

Scaling Data Diversity for Fine-Tuning Language Models in Human Alignment

Feifan Song¹, Bowen Yu², Hao Lang², Haiyang Yu²
Fei Huang², Houfeng Wang^{1*}, Yongbin Li^{2*}

¹National Key Laboratory of Multimedia Information Processing
School of Computer Science, Peking University
songff@stu.pku.edu.cn; wanghf@pku.edu.cn

²Alibaba Group
yubowen.ybw, hao.lang, yifei.yhy, f.huang, shuide.lyb@alibaba-inc.com

Abstract

Alignment with human preference prevents large language models (LLMs) from generating misleading or toxic content while requiring high-cost human feedback. Assuming resources of human annotation are limited, there are two different ways of allocating considered: more diverse **PROMPTS** or more diverse **RESPONSES** to be labeled. Nonetheless, a straightforward comparison between their impact is absent. In this work, we first control the diversity of both sides according to the number of samples for fine-tuning, which can directly reflect their influence. We find that instead of numerous prompts, more responses but fewer prompts better trigger LLMs for human alignment. Additionally, the concept of diversity for prompts can be more complex than responses that are typically quantified by single digits. Consequently, a new formulation of prompt diversity is proposed, further implying a linear correlation with the final performance of LLMs after fine-tuning. We also leverage it on data augmentation and conduct experiments to show its effect on different algorithms.

Keywords: Human Alignment, Large Language Model, Scaling Law

1. Introduction

Large Language Models (LLMs) have gained widespread recognition for their proficiency in many domains, including instruction following, intimation and knowledge utilization (Brown et al., 2020; Chung et al., 2022; Muennighoff et al., 2022; Wei et al., 2022; Wang et al., 2022a; Zhou et al., 2022; Von Oswald et al., 2023; Dai et al., 2023; Yang et al., 2023a; Zhong et al., 2023; Schick et al., 2023; Li et al., 2023; Song et al., 2023b; Qin et al., 2023; Wang et al., 2023a; Yang et al., 2023b; Lyu et al., 2024). However, they can reveal toxic or offensive content either inadvertently or intentionally, underscoring the importance of aligning them with human values (Bai et al., 2022b). The transition of paradigm from model-centric to data-centric (Zha et al., 2023b,a) has led to the development of products that are refined using abundant data with human feedback (e.g., ChatGPT, Claude). These products show remarkable capabilities in delivering reliable responses, which prioritizes data collection for LLM fine-tuning aimed at human alignment.

In this field, a natural challenge is the huge expense of high-quality human annotation for diverse samples (Casper et al., 2023). The greater the diversity within the dataset is, the higher upper bound of performance can be achieved. Nevertheless, this diversity also results in higher costs. In detail, LLMs are forced to generate responses in line with

human preference, based on provided prompts. When the annotation resources are limited, a decision must be made regarding the allocation of these resources between a broader range of prompts or a larger number of responses to be annotated, as illustrated in Figure 1.

The well-known LLaMA-2 (Touvron et al., 2023b) chooses to utilize samples of human alignment, each containing one prompt and two responses to maximize the prompt diversity. On the contrary, various studies (Ouyang et al., 2022; Yuan et al., 2023b; Song et al., 2023a) concentrate on providing each prompt with more responses, enabling LLMs to distinguish subtle differences among various candidates. Although both sides are intuitively reasonable, there is currently a lack of direct comparison and comprehensive analyses between them.

In this paper, we investigate the impact of both prompts and responses for LLM fine-tuning in human alignment. We first design a quantitative experiment to assess the effect of the two strategies. A series of sub-datasets are created from the raw dataset, some of which contain more prompts but fewer responses for each prompt, while others have more responses for each prompt but fewer prompts. These subsets maintain a certain proportional relation to maintain a constrained total number of annotations, on which we fine-tune LLMs and compare their performances for comprehensive analyses.

While Song et al. (2023a) has demonstrated the effect of increased responses, a scaling law be-

* Corresponding authors

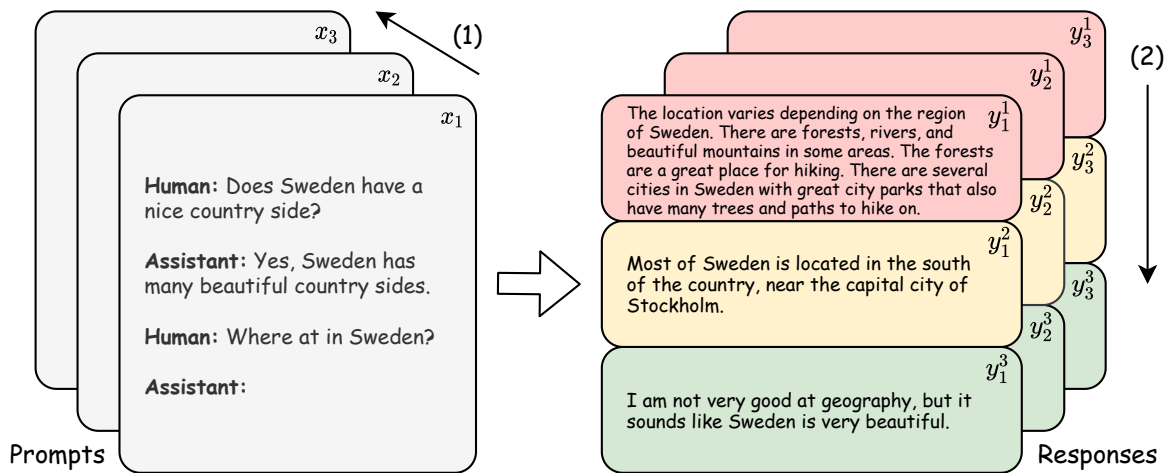


Figure 1: Different directions of data expansion for human alignment: (1) Expanding more prompts; (2) Expanding more responses for each prompt.

tween prompts diversity and the final performance in human alignment is yet to be established. Similar to Kaplan et al. (2020) and Muennighoff et al. (2023) exploring the correlation token statistics and evaluation metrics, the aforementioned quantitative experiment manipulates prompt diversity by adjusting training set sizes only, overlooking the influence of token combinations representing syntax and contextual information. To address this gap, we introduce a novel formulation to empirically define prompt diversity based on N-grams. Furthermore, we uncover a linear relationship between this diversity and the acquired reward scores by examining various scales of training sets, different base models, and algorithms.

We also try to enhance data diversity by employing this new formulation to guide a data augmentation process. Beginning with existing samples, we sample multiple new prompts and corresponding responses, then assess them based on N-gram overlap with the given demonstrations to determine their acceptance. Implementing this method leads to an improvement in performance compared to randomly sampled data.

We conclude from all experiments that:

(1) Expanding responses yields more benefit than prompts. We attribute it to two reasons: just a few prompts can activate LLMs in human alignment, as explained in Zhou et al. (2023), while more responses offer clearer signals for fine-tuning, thus proving more help.

(2) The empirical formulation of prompt diversity can establish a linear correlation with the final performance of LLMs.

(3) Directed by the proposed formulation of prompt diversity, the new process of data augmentation can promote the performance of LLMs.

2. Related Work

2.1. Fine-tuning for Human Alignment

Despite their promising potential, large language models carry the risk of generating toxic or offensive content without human alignment. One approach that has gained considerable attention in addressing this issue is Reinforcement Learning from Human Feedback (RLHF) (Stiennon et al., 2020; Ouyang et al., 2022; Bai et al., 2022a; Zhu et al., 2023a,b; Yu et al., 2023). For instance, InstructGPT (Ouyang et al., 2022) builds a three-step pipeline of RLHF, which includes supervised fine-tuning (SFT), reward model (RM) training, and reinforcement learning using PPO (Schulman et al., 2017). This process involves collecting numerous samples, each consisting of one prompt and multiple candidate responses ranked by human annotators. These annotated rankings are then segmented into pairs to enhance computational efficiency. Touvron et al. (2023b) allocate more resources to the prompt collection to maximize its diversity while featuring only two responses per prompt. Conversely, some works introduce fine-grained distinctions to LLMs by incorporating list-wise comparisons among responses, or dynamically sampling better candidates for SFT (Yuan et al., 2023b; Dong et al., 2023; Song et al., 2023a), also leading to improved performance.

While more prompts can cover a wider range of domains and topics, limitations in annotation resources often force researchers to choose one side between diverse prompts and longer rankings with more responses. In our study, we investigate the impact of prompt diversity and compare it quantitatively with that of responses. We also establish empirical relations between prompt diversity and the final performance of tuned LLMs.

2.2. Scaling Analyses of LLMs

As LLMs continue to increase in scale, leading to higher training costs, it becomes crucial to make initial predictions regarding their performance. Various key factors of LLMs can be scaled to predict the ultimate performance. From a micro perspective, Kaplan et al. (2020) and OpenAI (2023) try to formulate power laws from the model size or the amount of computation of LLMs to their converged loss values during pre-training. In contrast, Lee et al. (2023) examine the impact of different training paradigms for human alignment from a macro perspective. Additionally, Zhang et al. (2023) explore how the assembly of LLMs can influence the final performance, and Yuan et al. (2023a) show that loss values can even indicate the accuracy of mathematical reasoning.

The impact of data used in the pre-training or fine-tuning stages can also be investigated. Kaplan et al. (2020) and Muennighoff et al. (2023) scale the total number of tokens associated with the performance levels achievable by LLMs. However, token count may not perfectly represent the diversity of data distribution. Zhao et al. (2023) accordingly propose a tree-like structure for instruction alignment and study the scaling relationship between the complexity of instructions and final success rates. Lu et al. (2023) and Wei et al. (2023) propose different metrics for the estimation data diversity to label or filter training samples, building upon the observation (Zhou et al., 2023) that a small dataset can unlock the specific capabilities of LLMs through fine-tuning. Building on these studies, we concentrate on the distribution of split prompts and responses for human alignment, and provide detailed analyses of performance improvement influenced by dataset sizes and diversity.

3. Quantitative Experiments

3.1. Background

Different from pre-training, individual samples are typically divided into a prompt x and a response y in LLM fine-tuning. Specifically in human alignment, a single prompt can be associated with multiple responses $y^{1:n} = y^1, y^2, \dots, y^n$, ranked according to varying levels of preference, which are learned by LLMs to enhance their outputs.

Intuitively, a broad range of prompts can potentially enhance the generalization ability of LLMs, thus improving their final performance. Likewise, using diverse candidate responses can be beneficial by enabling LLMs to capture subtle distinctions reflecting different preferences. It is difficult to determine the ideal number of samples and the optimal length of response rankings for LLMs to align with human preference. However, the fact

is that human annotations are always costly, and the total amount of annotations can be accordingly limited. Therefore, given a fixed amount of human annotations, there has to be a trade-off between increasing prompts (while reducing the length of response rankings), or associating each prompt with more responses (but fewer prompts in total) in the dataset. Researchers need to make a choice between these two directions (Ouyang et al., 2022; Yuan et al., 2023b; Dong et al., 2023; Song et al., 2023a; Touvron et al., 2023b).

In this section, we design a quantitative experiment aimed at conducting preliminary comparisons of their effects. We select a series of subsets for fine-tuning, all sharing the same total annotation volume, some emphasizing more prompts while others prioritize more responses. Subsequently, we apply two well-known algorithms to these subsets and aggregate their performance results to assess the impact of different configurations.

3.2. Dataset Construction

Similar to Yuan et al. (2023b), Rafailov et al. (2023) and Song et al. (2023a), we utilize the **Human Preference Data on Helpfulness and Harmlessness**, referred to as *HH-RLHF* (Bai et al., 2022), as the foundational dataset. Each original sample consists of a common prompt and two candidate responses (named 2-ranking), one chosen by human annotators and the other rejected. We extend each 2-ranking into 4-ranking through zero-shot augmentation using Curie (Brown et al., 2020) and Alpaca (Taori et al., 2023), neither of which has been fine-tuned for human alignment previously.

Assuming the total volume of human annotations is $2N$, there are various subsets with different prompt sizes and response ranking lengths. For example, each subset may consist of N prompts with 2-rankings, $2N/3$ prompts with 3-rankings, or $N/2$ prompts with 4-rankings, which all maintain $2N$ annotations ($2N=2 \times N=3 \times 2N/3=4 \times N/2$).

We also attach additional subsets containing N prompts, $2N/3$ prompts, and $N/2$ prompts with 2/3/4-rankings, to present comprehensive results for further analyses.

3.3. Metrics

Unlike some tasks that can be easily measured, human preference can be more abstract and hard to estimate. Both Yuan et al. (2023b) and Song et al. (2023a) utilize RMs to evaluate the performance of fine-tuned LLMs, while the emerging GPT-4-as-a-judge in human alignment (Rafailov et al., 2023) can be also convincing. Our evaluation predominantly relies on public reward models, employing distinct reward models RM_{train} and RM_{test} for training and testing phases, respectively. The outcomes are

Settings (Algorithm, Backbone, Domain)	# Candidate Responses	# Prompts (N=24000)			# Prompts (N=3000)		
		N/2	2N/3	N	N/2	2N/3	N
PRO, OPT-1.3B, Harmless	2	55.58	55.58	57.01	50.42	51.37	53.65
	3	57.11	57.29	59.28	53.28	54.80	56.47
	4	59.24	58.98	59.92	55.36	56.73	58.28
PRO, OPT-1.3B, Helpful	2	49.05	49.09	50.06	44.69	45.37	46.49
	3	49.98	51.00	51.43	45.74	47.01	49.73
	4	51.35	51.04	51.74	48.85	48.69	50.33
PRO, OPT-1.3B, Global	2	52.78	52.73	53.78	47.64	48.57	50.13
	3	53.91	54.67	55.51	49.56	50.76	53.28
	4	55.30	55.12	55.96	52.28	52.54	54.19
SFT, OPT-1.3B, Global	2	52.25	52.78	52.63	49.85	49.33	50.47
	3	53.60	54.18	54.20	51.59	51.26	51.35
	4	55.06	55.00	56.27	51.93	51.97	53.03
PRO, LLaMA-7B, Global	2	54.53	54.68	55.16	53.16	53.06	53.31
	3	56.58	55.89	56.26	54.74	53.00	56.05
	4	57.50	56.42	57.26	55.25	55.88	55.35
SFT, LLaMA-7B, Global	2	53.18	53.29	54.63	52.78	53.03	53.80
	3	56.01	55.53	56.07	53.98	53.74	54.50
	4	56.13	56.30	57.30	55.05	55.00	56.14

Table 1: Results of quantitative experiments. LLMs can acquire better performance with either diverse prompts or responses used for fine-tuning, while increasing responses benefits LLMs more than increasing prompts with the same amount of annotations (highlighted in red bold).

then cross-validated by GPT-4 assessments. For fine-tuning, we combine all 4 subsets in *HH-RLHF* for LLM fine-tuning and present the outcomes for two representative subsets, namely *Harmless*_{base} and *Helpful*_{base}. Furthermore, we provide the overall scores for all test samples across the 4 subsets.

3.4. Benchmark Algorithms

For each dataset, we select representative supervised methods as benchmark algorithms, because supervised training can directly reflect the impact of the datasets involved. Specifically, we opt for two widely-used algorithms, namely **Supervised Fine-tuning** and **Preference Ranking Optimization** (Song et al., 2023a), denoted SFT and PRO, to represent other methods that are either sensitive or insensitive to response rankings.

To elaborate, SFT is similar to the pre-training process but exerts supervision solely on the top candidate y^1 ,

$$\mathcal{L}_{\text{SFT}}(y^1 | x) = - \sum_{i=1}^{|y^1|} \log p_{\Theta}(y_i^1 | x, y_{<i}^1) \quad (1)$$

Instead, PRO forces the LLM to distinguish the best one from multiple candidates. It utilizes the whole ranking y^1, y^2, \dots, y^n through multiple one-to-many contrasts, implemented as:

$$\mathcal{L}(y^{1:n} | x) = - \sum_{k=1}^{n-1} \log \frac{\exp\left(\frac{\log p_{\Theta}(x, y^k)}{\mathcal{T}_k}\right)}{\sum_{i=k}^n \exp\left(\frac{\log p_{\Theta}(x, y^i)}{\mathcal{T}_k}\right)} \quad (2)$$

$$\mathcal{T}_k^{i>k} = \frac{1}{r_{\phi}(x, y^k) - r_{\phi}(x, y^i)} \quad (3)$$

$$\mathcal{T}_k^k = \min_{i>k} \mathcal{T}_k^i \quad (4)$$

and the final objective appends \mathcal{L}_{SFT} for a balance between text quality and human preference,

$$\mathcal{L}_{\text{PRO}}(y^{1:n} | x) = \beta \mathcal{L}_{\text{SFT}}(y^1 | x) + \mathcal{L}(y^{1:n} | x) \quad (5)$$

3.5. Implementation Details

The experiments are conducted with different N (24000 and 3000). We mainly utilized OPT-1.3B (Zhang et al., 2022) as the base LLM and tested it with three different seeds, while incorporating LLaMA-7B (Touvron et al., 2023a) with just one seed due to computational constraints. For the fine-tuning process of LLMs, we configured the total training steps to 4000 for each dataset, performing validation every 500 steps. Both $\text{RM}_{\text{train}}^1$ and $\text{RM}_{\text{test}}^2$ are publicly available checkpoints.

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

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

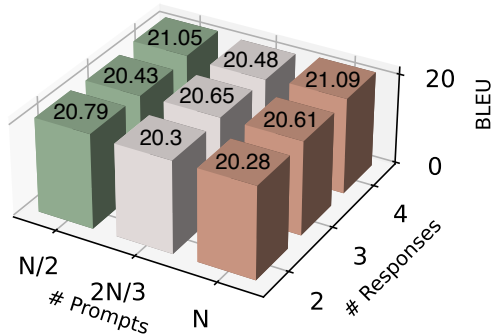


Figure 2: Distribution of BLEU scores with different settings.

The original dataset comprises newly added data without human annotations. Therefore, we first score all responses using RM_{train} , then re-rank them based on their scores. Furthermore, we ensure that datasets of larger sizes will encompass their smaller counterparts with rankings of the same length, while for datasets of the same size, longer rankings will include the shorter ones for each sample. More details are available in the code³.

3.6. Results of Automatic Evaluation

As demonstrated in Table 1, we gathered the average reward scores corresponding to various settings, including algorithms, backbones, and domains. These results aim to address the following research questions (RQs).

RQ1: More Diverse Prompts or Responses?

Longer response rankings and more prompts are both beneficial, but their effects are different. Generally, with an equal quantity of annotations, extending each ranking of responses leads to **better** enhancement compared to expanding prompts, regardless of the backbones used or the values of N . These outcomes are highlighted in red bold within each grid, where scores in the **lower left** side indicate longer response rankings but fewer prompts, while scores in the **upper right** side represent the opposite scenario. This observation is compatible with the hypothesis that LLMs possess the potential for human alignment that can be activated with a small number of samples (Zhou et al., 2023). However, more responses for each prompt offer clearer alignment signals through comparisons, leading to a more significant optimization of LLMs.

In detail, an increased number of responses benefits PRO more than SFT, as the former emphasizes the importance of response rankings lacking

³<https://github.com/F2-song/ScalingAlignment>

in the latter. Nonetheless, a longer ranking represents more samplings from the whole linguistic space, where the preferred response is more likely to be identified. This explains why SFT methods also benefit from expanding responses. Further examination of the results in the domains of harmlessness and helpfulness with OPT-1.3B and PRO confirms the validity of the above statement.

RQ2: How Does the Allocation of Annotations Impact the Quality of Output Texts?

Changes in the allocation of annotations (more prompts or more responses) appear to have no impact on the outcome. We hereby plot the BLEU distribution generated by the OPT-1.3B model fine-tuned with PRO and 24000 samples in Figure 2, where the distribution does not present a consistent pattern but fluctuates randomly across different configurations. The variations observed may be explained by the fact that LLMs possess robust language modeling abilities, consequently not requiring too many samples. It underscores the importance of allocating more annotation resources to prepare responses when a certain number of prompts have been guaranteed.

RQ3: How Many Samples Are Sufficient for Human Alignment Fine-tuning?

Intuitively, the more samples for fine-tuning, the more diverse they are, leading to potentially greater improvements for the tuned LLMs. However, determining the adequate quantity is a complex task, as it depends on factors like the algorithms, base models, and the number of responses. For instance, while LLaMA-7B demonstrates notably high scores with 3000 samples, surpassing OPT-1.3B with an equal amount of training data, it shows a slower increase in performance compared to OPT-1.3B when more samples are included. Moreover, the degree of improvement achieved by increasing the dataset size usually decreases. This is because the dataset is more likely to contain duplicate or similar content as the number of samples grows, making it less efficient to continually invest resources in comparison to potential performance gains.

3.7. GPT-4 Evaluation

Apart from the automatic evaluation above with RM_{test} , we also take GPT-4-as-a-judge into consideration, since it has been widely recognized as an efficiently human-like tool to give fair judgment, especially for abstract concepts like human preference (Zheng et al., 2023; Song et al., 2023a; Dubois et al., 2023). It further validates the statistical findings in Table 1 by directly comparing the three settings as described below:

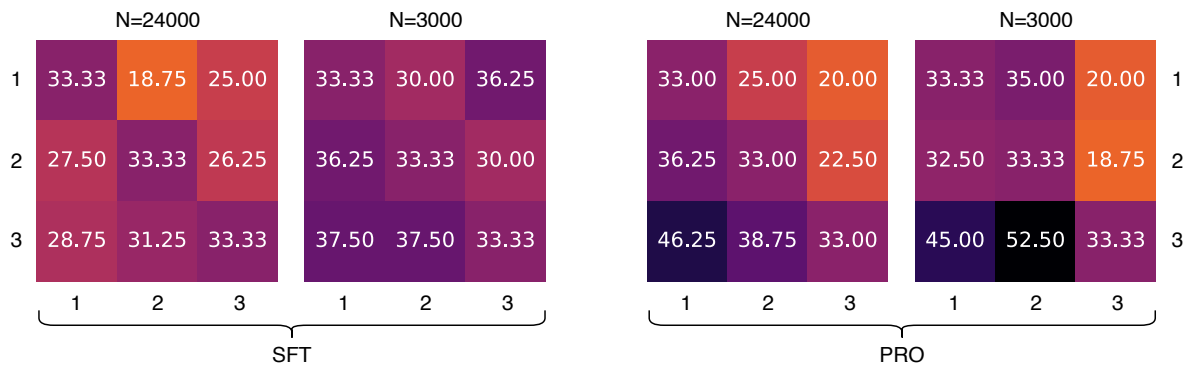


Figure 3: GPT-4 Evaluation

Setting 1: N prompts, i.e., a total of N samples, each with 2 responses.

Setting 2: 2N/3 prompts, each with 3 responses.

Setting 3: N/2 prompts, each with 4 responses.

All three settings have a total of 2N annotations for fine-tuning.

We use LLaMA-7B for fine-tuning under these three settings since it can yield high-quality outputs. We randomly select prompts from the test sets of $Harmless_{base}$ and $Helpful_{base}$ for evaluation. The outputs of each tuned LLaMA are compared with those under other settings, scored directly by GPT-4 through bi-directional comparisons to enhance fairness (Wang et al., 2023b), and the win-lose rates of each comparison are depicted in Figure 3. In each matrix M , each row i (or column j) corresponds to Setting i (or Setting j), with the element $M_{i,j}$ indicating the win rate of LLM outputs tuned with Setting i against those with Setting j . The diagonal elements in M are uniformly set at 33.33 for comparisons between two identical contents, distributed evenly among [Win, Lose, Tie].

Figure 3 illustrates that $M_{i,j}$ always surpasses $M_{j,i}$ along the main diagonal. This implies that the win rate of Setting i against Setting j is always higher than its loss rate, consistent with the results in Table 1. This also proves that RM_{test} can be a reliable evaluator. In general, it reaffirms the conclusion that increasing annotations for responses improves LLMs to better align with human preference than prompts.

4. The Scaling Law between Prompt Diversity and LLMs Preference

While the diversity of prompts or responses can both be beneficial to LLMs fine-tuning, increasing prompts is less effective compared to increasing responses. This difference can be attributed to the inefficiency of using quantity alone to measure prompt diversity. In this section, we explore the concept of prompt diversity. We first discuss the impor-

tance of controlling prompt diversity and then propose a new empirical formulation for it. Additionally, assuming all other factors remain constant (such as base models, fine-tuning algorithms, annotation sources, and lengths of response rankings), a **linear correlation** between the final performances of fine-tuned LLMs and the calculated diversity from various subsets can be illustrated.

4.1. Diversity Formulation

Utilizing quantity to control response diversity appears to be rational. Given that new responses typically originate from various sources, and the range of response rankings is relatively limited, augmenting the number of responses can lead to significant variations. The quantitative experiments also demonstrate that expanding response rankings contributes to improvements for LLMs.

However, the approach becomes oversimplified when applied to prompt diversity. Here the diversity is based on all prompts, and minor adjustments in quantity may not have a noticeable impact, as evidenced by the marginal improvement from N/2 to 2N/3 (for N=3000) in the quantitative experiments, while potential duplication in prompts can be another factor. Furthermore, fine-grained features within utterances, such as semantics, contexts, and even syntax, are crucial to prevalent LLMs that depend on tokenization followed by causal modeling. They should also be taken into consideration.

Moreover, the redundant prompts mentioned earlier do not contribute significantly to enhancing overall diversity and should be initially removed. This can be measured by assessing the proportion of distinct N-grams within the dataset.

Different from Kaplan et al. (2020) and Muenighoff et al. (2023), we leverage N-grams instead of individual tokens as the basic element for calculation, because N-grams inherently capture contextual details beyond the meaning of single tokens. In addition, the aforementioned duplicated data essentially do not provide any extra advantages in

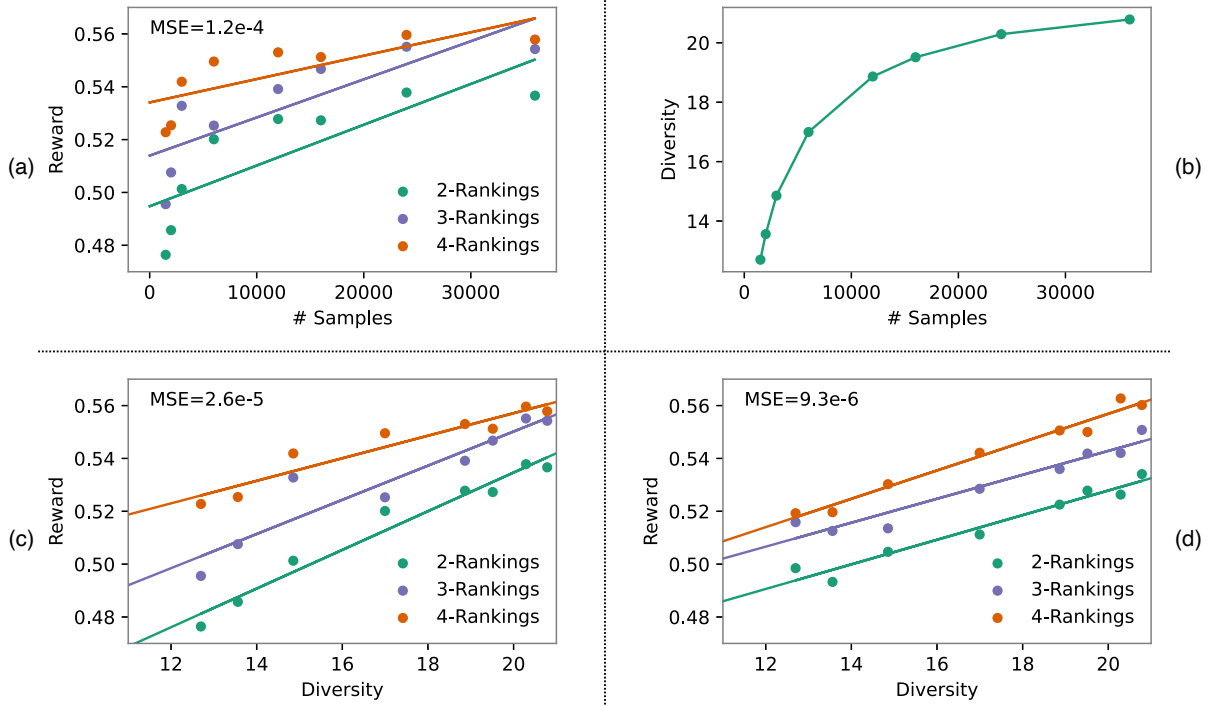


Figure 4: (a) Linear fitting from different sample amounts to finally acquired rewards of LLMs tuned with **PRO**. (b) The trend of diversity with the increasing sample amount. (c) Linear fitting from the proposed diversity metric to finally acquired rewards of LLMs tuned with **PRO**. (d) Linear fitting from the proposed diversity metric to finally acquired rewards of LLMs tuned with **SFT**.

terms of overall diversity and should be filtered out initially. We define it as the rate of unique N-grams present in the dataset,

$$r_{\text{unique}} = \frac{|\text{Filter}(G)|}{|G|} \quad (6)$$

where G is the collection of all N-grams derived from the tokenized corpus and $\text{Filter}(G)$ denotes the removal of repeated elements. Subsequently, the diversity metric d can be formulated as product of r_{unique} and the total number of prompts m ,

$$d = r_{\text{unique}} * m \quad (7)$$

Empirically, as the number of prompts increases, the marginal effect decreases gradually, and the diversity should follow the same pattern. Therefore, we introduced a decay index p to the sample quantity m to incorporate decay into its growth rate. The concept of prompt diversity is formulated as,

$$d = r_{\text{unique}} * m^p \quad (8)$$

4.2. Analysis

To examine the connection between the pre-defined diversity metric and the final performance of fine-tuned LLMs, we use 2-grams for calculation and collect {1500, 2000, 3000, 6000, 12000, 16000, 24000, 36000} samples from the original dataset

with rankings of lengths 2/3/4. This resulted in 24 subsets for LLMs fine-tuning. In this part, we empirically set the value of p to 0.5 for *HH-RLHF*, although it could also be found using grid search.

We start our analysis by visualizing the results of above 24 subsets in Figure 4(a). A discernible positive correlation between enhanced performance and the increasing quantity is observed. Furthermore, improved scores can also be achieved with longer response sequences. Nevertheless, the growth in reward scores and quantity of prompts is not synchronized. The former shows a gradual decline in speed, while the latter maintains a consistent pace. More precisely, the performance experiences a sharp increase with a rising number of samples at the beginning, yet tends to plateau with a larger volume of samples. Even when we convert the X-axis from sample quantity to token quantity, this conclusion still holds.

A commonly accepted concept is that the diversity of a dataset may not continue to increase indefinitely. As the size of the dataset expands, new content often contains complete or partial duplications of earlier material. By analyzing actual datasets, we have graphed in Figure 4(b) the evolution of diversity outlined in Equation 8 as the sample size grows. This graph aligns with the idea that the rate of diversity growth should gradually decrease. Based on the similar patterns observed in performance and

Algorithm 1: Determining the decay index with grid search.

Input: Fine-tuning algorithm FT ,
 Datasets $\{D_i\}$ of ascending sizes,
 Language Model π , Step Length l

Output: Decay Index p

// Fine-tuning and evaluation

- 1 Let S be an empty set
- 2 **for** $D_i \in \{D_1, \dots, D_n\}$ **do**
- 3 Let $\pi_i = FT(\pi, D_i)$
- 4 Evaluate π_i
- 5 Let s_i be the performance of π_i
- 6 Add s_i to S
- 7 **end for**
- // Searching the decay index
- 8 Let $p = 0, \hat{p} = 0, L = \text{inf}$
- 9 **while** $\hat{p} < 1$ **do**
- 10 Let $\hat{p} = \hat{p} + l$
- 11 Compute diversity degrees $\{d_i\}$ based on $\hat{p}, \{D_i\}$ and Equation 8
- 12 Compute the MSE \hat{L} using linear fitting on $\{d_i\}$ and S
- 13 **if** $\{d_i\}$ is ascending && $\hat{L} < L$ **then**
- 14 | Let $L = \hat{L}, p = \hat{p}$
- 15 **end if**
- 16 **end while**
- 17 **return** p

diversity trends, a linear correlation between these them may exist,

$$r = \alpha * d + \beta + \epsilon \quad (9)$$

where α and β are coefficients, while r and ϵ denote the reward score and error term, respectively.

We then gather the performance (i.e., average reward) achieved by OPT-1.3B under the supervision of PRO and SFT, respectively, and then apply linear fitting to correlate each final score with the computed diversity of the specific subset used. The outcomes are presented in Figure 4(c) and (d). Additionally, we compare these results with the linear fitting between performance and the sample quantity, as shown in Figure 4(a), while it leads to a significantly higher mean squared error (MSE) of $1.2e-4$ compared to $2.8e-5$ and $8.8e-6$ corresponding to Figure 4(c) and (d), respectively. This validates the linear correlation between our proposed d and the final performance. We also compute the MSE values with LLaMA-7B, which are marginally higher than those with OPT-1.3B, possibly due to fluctuations in a single seed ($4.6e-5$ for PRO and $1.7e-5$ for SFT).

5. Sampling with Diversity Check

In this section, we present a technique for data augmentation using the existing samples. Fresh samples are first collected and then selected to enrich the overall variety of prompts. This selection is aimed to optimize the local diversity between the new and existing samples. We initially demonstrate the effectiveness of this technique in enhancing prompt diversity. By fine-tuning LLMs on the augmented datasets, there is also a slight performance enhancement along with increased diversity.

5.1. Augmentation

We design the data augmentation as where there are existing samples that constitute a seed set. In this setting, n samples are randomly selected to support each augmentation iteration. To simplify the experiment, we reuse one subset D in the last section as the aforementioned seed set, and select new samples from the original *HH-RLHF* to simulate the process of augmentation. It is ensured that each new sample is selected from the non-overlapping portion of *HH-RLHF* concerning D .

5.2. Filtering with Diversity

We start by revisiting the concept of prompt diversity. The proposed metric above can be affected by two factors: r_{unique} , representing the ratio of unique N-grams, and the total number of prompts, which rises with decreased speed. Consequently, by simultaneously increasing r_{unique} during data augmentation, the diversity metric can experience a more rapid growth.

However, identifying a batch of new samples that maximizes the diversity of the total D can be challenging. Therefore, we introduce a locally greedy search process to filter new samples based on the supporting samples. Specifically, by computing the Jaccard Index between the set X of supporting samples and the i -th element Y_i of set Y ,

$$\text{Jaccard}_i = \frac{\text{Filter}(G_X) \cap \text{Filter}(G_{Y_i})}{\text{Filter}(G_X) \cup \text{Filter}(G_{Y_i})} \quad (10)$$

where Y_k with the lowest Jaccard_i can be selected to enhance the local r_{unique} , thereby boosting the overall diversity.

5.3. Results

We set n as 2 and utilize the subset D_{6000} containing 6000 samples (for all versions with 2/3/4 rankings), to which 6000 new samples are then added, forming \hat{D}_{12000} . Meanwhile, we treat the subset D_{12000} as its counterpart, comprising 12000 randomly sampled samples without filtration and

Method	# Responses	D_{12000}	\Rightarrow	\hat{D}_{12000}
PRO	2	52.78		52.85
	3	53.91	\Rightarrow	54.55
	4	55.30		55.50
SFT	2	52.25		51.81
	3	53.60	\Rightarrow	54.47
	4	55.06		55.20
Diversity d of Prompts		18.86	\Rightarrow	19.75

Table 2: Results of Data Augmentation.

covering D_{6000} . We apply PRO and SFT on OPT-1.3B using these subsets and present the outcomes in Table 2. The diversity of prompts increases from 18.86 to 19.75, and slight enhancements are shown with few exceptions, which can be normal fluctuations. This primarily demonstrates the impact of the proposed filtering method. It may potentially be amplified with larger n , a direction we leave for future research.

6. Conclusion

This study focuses on the impact of data diversity on human alignment fine-tuning. Given the common limitation of available annotations in most scenarios, we investigate the effect of distributing them to enhance diversity in different ways, such as increasing prompts or responses. Our extensive experiments show that increasing the number of responses generally leads to greater enhancements in human alignment compared to expanding prompts. Additionally, we design an empirical metric to measure prompt diversity and reveal a linear correlation between it and the final performance of LLMs. Finally, we propose a straightforward method to boost diversity in data augmentation, resulting in better performance of fine-tuned LLMs.

7. Ethical Statement

The presence of sensitive and offensive content within the datasets should be acknowledged. It is important to highlight that these contents do not reflect our views or beliefs, but are solely intended for research purposes.

Moreover, Curie inference and GPT-4 evaluation are utilized where they are available to adhere to legal requirements.

8. Acknowledgement

This work was supported by National Science and Technology Major Project (2022ZD0116308) and National Natural Science Foundation of China (62036001).

9. Bibliographical References

- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2022. [What learning algorithm is in-context learning? investigations with linear models](#). *arXiv preprint arXiv:2211.15661*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. [Qwen technical report](#). *arXiv preprint arXiv:2309.16609*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022a. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#).
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022b. [Constitutional ai: Harmlessness from ai feedback](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in neural information processing systems*, 33:1877–1901.
- Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. 2023. [Open problems and fundamental limitations of reinforcement learning from human feedback](#). *arXiv preprint arXiv:2307.15217*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. [Scaling instruction-finetuned language models](#). *arXiv preprint arXiv:2210.11416*.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. 2023. [Why can gpt learn in-context? language models implicitly perform gradient descent as meta-optimizers](#). In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- Hanze Dong, Wei Xiong, Deepanshu Goyal, Rui Pan, Shizhe Diao, Jipeng Zhang, Kashun Shum, and Tong Zhang. 2023. [Raft: Reward ranked finetuning for generative foundation model alignment](#).

- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#).
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *arXiv preprint arXiv:2001.08361*.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. 2023. [Rlaif: Scaling reinforcement learning from human feedback with ai feedback](#). *arXiv preprint arXiv:2309.00267*.
- Minghao Li, Feifan Song, Bowen Yu, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023. [Api-bank: A benchmark for tool-augmented llms](#). *arXiv preprint arXiv:2304.08244*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, and Chang Zhou. 2023. [# instag: Instruction tagging for diversity and complexity analysis](#). *arXiv preprint arXiv:2308.07074*.
- Youngang Lyu, Lingyong Yan, Shuaiqiang Wang, Haibo Shi, Dawei Yin, Pengjie Ren, Zhumin Chen, Maarten de Rijke, and Zhaochun Ren. 2024. [Knowtuning: Knowledge-aware finetuning for large language models](#). *arXiv preprint arXiv:2402.11176*.
- Niklas Muennighoff, Alexander M Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Noumane Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. 2023. [Scaling data-constrained language models](#). *arXiv preprint arXiv:2305.16264*.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, et al. 2022. [Crosslingual generalization through multitask finetuning](#). *arXiv preprint arXiv:2211.01786*.
- OpenAI. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. [Toolm: Facilitating large language models to master 16000+ real-world apis](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#).
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). *arXiv preprint arXiv:2302.04761*.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#).
- Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. 2023a. [Preference ranking optimization for human alignment](#).
- Yifan Song, Weimin Xiong, Dawei Zhu, Cheng Li, Ke Wang, Ye Tian, and Sujian Li. 2023b. [Restgpt: Connecting large language models with real-world applications via restful apis](#). *arXiv preprint arXiv:2306.06624*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Simeng Sun, Yang Liu, Dan Iter, Chenguang Zhu, and Mohit Iyyer. 2023. [How does in-context learning help prompt tuning?](#) *arXiv preprint arXiv:2302.11521*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava,

- Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. [Transformers learn in-context by gradient descent](#). In *International Conference on Machine Learning*, pages 35151–35174. PMLR.
- Peiyi Wang, Lei Li, Liang Chen, Feifan Song, Binghuai Lin, Yunbo Cao, Tianyu Liu, and Zhifang Sui. 2023a. [Making large language models better reasoners with alignment](#). *arXiv preprint arXiv:2309.02144*.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023b. [Large language models are not fair evaluators](#). *arXiv preprint arXiv:2305.17926*.
- Xinyi Wang, Wanrong Zhu, and William Yang Wang. 2023c. [Large language models are implicitly topic models: Explaining and finding good demonstrations for in-context learning](#). *arXiv preprint arXiv:2301.11916*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022a. [Self-consistency improves chain of thought reasoning in language models](#). *arXiv preprint arXiv:2203.11171*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022b. [Self-instruct: Aligning language model with self generated instructions](#). *arXiv preprint arXiv:2212.10560*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Lai Wei, Zihao Jiang, Weiran Huang, and Lichao Sun. 2023. [Instructiongpt-4: A 200-instruction paradigm for fine-tuning minigpt-4](#). *arXiv preprint arXiv:2308.12067*.
- Jiayi Yang, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. 2023a. [Iterative forward tuning boosts in-context learning in language models](#). *arXiv preprint arXiv:2305.13016*.
- Zhe Yang, Damai Dai, Peiyi Wang, and Zhifang Sui. 2023b. [Not all demonstration examples are equally beneficial: Reweighting demonstration examples for in-context learning](#). *arXiv preprint arXiv:2310.08309*.
- Tianshu Yu, Ting-En Lin, Yuchuan Wu, Min Yang, Fei Huang, and Yongbin Li. 2023. [Constructive large language models alignment with diverse feedback](#). *arXiv preprint arXiv:2310.06450*.
- Zheng Yuan, Hongyi Yuan, Chengpeng Li, Guanting Dong, Chuanqi Tan, and Chang Zhou. 2023a. [Scaling relationship on learning mathematical reasoning with large language models](#). *arXiv preprint arXiv:2308.01825*.
- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023b. [Rrhf: Rank responses to align language models with human feedback without tears](#).
- Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, and Xia Hu. 2023a. [Data-centric ai: Perspectives and challenges](#). In *SDM*.
- Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. 2023b. [Data-centric artificial intelligence: A survey](#). *arXiv preprint arXiv:2303.10158*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuhui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. [Opt: Open pre-trained transformer language models](#). *arXiv preprint arXiv:2205.01068*.
- Xinghua Zhang, Bowen Yu, Haiyang Yu, Yangyu Lv, Tingwen Liu, Fei Huang, Hongbo Xu, and Yongbin Li. 2023. [Wider and deeper llm networks are fairer llm evaluators](#). *arXiv preprint arXiv:2308.01862*.
- Yingxiu Zhao, Bowen Yu, Binyuan Hui, Haiyang Yu, Fei Huang, Yongbin Li, and Nevin L Zhang. 2023. [A preliminary study of the intrinsic relationship between complexity and alignment](#). *arXiv preprint arXiv:2308.05696*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- Wanjuan Zhong, Yifan Gao, Ning Ding, Zhiyuan Liu, Ming Zhou, Jiahai Wang, Jian Yin, and Nan Duan. 2023. [Improving task generalization via unified schema prompt](#). *AI Open*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. [Lima: Less is more for alignment](#). *arXiv preprint arXiv:2305.11206*.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2022. *Large language models are human-level prompt engineers*. *arXiv preprint arXiv:2211.01910*.

Banghua Zhu, Jiantao Jiao, and Michael I Jordan. 2023a. *Principled reinforcement learning with human feedback from pairwise or k -wise comparisons*.

Banghua Zhu, Hiteshi Sharma, Felipe Vieira Fruleri, Shi Dong, Chenguang Zhu, Michael I Jordan, and Jiantao Jiao. 2023b. *Fine-tuning language models with advantage-induced policy alignment*.

10. Language Resource References

Bai, Yuntao and Jones, Andy and Ndousse, Kamal and Askell, Amanda and Chen, Anna and DasSarma, Nova and Drain, Dawn and Fort, Stanislav and Ganguli, Deep and Henighan, Tom and others. 2022. *Training a helpful and harmless assistant with reinforcement learning from human feedback*. PID <https://github.com/anthropics/hh-rlhf>.