

Beyond Under-Alignment: Atomic Preference Enhanced Factuality Tuning for Large Language Models

Hongbang Yuan^{1,2}, Yubo Chen^{1,2}[†], Pengfei Cao^{1,2}, Zhuoran Jin^{1,2},
Kang Liu^{1,2}, Jun Zhao^{1,2}

¹The Key Laboratory of Cognition and Decision Intelligence for Complex Systems,
Institute of Automation, Chinese Academy of Sciences, Beijing, China

²School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China
{hongbang.yuan, yubo.chen, pengfei.cao, zhuoran.jin, kliu, jzhao}@nlpr.ia.ac.cn

Abstract

Large language models (LLMs) have achieved remarkable success but still tend to generate factually erroneous responses, a phenomenon known as hallucination. A recent trend is to use preference learning to fine-tune models to align with factuality. However, existing work primarily evaluates fine-tuned models on in-domain (ID) datasets and the factuality on out-of-domain (OOD) datasets remains underexplored. In this paper, we conduct a comprehensive evaluation of the factuality of different models tuned by various preference learning algorithms and demonstrate that their performance on OOD datasets either increases minimally or decreases. Subsequently, we reveal that the main cause of model’s failure to uphold factuality under a distribution shift is **under-alignment**, rather than **over-alignment**, by analyzing the token distribution shift of the models before and after tuning. Finally, we propose **APEFT** (Atomic Preference Enhanced Factuality Tuning), a framework that enhances model’s awareness of factuality at the granularity of individual facts. Extensive experiments demonstrate that APEFT improves model performance by an average of 3.45% on both ID and OOD datasets, which is highly effective.

1 Introduction

Large language models (LLMs) have demonstrated surprising abilities (Chen et al., 2024b; Sun et al., 2024) and have achieved impressive advancements in many tasks (Bubeck et al., 2023; Tashia et al., 2024). However, they are still troubled by the issue of **hallucinations**, a phenomenon wherein LLMs generate seemingly convincing but factually erroneous responses (Zhang et al., 2023; Huang et al., 2023), which greatly hinders their deployment in practical scenarios (Cui et al., 2023).

Recently, some studies indicate that LLMs have the potential to avoid generating hallucinated an-

[†]Corresponding author

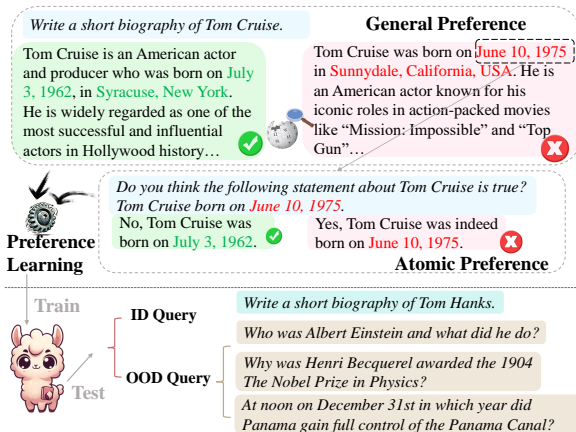


Figure 1: Illustration of our work. The combination of general preferences and atomic preferences can enhance model factuality on OOD queries across various preference learning algorithms.

swers. For example, LLMs mostly ‘know’ the correct answer but fail to ‘tell’ it (Li et al., 2023; Saunders et al., 2022) and they show intrinsic uncertainty during hallucination occurrences (Manakul et al., 2023; Chen et al., 2024a). Therefore, a recent trend is to finetune LLMs to fully exploit their potential, typically using **preference learning** (or **preference tuning**) to steer the models towards desired behaviours and mitigate hallucinations¹ (Tian et al., 2023; Lin et al., 2024). Specifically, it involves collecting diverse model completions, annotating each with an overall preference score related to factuality, and refining models through the preference feedback.

Despite being successful, existing work (Chen et al., 2024c; Tian et al., 2023) mainly uses test queries sourced from the same domain as the training dataset to evaluate the factuality of the tuned model, which is impractical in real-world applications. For example, as shown in Figure 1, when training the model to be more factual in generating

¹In our work, ‘mitigate hallucination’ and ‘increase factuality’ both refer to the same concept.

a biography of Tom Cruise, we expect performance improvements not only on in-domain (ID) queries, like generating a biography of Tom Hanks, but also on out-of-domain (OOD) queries, such as answering questions about Albert Einstein or the Panama Canal. Thus, extensive exploration of the factuality of the tuned models on OOD queries still remains incomplete.

To analyze this generalization problem, we first conduct a comprehensive quantitative study on the OOD performance of the models tuned by various preference learning algorithms. In particular, the training data is constrained to the biography generation task and we train two LLMs on these data using five preference learning algorithms. Subsequently, we evaluate their performance on one ID dataset and OOD datasets. The OOD datasets incorporate various aspects related to model factuality, including generating factually accurate long-form responses consisting of several paragraphs, identifying false premises in user queries and answering with short responses on a scale of several phrases. In the end, the performance either increases minimally, or actually decreases, with the maximum drop in performance reaching 8.47%.

Therefore, we investigate the reasons for the ineffectiveness of the tuned models on OOD queries. We posit that there are two potential failure modes. One possibility is that the tuning process is particularly superficial, resulting in no significant change in the tuned model’s behaviour in OOD settings (**under-alignment**). Alternatively, the tuning process may be adversely affected by spurious features in the training data, resulting in unintended changes in the tuned model’s behaviour (**over-alignment**). For example, when asked about the contribution of Albert Einstein, the tuned models may produce nearly identical responses to those of their pre-tuned counterpart, or tend to generate vague and subjective sentences like ‘his insights are unparalleled’, which is undesirable in this context. To distinguish them, we measure the behaviour change of the model by directly comparing the shift in token distributions before and after tuning. We discover that the number of shifted tokens on OOD queries is less than a third of those on ID queries. *It indicates that the primary cause of the ineffectiveness is under-alignment rather than over-alignment.*

Furthermore, we posit that under-alignment may result from the the original factual preference feedback, which is predominantly based on paragraphs and does not adequately inform the model

about the factuality of individual facts. Therefore, we propose **Atomic Preference Enhanced Factuality Tuning (APEFT)**, a framework that enhances model’s awareness of factuality at the granularity of individual facts. Initially, we define the preferences collected on one specific task as **general preferences**. Firstly, we break the responses in general preferences into single sentences containing only one piece of knowledge. Subsequently, we assess the extent to which the model knows these pieces of knowledge by collecting stochastically sampled responses to a knowledge detection prompt. Then, we select the contradicted responses along with the knowledge detection prompt to construct preference pairs, denoted as **atomic preferences**. Finally, we train the models based on both the atomic preferences and the general preferences. Experimental results demonstrate performance improvements across almost all the preference learning algorithms. Additionally, we briefly explore whether merely increasing the quantity and improving the quality of training preference pairs enhances factuality and find that they do not necessarily lead to performance gains.

Our contributions can be summarized as follows:

- We conduct a comprehensive evaluation of the factuality of different models tuned by various preference learning algorithms. We demonstrate that the model’s performance on OOD queries either increases minimally or even decreases after the tuning.
- We provide two potential failure modes explaining why the models fail to uphold their factuality under a distribution shift: under-alignment and over-alignment. We reveal that under-alignment is the primary cause by analyzing the differences in token distributions of the models before and after tuning.
- We propose **APEFT (Atomic Preference Enhanced Factuality Tuning)**, a framework that enhances model’s awareness of factuality at the granularity of individual facts. Extensive experiments demonstrate its effectiveness and the model performance increases an average of **3.45%** on both the OOD and ID datasets.²

²The code and datasets will be available at <https://github.com/HongbangYuan/APEFT>.

2 Related Work

Hallucination Hallucinations greatly impede the practical applications of LLMs (Liu et al., 2024a; Mündler et al., 2023; Jin et al., 2024a) and many studies have been conducted to address the issue (Trivedi et al., 2023; Gou et al., 2023). Since LLMs have shown internal awareness when hallucination occurs (Azaria and Mitchell, 2023; Manakul et al., 2023; Men et al., 2024), some methods propose to intervene the internal representations during inference to mitigate hallucinations (Li et al., 2023; Chuang et al., 2023; Zou et al., 2023; Jin et al., 2024c) while others directly fine-tune the LLMs to be more factual in the first place (Tian et al., 2023; Chen et al., 2024c; Lin et al., 2024). However, they primarily use test data sourced from the same domain as the training data to evaluate factuality, and fail to consider the generalization of their methods on out-of-domain datasets.

Alignment Aligning LLMs with human preferences is a crucial step to steer already-capable models towards desired behaviors (Burns et al., 2023; Liu et al., 2024c). It typically involves instruction tuning and reinforcement learning from human feedback (Sanh et al., 2022; Ouyang et al., 2022; Bai et al., 2022; Zhou et al., 2023). Since relative preference scores are more feasible than expert demonstrations (Rafailov et al., 2023), many studies develop methods to fine-tune models with preferences (Azar et al., 2023; Ethayarajh et al., 2024; Xu et al., 2024a). However, as human preferences in many datasets tend to under-represent the factuality of the responses (Hosking et al., 2024), it’s still unclear how to construct preference pairs that can enhance the factuality of LLMs.

3 A Comprehensive Evaluation

In this section, we conduct comprehensive quantitative study on the performance of models tuned by various preference learning algorithms. We demonstrate that nearly all the preference learning algorithms deliver unsatisfactory performance on the OOD datasets.

3.1 Data Construction

Initially, we construct a preference dataset that focuses on the factuality of the responses.

Formally, we prompt a pre-trained LLM π_θ with a prompt x to produce pairs of answers: $(y_1, y_2) \sim \pi_\theta(y|x)$. Then we compare which response aligns

more closely with external knowledge sources. The preferred response with fewer factual errors is denoted as y_w , while the dispreferred response with more factual errors is denoted as y_l .

Specifically, we prompt the LLMs to generate biographies of famous people, with names sampled from a list of the most popular individuals³. To measure the factuality of the responses, we employ FActScore (Min et al., 2023), which computes the percentage of individual facts in the generated responses that are supported by the knowledge source. For each prompt, we generate N responses, create $\binom{N}{2}$ preference pairs, and then filter out those pairs that have identical factuality scores. Finally, we obtain a dataset of N preference pairs $D = \{x^i, y_w^i, y_l^i\}_{i=1}^N$.

We employ LLaMA-2-7B-Chat (Touvron et al., 2023) and LLaMA-3-8B-Instruct⁴ as the base model of our experiments. Ultimately, we construct 2777 preference pairs for LLaMA-3-8B-Instruct and 2730 preference pairs for LLaMA-2-7B-Chat.

3.2 Training

We primarily use reward-free preference learning algorithms to avoid the complex process of reinforcement learning without the loss of generality. Specifically, we employ the following representative methods:

(1) **DPO** (Rafailov et al., 2023), which directly optimizes a simple classification loss function without the explicit usage of reward model.

(2) **RSO** (Liu et al., 2024b), which proposes to employ a hinge loss to replace the sigmoid loss in the classification loss function in DPO.

(3) **IPO** (Azar et al., 2024), which designs a new general objective to avoid over-fitting to the preference dataset.

(4) **KTO** (Ethayarajh et al., 2024), which directly maximizes the utility of generated responses instead of the log-likelihood of preferences.

(5) **CPO** (Xu et al., 2024b), which derives an approximation of the DPO loss that is more memory-efficient and speed-efficient.

Further details about these methods are provided in Appendix A.

For the training hyperparameters, we set training epochs to 3, learning rate to $1e^{-6}$, batch size to 4 and gradient accumulation step to 4. We use

³<https://today.yougov.com/ratings/entertainment/popularity/people/all>

⁴<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

Dataset	Samples	Example	Task	Type
Bio	44	Write a short biography of Robert Duvall.	Biography Generation	ID
FAVA	100	Explain Hypermarcas, including information about industry, key person.	Open-ended Generation	OOD
FPQA	986	Why was Albert Einstein awarded the 1922 The Nobel Prize in Physics?	False Premise QA	OOD
KUQA	946	What languages did Mila Kunis speak as a child?	Knowledge-based QA	OOD

Table 1: Statistics and examples of the evaluation datasets.




Figure 2: Performance changes of LLaMA-3-8B-Instruct before and after tuning using various preference learning techniques.

LLaMA-Factory⁵ to train the models using full parameter fine-tuning. All of our experiments are conducted on 4 Nvidia A100-80G GPUs.

3.3 Evaluation

We evaluate the factuality of the tuned models from two aspects: in-domain and out-of-domain. For the in-domain evaluation, we use the same biography generation task as the training data. For the out-of-domain evaluation, we choose the following datasets:

(1) **FAVA** (Mishra et al., 2024), which contains information-seeking queries on open-ended topics, aims to benchmark model’s long-form factuality in a variety of domains.

(2) **FPQA** (Yuan et al., 2024), which contains questions that have false premises, aims to evaluate the model’s factuality when faced with such unanswerable questions.

(3) **KUQA** (Jin et al., 2024b), which contains questions about particular entities and their neighbours, aims to assess the model’s knowledge through its ability to generate short-form phrases.

To facilitate evaluation, we select a subset of each dataset. The statistics and concrete examples of each dataset are shown in Table 1.

⁵<https://github.com/hiyouga/LLaMA-Factory>

3.4 Results and Analysis

We train the models using various preference learning algorithms using our constructed preference pairs. Then, we depict the performance changes of LLaMA-3-8B-Instruct before and after tuning in Figure 2. Additional results are presented in Appendix B. We can draw the following observations: (1) Our constructed preference pairs can improve model factuality on in-domain datasets. For example, apart from the IPO, nearly all the preference learning algorithms demonstrate an improvement in the FActScore on Bio dataset. (2) Nearly all the preference learning algorithms perform poorly on OOD datasets. For example, the performance on FAVA increases minimally while the performance actually decreases on FPQA and KUQA.

4 Under-Aligned or Over-Aligned?

Given the unsatisfactory performance of the tuned models, we set out to better understand the reasons for the ineffectiveness. We start by proposing two potential failure modes, namely under-alignment and over-alignment, and then we conduct the token distribution analysis proposed by Lin et al. (2023) to empirically determine the primary cause.

4.1 Two Failure Modes

Intuitively, the failure may be caused either by insufficient learning from the preference pairs or by excessive attention to spurious features that are unrelated to factuality. We termed these two cases under-alignment and over-alignment. We describe them in detail in the following paragraphs:

Under-Alignment The preference learning process is particularly superficial, resulting in no significant change in the tuned model’s behaviour in OOD settings. For example, the tuned models may produce the same responses as their pre-tuned counterpart because they didn’t fully aware of the factuality conveyed by the training preference pairs.

Over-Alignment The preference learning process is adversely affected by spurious features in the training data, resulting in unintended changes

in model behaviour in OOD settings. For example, the model tends to generate vague and subjective queries after training on the biography generation task, thus failing to answer questions such as the specific contributions of Albert Einstein or when Panama gained full control of the Panama Canal.

Notably, it is significant to distinguish them as it influences how we proceed to address the generalization failure. If the under-alignment is the primary cause, enhancing the model’s awareness of factuality becomes necessary. Conversely, if the over-alignment is the primary cause, it would be imperative to revise the learning process or the current preference pairs to minimize the impact of spurious features.

4.2 Empirical Validation

To quantitatively measure the behaviour change during the preference learning, we directly compare the differences in token distribution of the models before and after tuning.

Method Intuitively, if the model’s behaviour changes significantly after the fine-tuning process, the predicted tokens using the same context will also change dramatically. Specifically, the response of the fine-tuned model is denoted as $\{t_1, t_2, \dots, t_N\}$ and at each token position, the next token prediction probability distribution is denoted as $P_{aligned}$. For each token position t_i , the context $\{t_1, t_2, \dots, t_{i-1}\}$ is input again to the pre-tuned counterpart and the resulting distribution is denoted as P_{base} . We measure how the rank of t_i in $P_{aligned}$ changes compared to its rank in P_{base} and define the difference as δ . The greater the difference, the more significant the change in model behaviour after the tuning. We refer to a tokens as ‘shifted token’ if its rank difference is greater than zero.

Results and Analysis We conduct experiments on the ID dataset Bio and the OOD dataset FAVA on LLaMA-2-7B-Chat trained by various learning algorithms. We calculate the frequency of various levels of the shifted tokens along with their relative positions in each sentence. The experimental results are shown in Figure 3. We draw the following observation: shifted tokens are more prevalent in the in-domain dataset than in the out-of-domain dataset. The frequency bars in the Bio figure are almost three times as high as those in the Fava figure. This indicates that the the model’s behaviour changes far less on OOD queries than on ID queries and the performance drop on OOD

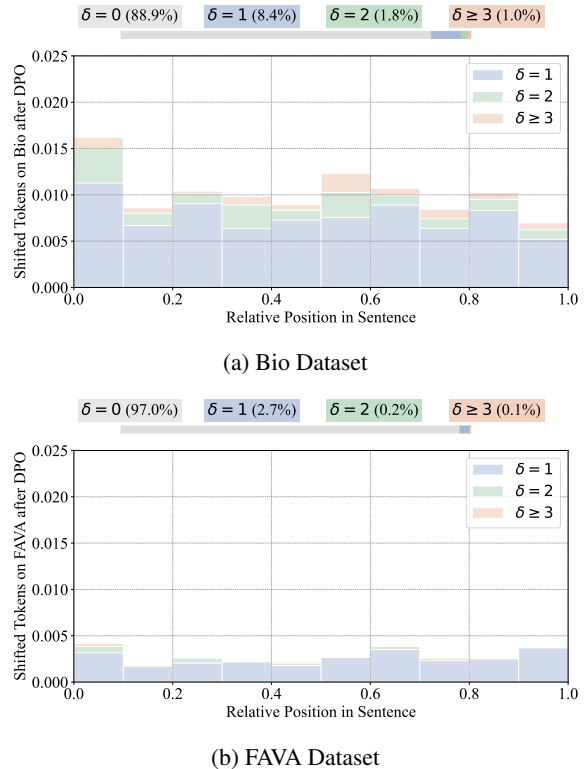


Figure 3: Token shift analysis on LLaMA-2 trained by DPO. More results are presented in Appendix C.

queries is more likely a result of under-alignment rather than over-alignment. Therefore, it’s essential to design methods to enhance model’s awareness of factuality.

5 Atomic Preference Enhanced Factuality Tuning

In this section, we introduce APEFT, a framework designed to enhance the model’s awareness of factuality at the granularity of individual facts. We start with a detailed description of our framework and then demonstrate its effectiveness by extensive quantitative experiments.

5.1 Framework Description

We posit that the original factual preference feedback, which is mainly based on paragraphs, is too coarse-grained and insufficient to convey the factuality of individual facts. For example, as shown in Figure 1, the general preferences contain too much information, such as the date and place of Tom Cruise’s birth, the films in which he starred. Consequently, the model is unable to discern exactly which fact is incorrect, leading to under-alignment on OOD queries. Thus it’s essential to construct preferences at the granularity of each individual

fact. In particular, as shown in Figure 4, our method APEFT is divided into three steps: atomic fact extraction, knowledge detection and atomic preference creation. We will describe each step in detail sequentially.

Atomic Fact Extraction Firstly, we extract individual atomic facts from the general preference dataset. Following (Min et al., 2023; Wei et al., 2024), we define the atomic facts as short sentences containing only one piece of information. We extract atomic facts of both the preferred response y_w and the dispreferred response y_l for every preference pair by prompting ChatGPT.

Knowledge Detection Subsequently, we assess the extent to which model knows the knowledge contained in each atomic fact. We ask the model to decide whether the provided atomic fact is true or false and detect the model’s internal knowledge through the percentage of correct responses among the stochastically sampled responses.

We employ multinomial sampling as the decoding strategy and set temperature to 1 to increase randomness and diversity. With the percentage of the correct responses after the sampling denoted as r , the **unknown**, **potentially-known** and **known** atomic facts are defined as when ($r = 0$), ($0 < r < 1$) and ($r = 1$) respectively. We propose that our focus should be on improving those potentially-known atomic facts.

Atomic Preference Creation Finally, we construct atomic preferences from previously identified potentially-known atomic facts. We direct use the knowledge detection prompt as the user query x and the two contradictory responses as preference (y_w, y_l). The atomic preferences are mixed together with the general preferences to enhance the factuality of LLMs. With the original general preferences denoted as D_g and the newly constructed atomic preferences denoted as D_a , we aim to optimize the following generic objective function:

$$O(\pi_\theta) = \mathbb{E}_{(\mathbf{y}, x) \sim D_g \cup D_a} \left[r(x, \mathbf{y}) - \beta \log \frac{\pi_\theta(\mathbf{y} | x)}{\pi_{\text{ref}}(\mathbf{y} | x)} \right]$$

where the data (x, \mathbf{y}) can either be (x, y_w) or (x, y_l) , r is the reward model and β is a regularisation term controlling the Kullback-Leibler divergence between the policy model π_θ and the reference model π_{ref} . Further details about the objective function are presented Appendix A.

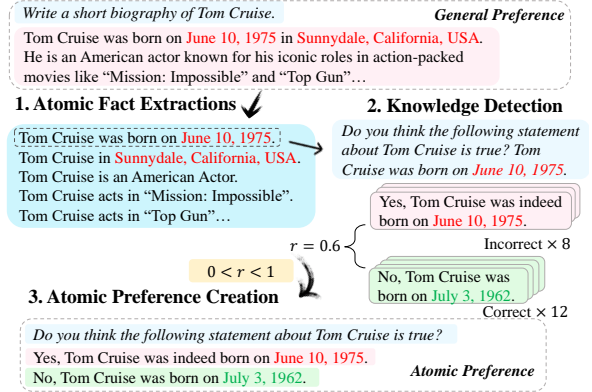


Figure 4: Illustration of our proposed APEFT. It creates atomic preferences from the responses in original general preferences to enhance model’s awareness of factuality at the granularity of individual facts.

5.2 Experimental Results

We create 2063 atomic preferences each for LLaMA-2-7B-Chat and LLaMA-3-8B-Instruct, and train the models using both these atomic preferences and the previously generated general preferences. The training process and the evaluation of the models follow the procedure outlined in Section 3. The experimental results are shown in Table 2 (our method is denoted as ‘ $w/atom$ ’) and we can draw the following observations:

(1) Add atomic preferences to the original general preferences gain performance increase among almost all the learning algorithms. For example, for each preference learning algorithm, our method achieves improvements compared with the general-preference-only baselines. This demonstrates the effectiveness of the atomic preferences in contributing to the enhancement of factuality.

(2) Despite containing knowledge of only one single factoid, the atomic preferences can enhance the factuality in long-form responses. For example, the factuality of LLaMA-2 on the FAVA dataset improves after the incorporation of atomic preferences compared to using only the general preferences.

5.3 Ablation Study

To demonstrate the effectiveness of the selection of atomic facts in the general preferences, we conduct the following ablation study. We randomly select questions from the TriviaQA dataset (Joshi et al., 2017). To align with the atomic preferences, we filter these questions to ensure that they are relevant to the famous individuals in the training data and also contain a single ‘potentially known’

Methods	Bio			FAVA			FPQA	KUQA	AVG(↑)
	FS(%) (↑)	NC(↑)	NE(↓)	FS(%) (↑)	NC(↑)	NE(↓)	Acc(↑)	Acc(↑)	
<i>LLaMA-2-7B-Chat</i>									
Vanilla	82.11	27.41	6.02	60.37	18.46	11.96	22.99	68.09	58.39
+DPO	87.35	29.23	4.30	62.37	19.14	11.53	36.18	65.11	62.75
<i>w/rand</i>	87.45	29.45	4.23	64.66	19.95	10.83	29.22	71.62	63.24
<i>w/atom</i>	88.08	28.23	3.86	63.66	18.90	11.08	39.03	69.56	65.08
+IPO	85.66	29.05	5.05	64.30	19.56	11.32	37.03	64.67	62.92
<i>w/rand</i>	85.55	28.77	4.77	63.25	19.15	11.32	11.29	70.93	57.75
<i>w/atom</i>	88.17	24.66	3.25	65.84	19.88	10.75	30.49	72.58	64.27
+KTO	82.88	28.14	5.75	61.25	18.63	12.12	43.99	60.92	62.26
<i>w/rand</i>	85.34	28.14	4.86	60.79	18.56	12.53	21.84	70.30	59.57
<i>w/atom</i>	84.62	28.43	5.20	63.98	20.12	11.27	40.40	60.52	62.38
+CPO	89.15	28.48	3.52	62.31	19.61	11.97	34.49	67.44	63.35
<i>w/rand</i>	89.42	29.59	3.55	64.92	19.93	11.02	30.70	72.92	64.49
<i>w/atom</i>	89.54	29.89	3.50	65.91	20.07	10.31	30.59	70.81	64.21
+RSO	87.63	28.30	4.02	62.18	19.22	11.69	35.76	65.08	62.66
<i>w/rand</i>	87.59	30.41	4.30	63.23	19.32	11.56	22.15	73.18	61.54
<i>w/atom</i>	89.06	30.16	3.57	65.11	20.07	10.83	39.13	67.57	65.22
<i>LLaMA-3-8B-Instruct</i>									
Vanilla	86.37	33.86	5.43	69.14	26.95	11.97	82.91	74.29	78.18
+DPO	91.94	35.77	3.20	72.32	29.19	10.78	56.01	70.21	72.62
<i>w/rand</i>	91.79	36.61	3.27	71.68	27.02	10.32	50.21	77.90	72.90
<i>w/atom</i>	90.69	35.16	3.64	69.49	25.36	10.53	67.30	77.27	76.19
+IPO	82.81	33.50	7.09	68.11	26.78	11.93	85.76	56.15	73.21
<i>w/rand</i>	85.55	28.77	4.77	67.32	24.25	11.68	81.22	81.11	78.80
<i>w/atom</i>	84.75	34.02	6.16	65.27	24.61	12.24	95.99	72.53	79.63
+KTO	86.17	32.70	5.32	69.24	27.22	12.10	65.93	73.95	73.82
<i>w/rand</i>	85.06	33.00	5.82	65.35	24.02	11.10	48.73	82.44	70.40
<i>w/atom</i>	84.86	33.41	6.02	64.73	24.37	10.29	96.80	77.35	80.93
+CPO	91.25	36.18	3.50	71.43	27.97	11.00	59.92	67.21	72.45
<i>w/rand</i>	91.04	35.34	3.48	69.40	26.22	10.12	56.54	79.47	74.11
<i>w/atom</i>	90.22	35.20	3.80	69.58	26.42	9.98	92.41	75.42	81.91
+RSO	89.82	37.09	4.25	70.93	28.49	11.53	47.26	70.84	69.71
<i>w/rand</i>	90.09	35.14	4.05	69.67	26.49	10.39	37.45	80.18	69.35
<i>w/atom</i>	88.11	33.64	4.64	69.40	26.22	10.12	71.20	80.23	77.23

Table 2: Experimental results of models trained by different preference learning techniques. ‘FS’ denotes FActScore, ‘NC’ denotes the number of correct atomic facts and ‘NE’ denotes the number of error atomic facts. ‘Acc’ stands for ‘Accuracy’ and ‘Rec’ stands for ‘Recall Score’. ‘Avg’ is averaged across all ‘FS’ and ‘Acc’ values. Cells in blue indicate a performance increase, while cells in orange indicate a performance decrease.

fact according to stochastically sampled responses. Therefore, we use each question and its corresponding contradicted responses to construct preference pairs directly. We train the models based on these preferences and the previously generated general preferences (denoted as ‘*w/random*’) and the experimental results are shown in Table 2.

(1) Preferences constructed with the randomly selected questions are not as effective as those constructed with atomic facts selected from the general preferences. For example, the average accuracy of the ‘*w/random*’ method does not exceed that of our ‘*w/atom*’ method across almost all the preference learning algorithms and both models.

(2) Randomly selected QA samples are more beneficial for short-form factuality but contribute

less to longer responses on the scale of several sentences or paragraphs. For example, while the accuracy on the KUQA dataset is generally slightly higher than that achieved with our method, it slightly drops on the other three datasets.

6 Discussion

In this section, we further explore whether merely increasing the quantity or improving the quality of training preference pairs facilitates the preference learning of factuality more effectively. Our main finding is that scaling or a greater difference in preference pairs does not necessarily result in improved factuality performance.




Figure 5: Experimental results of models trained by various number of training preferences.




Figure 6: Experimental results of models trained by preferences of different quality levels using DPO. More results are presented in Appendix D.

6.1 Data Quantity

Initially, we investigate the influence of increasing the number of the training preference pairs. We create different amount of preference pairs and train LLaMA-2-7B-Chat based on them. The experimental results are shown in Figure 5.

We can draw the following observation: scaling is ineffective when training models for factuality. For example, what can be clearly seen in the figure is the dramatic decline in performance on the KUQA dataset as the number of training preference pairs increase. For the Bio and FAVA datasets, the performance quickly reaches a peak and then struggles to improve further. It suggests that a moderate amount of preference pairs is sufficient for the model to learn factuality, and simply adding more data might even backfire.

6.2 Data Quality

Subsequently, we study the influence of data quality in training preference pairs. We define the preference data quality as the difference in fac-

tuality scores between the preferred and dispreferred responses. Formally, the data quality score q of a preference (x, y_w, y_l) is calculated as $q = f(y_w) - f(y_l)$ where $f(y)$ represents the factuality score of model response y .

We divide the general preferences generated in Section 3 into three levels of data groups according to their quality score: Level1 ($0 < q \leq 0.1$), Level2 ($1 < q \leq 0.2$) and Level3 ($q > 0.2$). We have also added a mixed data group encompassing data of all quality levels. To ensure a fair comparison, we ensure that each data group is filtered to have an equal quantity. Then we use DPO to train LLaMA-2-7B-Chat using each data group. The experimental results are shown in Figure 6.

We can draw the following observation: preference pairs of better quality do not necessarily ensure better performance. For example, data group with quality Level2 consistently outperforms others on OOD datasets. One possible reason is that the preference learning is actually aligning the knowledge inside the model parameters. Meanwhile, the quality of the preference pairs is assessed based on whether the responses align with the external knowledge sources, which are not necessarily relevant to the model itself.

7 Conclusion

In this work, we firstly conduct a comprehensive evaluation of the factuality of different models on OOD queries tuned by various preference learning methods and demonstrate their ineffective performance. Subsequently, we conduct a token distribution analysis and reveal that the primary cause of the ineffectiveness is under-alignment rather than over-alignment. Finally, we propose APEFT, a framework that can enhance model’s awareness of

factuality at the granularity of individual facts. Extensive experiments demonstrate the effectiveness of our proposed APEFT.

Limitations

While our work provides an in-depth analysis of how to enhance model’s factuality on OOD queries, it is still subject to the following limitations: (1) Our research primarily concentrates on how to enhance the factuality of LLMs through fine-tuning. However, it naturally raises concerns about whether the same process could adversely affect other capabilities that LLMs have already developed, such as mathematical reasoning, code completion and safety preservation. Future research could systematically investigate the changes in various capabilities of LLMs during the fine-tuning process. (2) Our work does not delve into the internals of LLMs. It would be interesting to develop a more fine-grained understanding of behaviour changes within the LLMs induced by fine-tuning. For instance, is there a tiny region within LLMs focusing on the concept of factuality, and do the primary changes in parameters occur in this area? We leave the exploration for future work.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No.U24A20335, No. 62176257, No. 62406321). This work is also supported by the Youth Innovation Promotion Association CAS and the China Postdoctoral Science Foundation under Grant Number 2024M753500.

References

Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Rémi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. 2024. [A general theoretical paradigm to understand learning from human preferences](#). In *International Conference on Artificial Intelligence and Statistics, 2-4 May 2024, Palau de Congressos, Valencia, Spain*, volume 238 of *Proceedings of Machine Learning Research*, pages 4447–4455. PMLR.

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

Amos Azaria and Tom Mitchell. 2023. [The internal state of an LLM knows when it’s lying](#). In *Findings of the Association for Computational Linguistics:*

EMNLP 2023, pages 967–976, Singapore. Association for Computational Linguistics.

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

Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. 2023. [Sparks of artificial general intelligence: Early experiments with gpt-4](#). *Preprint*, arXiv:2303.12712.

Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. 2023. [Weak-to-strong generalization: Eliciting strong capabilities with weak supervision](#). *Preprint*, arXiv:2312.09390.

Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024a. [Inside: LLMs’ internal states retain the power of hallucination detection](#). *Preprint*, arXiv:2402.03744.

Hailin Chen, Fangkai Jiao, Xingxuan Li, Chengwei Qin, Mathieu Ravaut, Ruo Chen Zhao, Caiming Xiong, and Shafiq Joty. 2024b. [Chatgpt’s one-year anniversary: Are open-source large language models catching up?](#) *Preprint*, arXiv:2311.16989.

Weixin Chen, Dawn Song, and Bo Li. 2024c. [Grath: Gradual self-truthifying for large language models](#). *Preprint*, arXiv:2401.12292.

Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. 2023. [Dola: Decoding by contrasting layers improves factuality in large language models](#). *arXiv preprint arXiv:2309.03883*.

Jiaxi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. [Chatlaw: Open-source legal large language model with integrated external knowledge bases](#). *Preprint*, arXiv:2306.16092.

Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. 2024. [Kto: Model alignment as prospect theoretic optimization](#). *Preprint*, arXiv:2402.01306.

Zhibin Gou, Zhihong Shao, Yeyun Gong, Yelong Shen, Yujia Yang, Nan Duan, and Weizhu Chen.

2023. [Critic: Large language models can self-correct with tool-interactive critiquing](#). *Preprint*, arXiv:2305.11738.
- Tom Hosking, Phil Blunsom, and Max Bartolo. 2024. [Human feedback is not gold standard](#). *Preprint*, arXiv:2309.16349.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *Preprint*, arXiv:2311.05232.
- Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, Xiaojian Jiang, Jiexin Xu, Qiuxia Li, and Jun Zhao. 2024a. [Tug-of-war between knowledge: Exploring and resolving knowledge conflicts in retrieval-augmented language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 16867–16878. ELRA and ICCL.
- Zhuoran Jin, Pengfei Cao, Chenhao Wang, Zhitao He, Hongbang Yuan, Jiachun Li, Yubo Chen, Kang Liu, and Jun Zhao. 2024b. [Rwku: Benchmarking real-world knowledge unlearning for large language models](#). *Preprint*, arXiv:2406.10890.
- Zhuoran Jin, Pengfei Cao, Hongbang Yuan, Yubo Chen, Jiexin Xu, Huaijun Li, Xiaojian Jiang, Kang Liu, and Jun Zhao. 2024c. [Cutting off the head ends the conflict: A mechanism for interpreting and mitigating knowledge conflicts in language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1193–1215, Bangkok, Thailand. Association for Computational Linguistics.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Kenneth Li, Oam Patel, Fernanda B. Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. [Inference-time intervention: Eliciting truthful answers from a language model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandrabhagavatula, and Yejin Choi. 2023. [The unlocking spell on base llms: Rethinking alignment via in-context learning](#). *Preprint*, arXiv:2312.01552.
- Sheng-Chieh Lin, Luyu Gao, Barlas Oguz, Wenhan Xiong, Jimmy Lin, Wen tau Yih, and Xilun Chen. 2024. [Flame: Factuality-aware alignment for large language models](#). *Preprint*, arXiv:2405.01525.
- Hanchao Liu, Wenyuan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024a. [A survey on hallucination in large vision-language models](#). *CoRR*, abs/2402.00253.
- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J. Liu, and Jialu Liu. 2024b. [Statistical rejection sampling improves preference optimization](#). *Preprint*, arXiv:2309.06657.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024c. [What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning](#). *Preprint*, arXiv:2312.15685.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Tianyi Men, Pengfei Cao, Zhuoran Jin, Yubo Chen, Kang Liu, and Jun Zhao. 2024. [Unlocking the future: Exploring look-ahead planning mechanistic interpretability in large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7713–7724, Miami, Florida, USA. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. 2024. [Fine-grained hallucination detection and editing for language models](#). *Preprint*, arXiv:2401.06855.
- Niels Mündler, Jingxuan He, Slobodan Jenko, and Martin Vechev. 2023. [Self-contradictory hallucinations of large language models: Evaluation, detection and mitigation](#). *Preprint*, arXiv:2305.15852.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.

- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- William Saunders, Catherine Yeh, Jeff Wu, Steven Bills, Long Ouyang, Jonathan Ward, and Jan Leike. 2022. [Self-critiquing models for assisting human evaluators](#). *Preprint*, arXiv:2206.05802.
- Tianxiang Sun, Xiaotian Zhang, Zhengfu He, Peng Li, Qinyuan Cheng, Xiangyang Liu, Hang Yan, Yunfan Shao, Qiong Tang, Shiduo Zhang, Xingjian Zhao, Ke Chen, Yining Zheng, Zhejian Zhou, Ruixiao Li, Jun Zhan, Yunhua Zhou, Linyang Li, Xiaogui Yang, Lingling Wu, Zhangyue Yin, Xuanjing Huang, Yutang Jiang, and Xipeng Qiu. 2024. [Moss: An open conversational large language model](#). *Machine Intelligence Research*, 21(5):888–905.
- Dorje Tashia, Tianying Shengc, Bingtian Chenb, Renzeng Duojiea, Rinchen Dongruba, Yongbin Yub, Xi-angxiang Wangb, and Nyima Tashia. 2024. [Construction of a tibetan handwriting khyug-yig dataset](#). *Data Intelligence*.
- Katherine Tian, Eric Mitchell, Huaxiu Yao, Christopher D. Manning, and Chelsea Finn. 2023. [Fine-tuning language models for factuality](#). *Preprint*, arXiv:2311.08401.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Harsh Trivedi, Niranjana Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). *Preprint*, arXiv:2403.18802.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024a. [Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation](#). *Preprint*, arXiv:2401.08417.
- Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024b. [Contrastive preference optimization: Pushing the boundaries of llm performance in machine translation](#). *Preprint*, arXiv:2401.08417.
- Hongbang Yuan, Pengfei Cao, Zhuoran Jin, Yubo Chen, Daojian Zeng, Kang Liu, and Jun Zhao. 2024. [Whispers that shake foundations: Analyzing and mitigating false premise hallucinations in large language models](#). *Preprint*, arXiv:2402.19103.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren’s song in the ai ocean: A survey on hallucination in large language models](#). *Preprint*, arXiv:2309.01219.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: less is more for alignment](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, Zico Kolter, and Dan Hendrycks. 2023. [Representation engineering: A top-down approach to ai transparency](#). *Preprint*, arXiv:2310.01405.

A Preference Learning Algorithms

In this section, we formally introduce the loss functions used by the various preference learning algorithms. Formally, given the preference dataset $D = \{x^i, y_w^i, y_l^i\}_{i=1}^N$, the reference policy model π_θ and the current policy model π_{ref} , the loss functions can be articulated as follows:

(1) DPO (Direct Preference Optimization)

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

where β is a hyperparameter controlling the deviation between the policy model π_θ and the reference model π_{ref} during the optimization process.

(2) IPO (Identity Preference Optimization)

$$\mathcal{L}_{\text{IPO}} = \mathbb{E}_{(y_w, y_l, x) \sim D} \left(h_\pi(y_w, y_l, x) - \frac{\tau^{-1}}{2} \right)^2$$

$$h_\pi(y, y', x) = \log \left(\frac{\pi(y | x) \pi_{\text{ref}}(y' | x)}{\pi(y' | x) \pi_{\text{ref}}(y | x)} \right)$$

where τ is a regularisation term. A lower value of τ means a higher value of the log-likelihood ratio of y_w to y_l .

(3) KTO (Kahneman-Tversky Optimization)

$$L_{\text{KTO}} = \mathbb{E}_{x, y \sim D} [w(y) (1 - v_{\text{KTO}}(x, y; \beta))]$$

where

$$\begin{aligned} r_{\text{KTO}}(x, y) &= \beta \log \frac{\pi_\theta(y | x)}{\pi_{\text{ref}}(y | x)} \\ z_{\text{ref}} &= \mathbb{E}_{x' \sim D} [\beta \text{KL}(\pi_\theta(y' | x') || \pi_{\text{ref}}(y' | x'))] \\ v_{\text{KTO}}(x, y; \beta) &= \begin{cases} \sigma(r_{\text{KTO}}(x, y) - z_{\text{ref}}) & \text{if } y \sim y_{\text{desirable}} | x \\ \sigma(z_{\text{ref}} - r_{\text{KTO}}(x, y)) & \text{if } y \sim y_{\text{undesirable}} | x \end{cases} \\ w(y) &= \begin{cases} \lambda_D & \text{if } y \sim y_{\text{desirable}} | x \\ \lambda_U & \text{if } y \sim y_{\text{undesirable}} | x \end{cases} \end{aligned}$$

Notably, KTO only requires one response and a binary signal indicating whether it is ‘desirable’ or ‘undesirable’. We set the `loss_type` parameter to `kto_pair` in `DPOTrainer` which is a highly simplified version and we leave a more sophisticated implementation for future work.




Figure 7: LLaMA-2-7B-Chat performance change before and after tuning using various preference learning techniques.

(4) CPO (Contrastive Preference Optimization)

$$\begin{aligned} \mathcal{L}_{\text{NLL}} &= -\mathbb{E}_{(x, y_w) \sim D} [\log \pi_\theta(y_w | x)] \\ \mathcal{L}_{\text{prefer}} &= -\mathbb{E}_{(x, y_w, y_l) \sim D} [\log \sigma(\beta \log \pi_\theta(y_w | x) \\ &\quad - \beta \log \pi_\theta(y_l | x))] \\ \mathcal{L}_{\text{CPO}} &= \mathcal{L}_{\text{prefer}} + \mathcal{L}_{\text{NLL}} \end{aligned}$$

where β is a regularisation term. Notably, CPO doesn’t require the reference model. Thus it is more memory-efficient and speed efficient.

(5) RSO (Statistical Rejection Sampling Optimization)

$$\begin{aligned} \mathcal{L}_{\text{rso}} &= \mathbb{E}_{(y_w, y_l, x) \sim D} [\max(0, \\ &\quad 1 - \left[\gamma \log \frac{\pi_\theta(y_w | x)}{\pi_{\text{st}}(y_w | x)} - \gamma \log \frac{\pi_\theta(y_l | x)}{\pi_{\text{st}}(y_l | x)} \right]) \end{aligned}$$


where γ is a regularisation term. In our work, the preferences are sampled and annotated directly from the models. Thus we use a simplified version by setting the supervised finetuning model to the original base reference policy model. This equals to substituting the sigmoid function in DPO loss with a hinge loss. We leave the original implementation for future work.

B Additional Results of General Preference Performance


We present the performance change of LLaMA-2-7B-Chat trained by general preferences using various preference learning techniques.

C Additional Result of Token Shift Analysis

In this section, we present the results of token analysis of LLaMA-2-7B-Chat trained by IPO (Figure 8), KTO (Figure 9), CPO (Figure 10) and RSO (Figure 11).



(a) Bio Dataset




(b) FAVA Dataset


Figure 8: Token shift analysis on LLaMA-2 trained by IPO.

D Additional Results of the Data Quality Experiments

In this section, we present the results demonstrating the impact of data quality on our experiments. The training methods are IPO (Figure 12), KTO (Figure 13), CPO (Figure 14) and RSO (Figure 15).




(a) Bio Dataset




(b) FAVA Dataset

Figure 9: Token shift analysis on LLaMA-2 trained by KTO.



(a) Bio Dataset



(b) FAVA Dataset

Figure 10: Token shift analysis on LLaMA-2 trained by CPO.

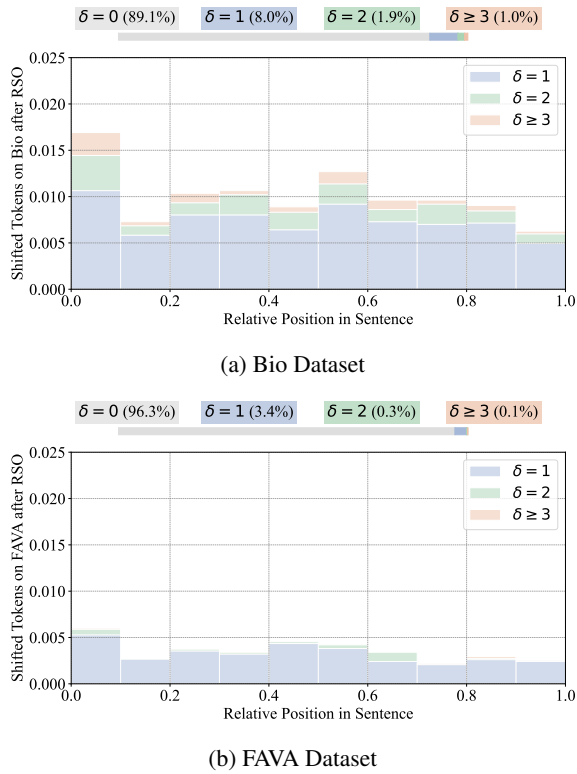


Figure 11: Token shift analysis on LLaMA-2 trained by RSO.




Figure 14: Experimental results of models trained by preferences of different quality levels using CPO.




Figure 12: Experimental results of models trained by preferences of different quality levels using IPO.




Figure 13: Experimental results of models trained by preferences of different quality levels using KTO.

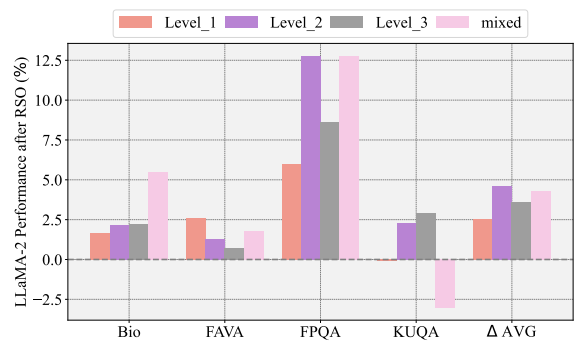


Figure 15: Experimental results of models trained by preferences of different quality levels using RSO.