.5. These results confirm that Vickrey-QA works well for both small and large datasets.

Meanwhile, Figures 3a, 3b and 4 offer evidence for the difference in data diversity between vanilla and Vickrey preferences. Figures 3a and 3b visualize the distribution of responses aggregated by their source LLMs when using 100% of the samples. It turns out that the response distribution in vanilla preferences is flatter than that in Vickrey preferences. Since different source LLMs are trained on different datasets and thus have their own biases, this observation implies that the responses in the vanilla preferences tend to be more diversified. Meanwhile, Fig. 4 visualizes the distribution of responses aggregated by overall scores, and it shows that there are more responses with overall scores larger than 3.75 in Vickrey preferences. These observations confirm that Vickrey preferences are less diversified than vanilla preferences.

Therefore, we argue that our VickreyFeedback indeed exchanges data diversity for quality. When using the DPO algorithm, VickreyFeedback leads to better performance on small datasets, but this advantage is outweighed by the loss of diversity on large datasets. Our proposed QA-DPO can mitigate this drawback and yield competitive performance for various sizes of datasets.

Table 1: Hyper-parameters used in our experiments. Our implementation is based on the TRL [33] and the PEFT [20] package. For hyper-parameters not listed here, we use the default values in the corresponding package.

DPO Hyper-parameters	
#epochs	2
batch size per GPU	7
#GPU	8
gradient checkpointing	True
optimizer	adamw_torch_fused
learning rate	$5 \times 10^{-5}$
max grad norm	0.3
warmup ratio	0.1
learning rate scheduler	cosine
bf16	True
tf32	True
$\beta$	0.1
max prompt length	4096
max length	4096
LoRA Hyper-parameters	
alpha	128
dropout	0.05
target module	all-liner
task type	CAUSAL_LLM

The proposed VickreyFeedback leads to better cost efficiency. Fig. 5a visualizes the costs of datasets when we increase the number of samples. It turns out that the average cost per sample is larger in the case of Vickrey preferences. This can be explained by the design of VickreyFeedback: we only include high-quality responses that are often longer than low-quality responses and thus incurs more costs. Meanwhile, Fig. 5b visualize models’ win rate as a function of data collection cost, which directly reveals the cost efficiency of different models.

To conclude, when the dataset owner has a limited budget (i.e.,  $1.5 \times 10^7$  in our example) for collecting data, Vickrey-DPO has the best performance, which is followed by Vickrey-QA and Vanilla. As the cost budget increases, Vanilla-DPO has a higher chance to achieve the best performance, but its advantage over Vickrey-QA diminishes as the cost increases. Therefore, when combined with our QA-DPO, VickreyFeedback is more cost-efficient than the competing protocol for collecting preference data.

## 6 Discussion

### 6.1 Diversity of Preference

As illustrated in Figures 3a, 3b and 4, preferences collected by VickreyFeedback are less diversified than those collected with the conventional approach.

The limited diversity of the sampled dataset is due to small data constraints, restricted by the dataset owner’s budget for collecting data. Certain ablation studies on data diversity, quality, and quantity have shown that data quantity is




Fig. 3: **Responses are sampled more evenly from different LLMs in vanilla preferences.** This figure shows the distributions of responses aggregated by their source LLMs in the vanilla and Vickrey preferences.





Fig. 4: **Vickrey preferences contain more high-quality responses.** This figure shows the distributions of responses aggregated by their overall scores in the vanilla and Vickrey preferences.

not a critical issue while diversity might affect the final performance [39]. Specifically, the proposed auction mechanism primarily focuses on the highest and second-highest bids, neglecting valuable information from other bidders. Bidders with significantly different responses from the top two bids might have their preferences underrepresented in the dataset. Despite DPO and QA-DPO can provide a computational trade-off between maximizing reward and minimizing KL-divergence, we argue that the loss of information diversity in VickreyFeed-



(a) **Vickrey preferences are more expansive on average.** This figure illustrates the costs of preference samples as a function of the number of samples.

(b) **Vickrey preferences are more cost-efficient.** This figure shows how the win rate against the base model changes with the cost of data collection.

Fig. 5: Analysis for model performance and data collection cost.

back is structural. This is quantitatively evaluated in Figures 3a, 3b and 4, where QA-DPO tends to sample from well-known LLM agents, e.g., GPT, llama, etc., that are good at generating responses with longer tokens. When comparing between Figures 3a and 3b, newer model versions that can generate higher-quality responses receive more attention from the mechanism, e.g., GPT-4 has a higher sampling ratio than GPT-3.5-turbo in Vickrey preferences while the older version GPT-3.5-turbo has a higher sampling ratio than GPT-4 in vanilla preference setting. As a result, the ratio of sampling higher response ratings is also higher in the case of Vickrey preference.

## 6.2 Cost-efficient Data Construction

To the best of our knowledge, this is the first research speculating on the economic aspects of dataset construction in RLHF.

In the context of Direct Preference Optimization (DPO) and dataset collection for fine-tuning Large Language Models (LLMs), there exists a critical trade-off among cost, preference diversity and fine-tuning performance [17]. DPO aims to refine models based on direct human preferences, necessitating a broad and diverse set of preference data to capture the nuanced and varied requirements of different users. However, obtaining such a diverse dataset is often expensive, as each instance involves a monetary cost. Therefore, striking the right trade-off is crucial and an auction mechanism, as proposed in this paper, can help navigate the resource allocation towards a maximized utility budgets for model fine-tuning.

Overall, the proposed **VickreyFeedback** offers certain advantages, such as cost-efficiency in constructing subsidiary LLM agents’ responses under a limited budget. Users of the proposed protocol should be aware that the proposed **QA-DPO** has a lower topline in data diversity, compared with the vanilla DPO algorithm where KL divergence is measured on a more granular dataset backed by all available LLM agents and resources.

## 7 Conclusion

This paper investigates a cost-efficient auction mechanism designed for aligning large language models (LLMs) with human preferences, a process integral to reinforcement learning from human feedback (RLHF). Traditional approaches rely heavily on procuring high-quality annotated preference datasets, neglecting the construction cost and economic utility of datasets. Considering the monetized cost of preference annotation and the complexities of intransitive or cyclic preference relations, which may fail to properly express preferences or degrade the marginal economic value during fine-tuning. This is particularly challenging when datasets need to be constructed over time in production environments.

To address these issues, we introduce an auction-based protocol utilizing a type of second-price auctions, i.e., Vickrey–Clarke–Groves (VCG) auction, to enhance cost efficiency in dataset construction. By treating the data collection process for fine-tuning LLMs as a monetized economy, the proposed mechanism incentivizes the data providers to provide high-quality responses. This helps the model owners who have a limited budget, in measuring the cost-efficiency and budgeting for data collection. Experimental results demonstrate that our proposed auction-based approach is not only cost-efficient for fine-tuning LLMs but also practical for online construction of datasets, while maintaining satisfactory model performance.

## References

1. Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al.: Training a helpful and harmless assistant with reinforcement learning from human feedback. arXiv preprint arXiv:2204.05862 (2022)
2. Bartolome, A., Martin, G., Vila, D.: Notus. <https://github.com/argilla-io/notus> (2023)
3. Chen, R.R., Roundy, R.O., Zhang, R.Q., Janakiraman, G.: Efficient auction mechanisms for supply chain procurement. *Management Science* **51**(3), 467–482 (2005)
4. Cui, G., Yuan, L., Ding, N., Yao, G., Zhu, W., Ni, Y., Xie, G., Liu, Z., Sun, M.: Ultrafeedback: Boosting language models with high-quality feedback (2024)
5. Duan, J., Li, J., Baba, Y., Kashima, H.: A generalized model for multidimensional intransitivity. In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. pp. 840–852. Springer (2017)
6. Duetting, P., Mirrokni, V., Paes Leme, R., Xu, H., Zuo, S.: Mechanism design for large language models. In: *Proceedings of the ACM on Web Conference 2024*. pp. 144–155 (2024)

7. Fu, T., Cai, D., Liu, L., Shi, S., Yan, R.: Disperse-then-merge: Pushing the limits of instruction tuning via alignment tax reduction. arXiv preprint arXiv:2405.13432 (2024)
8. Hong, J., Bhatia, K., Dragan, A.: On the sensitivity of reward inference to misspecified human models. arXiv preprint arXiv:2212.04717 (2022)
9. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2022)
10. Hu, S., Lu, C., Clune, J.: Automated design of agentic systems. arXiv preprint arXiv:2408.08435 (2024)
11. Hu, Y., Li, Q., Ouyang, S., Chen, G., Chen, K., Mei, L., Ye, X., Zhang, F., Liu, Y.: Towards comprehensive preference data collection for reward modeling. arXiv preprint arXiv:2406.16486 (2024)
12. Ji, J., Qiu, T., Chen, B., Zhang, B., Lou, H., Wang, K., Duan, Y., He, Z., Zhou, J., Zhang, Z., et al.: Ai alignment: A comprehensive survey. arXiv preprint arXiv:2310.19852 (2023)
13. Jiang, R., Chen, K., Bai, X., He, Z., Li, J., Yang, M., Zhao, T., Nie, L., Zhang, M.: A survey on human preference learning for large language models. arXiv preprint arXiv:2406.11191 (2024)
14. Lee, A., Miranda, B., Sundar, S., Koyejo, S.: Beyond scale: the diversity coefficient as a data quality metric demonstrates llms are pre-trained on formally diverse data. arXiv preprint arXiv:2306.13840 (2023)
15. Li, D., Wang, Z., Chen, Y., Jiang, R., Ding, W., Okumura, M.: A survey on deep active learning: Recent advances and new frontiers. IEEE Transactions on Neural Networks and Learning Systems (2024)
16. Li, J.: A comparative study on annotation quality of crowdsourcing and llm via label aggregation. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 6525–6529. IEEE (2024)
17. Lin, Y., Lin, H., Xiong, W., Diaoy, S., Liu, J., Zhang, J., Pan, R., Wang, H., Hu, W., Zhang, H., Dong, H., Pi, R., Zhao, H., Jiang, N., Ji, H., Yao, Y., Zhang, T.: Mitigating the alignment tax of rlhf (2024), <https://arxiv.org/abs/2309.06256>
18. Longpre, S., Hou, L., Vu, T., Webson, A., Chung, H.W., Tay, Y., Zhou, D., Le, Q.V., Zoph, B., Wei, J., Roberts, A.: The flan collection: Designing data and methods for effective instruction tuning. In: Proceedings of the Fortieth International Conference on Machine Learning. pp. 22631–22648. Proceedings of Machine Learning Research, PMLR (2023)
19. Lu, C., Lu, C., Lange, R.T., Foerster, J., Clune, J., Ha, D.: The ai scientist: Towards fully automated open-ended scientific discovery. arXiv preprint arXiv:2408.06292 (2024)
20. Mangrulkar, S., Gugger, S., Debut, L., Belkada, Y., Paul, S., Bossan, B.: Peft: State-of-the-art parameter-efficient fine-tuning methods. <https://github.com/huggingface/peft> (2022)
21. Matsushima, H., Noda, S.: Mechanism design with general ex-ante investments. Journal of Mathematical Economics **106**, 102831 (2023)
22. Nisan, N., Ronen, A.: Computationally feasible vcg mechanisms. Journal of Artificial Intelligence Research **29**, 19–47 (2007)
23. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. Advances in neural information processing systems **35**, 27730–27744 (2022)

24. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. In: *Advances in Neural Information Processing Systems*. pp. 53728–53741. Curran Associates, Inc. (2023)
25. Rosset, C., Cheng, C.A., Mitra, A., Santacrose, M., Awadallah, A., Xie, T.: Direct nash optimization: Teaching language models to self-improve with general preferences. *arXiv preprint arXiv:2404.03715* (2024)
26. Shi, C., Liang, P., Wu, Y., Zhan, T., Jin, Z.: Maximizing user experience with llmops-driven personalized recommendation systems. *arXiv preprint arXiv:2404.00903* (2024)
27. Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., Christiano, P.F.: Learning to summarize with human feedback. *Advances in Neural Information Processing Systems* **33**, 3008–3021 (2020)
28. Sun, H., Chen, Y., Wang, S., Chen, W., Deng, X.: Mechanism design for llm fine-tuning with multiple reward models. *arXiv preprint arXiv:2405.16276* (2024)
29. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023)
30. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* **16**(1), 8–37 (1961)
31. Wang, B., Zheng, R., Chen, L., Liu, Y., Dou, S., Huang, C., Shen, W., Jin, S., Zhou, E., Shi, C., et al.: Secrets of rlhf in large language models part ii: Reward modeling. *arXiv preprint arXiv:2401.06080* (2024)
32. Wang, Z., Bi, B., Huang, C., Pentylala, S.K., Zhu, Z.J., Asur, S., Cheng, N.C.: Una: Unifying alignments of rlhf/ppo, dpo and kto by a generalized implicit reward function. *arXiv preprint arXiv:2408.15339* (2024)
33. von Werra, L., Belkada, Y., Tunstall, L., Beeching, E., Thrush, T., Lambert, N., Huang, S.: Trl: Transformer reinforcement learning. <https://github.com/huggingface/trl> (2020)
34. Xu, C., Sun, Q., Zheng, K., Geng, X., Zhao, P., Feng, J., Tao, C., Lin, Q., Jiang, D.: WizardLM: Empowering large pre-trained language models to follow complex instructions. In: *Proceedings of the Twelfth International Conference on Learning Representations* (2024)
35. Zhang, G., Kashima, H.: Batch reinforcement learning from crowds. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 38–51. Springer (2022)
36. Zhang, H., Dong, Y., Xiao, C., Oyamada, M.: Jellyfish: A large language model for data preprocessing. *arXiv preprint arXiv:2312.01678* (2023)
37. Zheng, L., Chiang, W.L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E.P., Zhang, H., Gonzalez, J.E., Stoica, I.: Judging llm-as-a-judge with mt-bench and chatbot arena (2023)
38. Zheng, R., Dou, S., Gao, S., Hua, Y., Shen, W., Wang, B., Liu, Y., Jin, S., Liu, Q., Zhou, Y., et al.: Secrets of rlhf in large language models part i: Ppo. *arXiv preprint arXiv:2307.04964* (2023)
39. Zhou, C., Liu, P., Xu, P., Iyer, S., Sun, J., Mao, Y., Ma, X., Efrat, A., Yu, P., Yu, L., et al.: Lima: Less is more for alignment. *Advances in Neural Information Processing Systems* **36** (2024)
40. Zhou, W., Agrawal, R., Zhang, S., Indurthi, S.R., Zhao, S., Song, K., Xu, S., Zhu, C.: Wpo: Enhancing rlhf with weighted preference optimization. *arXiv preprint arXiv:2406.11827* (2024)