

# D<sub>3</sub>: Diversity, Difficulty, and Dependability-Aware Data Selection for Sample-Efficient LLM Instruction Tuning

Jia Zhang<sup>1,2,3\*</sup>, Chen-Xi Zhang<sup>1,2,3\*</sup>, Yao Liu<sup>3</sup>, Yi-Xuan Jin<sup>3</sup>, Xiao-Wen Yang<sup>1,2</sup>,  
Bo Zheng<sup>3</sup>, Yi Liu<sup>3</sup> and Lan-Zhe Guo<sup>1,4</sup>

<sup>1</sup>National Key Laboratory for Novel Software Technology, Nanjing University

<sup>2</sup>School of Artificial Intelligence, Nanjing University

<sup>3</sup>Algorithm Tech, Taobao & Tmall Group of Alibaba

<sup>4</sup>School of Intelligence Science and Technology, Nanjing University  
{zhangjia, zhangcx}@lamda.nju.edu.cn, guolz@nju.edu.cn

## Abstract

Recent advancements in instruction tuning for large language models (LLMs) suggest that a small, high-quality dataset can significantly equip LLMs with instruction-following capabilities, outperforming large datasets often burdened by quality and redundancy issues. However, the challenge lies in automatically identifying valuable subsets from large datasets to boost both the effectiveness and efficiency of instruction tuning. In this paper, we first establish data selection criteria based on three distinct aspects of data value: diversity, difficulty, and dependability, and then propose the D<sub>3</sub> method comprising two key steps of scoring and selection. Specifically, in the scoring step, we define the diversity function to measure sample distinctiveness and introduce the uncertainty-based prediction difficulty to evaluate sample difficulty by mitigating the interference of context-oriented generation diversity. Additionally, we integrate an external LLM for dependability assessment. In the selection step, we formulate the D<sub>3</sub> weighted coresets objective, which jointly optimizes three aspects of data value to solve for the most valuable subset. The two steps of D<sub>3</sub> can iterate multiple rounds, incorporating feedback to refine the selection focus adaptively. Experiments on both public datasets and the real-world Taobao Live application demonstrate the effectiveness of D<sub>3</sub> in endowing LLMs with competitive or even superior instruction-following capabilities using less than 10% of the entire dataset.

## 1 Introduction

With the success of large language models (LLMs) [Brown *et al.*, 2020; Touvron *et al.*, 2023], the application of data-driven AI has been significantly advanced across various industries and scenarios. In the LLM paradigm, large-scale pre-training on massive corpora equips models with knowledge and text generation capabilities based on next-token prediction [Naveed *et al.*, 2024]. Moreover, several post-training

methods and techniques [Bai *et al.*, 2022; Wang *et al.*, 2024b; Rafailov *et al.*, 2024] could further enhance the model’s alignment and performance on downstream tasks. In particular, instruction tuning [Ouyang *et al.*, 2022; Taori *et al.*, 2023], which fine-tunes models with labeled pairs of instruction and response, plays a crucial role in enabling LLMs’ instruction-following capabilities across various tasks.

Early studies on instruction tuning concentrated on creating extensive and cross-task datasets to enhance LLMs at a significant cost [Wang *et al.*, 2022; Xu *et al.*, 2023; Wang *et al.*, 2023]. However, recent work [Zhou *et al.*, 2023] has shown that a small, carefully curated set of high-quality data can effectively equip LLMs with instruction-following capabilities instead of relying on large datasets. This finding has prompted initial efforts in automated data selection [Xia *et al.*, 2024; Cao *et al.*, 2024; Du *et al.*, 2023; Li *et al.*, 2024a; Li *et al.*, 2024b], which attempt to identify the valuable subset of instruction data for fine-tuning. Despite their potential, these methods generally face three limitations: 1) Many approaches rely on a single metric for data selection, neglecting comprehensive assessments of instruction data. 2) Due to the inherent difference between inter- and intra-sample metrics like diversity and quality, existing methods often heuristically consider each in separate stages with manual thresholds to select data, hindering the dynamic balance of different aspects. 3) Lack of feedback to adaptively refine the selection focus.

To address the challenges above, we first establish data selection criteria based on three distinct aspects: diversity, difficulty, and dependability, corresponding to inter-sample, sample-model, and intra-sample perspectives of data value, respectively. Building upon these criteria, we propose D<sub>3</sub>, an effective data selection method consisting of two key steps: scoring and selection. In the scoring step, specifically, to cope with the commonly observed context-oriented generation diversity (CoGD) in text generation, we introduce the uncertainty-based prediction difficulty (UPD), a novel score that quantifies the prediction difficulty of each sample by mitigating the interference caused by CoGD. For diversity and dependability, we define the diversity function based on sample embeddings and use a teacher model to assess the sufficiency and reliability of both the sample’s instruction and its annotated response. In the selection step, we formulate

\*Work done during the authors’ internship at Alibaba Group.

the  $D_3$  weighted coreset object, which jointly optimizes the three aspects of the data value to solve for the most valuable data subset. Additionally, the scoring and selection steps can iterate through multiple rounds. This allows for adaptive refinement of the selection focus by incorporating sample-level feedback from the model during instruction tuning.

We conducted experiments on public datasets and real-world application to validate the  $D_3$  method. Remarkably, with only 5% to 10% of the entire dataset, our approach achieves competitive or even superior performance compared to full-dataset fine-tuning. This highlights the feasibility of  $D_3$  for sample-efficient instruction tuning, which enhances both efficacy and efficiency, thereby advancing related research. To summarize, our main contributions are as follows:

- (1) **Problem:** We establish data selection criteria based on three distinct aspects: diversity, difficulty, and dependability, corresponding to inter-sample, sample-model, and intra-sample perspectives of data value, respectively, which offers systematic guidance for the problem of data selection to enable sample-efficient instruction tuning.
- (2) **Method:** We propose the  $D_3$  approach, consisting of two key steps. In scoring, we define the diversity function, introduce the uncertainty-based prediction difficulty and integrate a teacher model for dependability assessment. In selection, we formulate the  $D_3$  weighted coreset object to jointly optimize three criteria and solve for the most valuable data subset. The two steps can iterate multiple rounds to adaptively refine the selection focus.
- (3) **Evaluation:** Extensive experiments validate the effectiveness of our proposal, demonstrating its potential for sample-efficient instruction tuning to enhance both efficacy and efficiency, thereby facilitating related research.

## 2 Related Work

### 2.1 Instruction Tuning for LLMs

Instruction tuning [Brown *et al.*, 2020; Touvron *et al.*, 2023] has emerged as a widely adopted post-training strategy to equip language models with instruction-following capabilities for various downstream tasks. Early research [Khashabi *et al.*, 2020; Ye *et al.*, 2021] primarily focused on manually creating task-specific instruction tuning datasets. More recent approaches [Honovich *et al.*, 2022; Wang *et al.*, 2022; Xu *et al.*, 2023; Wang *et al.*, 2023; Taori *et al.*, 2023] shift toward constructing instruction tuning datasets by leveraging powerful LLMs to generate data. These studies facilitate the automatic data creation via distilling knowledge from the teacher model [Wang *et al.*, 2024a; Long *et al.*, 2024]. These foundational studies demonstrate that large task-specific or cross-task instruction tuning datasets can boost the instruction-following performance of language models.

### 2.2 Sample-Efficient Tuning

However, large instruction tuning datasets are often plagued by concerns over quality and redundancy. Seminal studies [Zhou *et al.*, 2023] have suggested that a small, manually curated set of high-quality data can endow the model with powerful instruction-following capabilities, motivating

recent research to select valuable subsets for training from large datasets. These efforts generally focus on quality and diversity assessments. [Du *et al.*, 2023] employs a reward model to assess sample quality, while [He *et al.*, 2024] introduces the use of Shapley value to measure the marginal contribution of samples. [Li *et al.*, 2024b; Li *et al.*, 2024a] propose the IFD score to measure the model’s difficulty in instruction following to samples. [Kung *et al.*, 2023] adopts an active learning framework to select samples based on prompt uncertainty. For diversity, k-means or k-center [Sener and Savarese, 2018] are typically employed to select data [Ge *et al.*, 2024; Du *et al.*, 2023]. However, many approaches rely on a single measurement for data selection, neglecting a comprehensive assessment of instruction samples. Additionally, due to the inherent difference between inter-sample metrics like diversity and intra-sample metrics like quality, existing methods typically consider each aspect separately to select data with manual thresholds, lacking a unified way to jointly optimize multiple evaluation criteria. Besides, most methods lack feedback mechanisms to refine the selection focus. Moreover, the evaluation process in some methods is computationally expensive, such as using LLMs to generate responses for all training data, resulting in inefficiencies.

## 3 Methodology: $D_3$

### 3.1 Problem Formalization

In the task of data selection for the sample-efficient LLM instruction tuning, we are provided with an instruction tuning data pool  $\mathcal{D} = \{z^{(i)}\}$  consisting of paired instruction and response samples  $z^{(i)} = (x^{(i)}, y^{(i)})$ , and a pre-trained LLM  $F$  to be fine-tuned on  $\mathcal{D}$ . Our goal is to select a data coreset  $\mathcal{Q}$  from the entire dataset  $\mathcal{D}$  under a selection budget constraint of  $|\mathcal{Q}|/|\mathcal{D}| = k$ , such that the model  $F_{\mathcal{Q}}$ , fine-tuned from  $F$  on  $\mathcal{Q}$ , achieves optimal performance on the test dataset  $\mathcal{D}_t$ .

### 3.2 Overall Framework

To tackle the challenge of identifying and selecting the most valuable coreset for instruction tuning, we first establish data selection criteria based on three distinct aspects: diversity, difficulty, and dependability, corresponding to inter-sample, sample-model, and intra-sample perspectives of data value, respectively. We then propose the novel  $D_3$  method for selecting data based on these criteria. As illustrated in Figure 1, our method comprises a warm-up step and two key steps of scoring and selection. We first utilize a tiny portion of data to obtain the initially acclimated model  $F_0$ . In the scoring step, we evaluate samples in  $\mathcal{D}$  across three criteria. Specifically, we define the diversity function to measure sample distinctiveness and introduce the uncertainty-based prediction difficulty (UPD) to assess sample difficulty by mitigating the interference of context-oriented generation diversity. Additionally, we incorporate a teacher LLM  $F^*$  to evaluate sample dependability. In the selection step, we formulate the  $D_3$  weighted coreset objective to solve for the most valuable subset for LLM tuning within the selection budget. We integrate the scoring and selection steps into an iterative process that leverages the feedback from LLM to enable the adaptive refinement of the selection focus over multiple rounds.

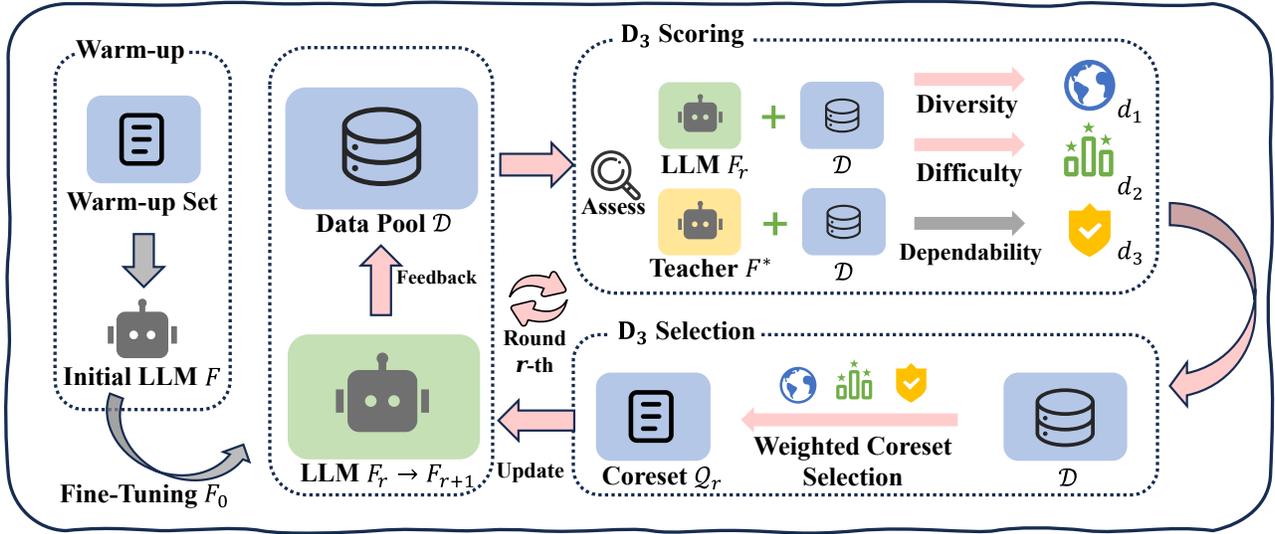


Figure 1: The overall framework of  $D_3$  data selection method, including the warm-up and two key steps of scoring and selection.

### 3.3 $D_3$ Criteria and Scoring

Unlike large-scale pre-training on vast data, we seek to equip LLMs with instruction-following capabilities using much fewer labeled samples. To identify the most valuable data subset, we establish data selection criteria based on three distinct aspects of diversity, difficulty, and dependability. Diversity, from an inter-sample perspective, quantifies the distinctiveness of each sample against others. Difficulty, a sample-model metric, gauges the challenge for the language model in fitting and predicting a given sample. Dependability serves as an intra-sample criterion to assess whether the sample itself is sufficient and reliable enough for training. These criteria offer a comprehensive evaluation of instruction samples from three independent perspectives, providing the systematic guidance for data selection to identify the most valuable data coreset with maximal diversity, difficulty, and dependability.

#### Sample Diversity

The selection of the instruction tuning data coreset should account for its diversity characteristics and representativeness, ensuring that the selected samples accurately capture the overall distribution of the entire dataset  $\mathcal{D}$ . This helps mitigate the risk of the model over-fitting to specific task types. We define the diversity function  $d_1$  of any sample  $z^{(i)}$  as the distance to its closest point in the coreset  $\mathcal{Q}$ , which quantifies the distinctiveness of the sample:

$$d_1(z^{(i)}, \mathcal{Q}) = \min_{z' \in \mathcal{Q}} \mathbf{d}(z^{(i)}, z'), \quad (1)$$

where  $\mathbf{d}(\cdot, \cdot)$  denotes the distance between two samples. Therefore, selecting a coreset with maximal diversity and representativeness with respect to  $\mathcal{D}$  is equivalent to identifying  $\mathcal{Q}$  that minimizes the maximum diversity across all samples:

$$\arg \min_{\mathcal{Q} \subseteq \mathcal{D}} \max_{z \in \mathcal{D}} d_1(z, \mathcal{Q}) = \max_{z \in \mathcal{D}} \min_{z' \in \mathcal{Q}} \mathbf{d}(z, z'). \quad (2)$$

Minimizing  $\max_{z \in \mathcal{D}} d_1(z, \mathcal{Q})$  ensures the coreset  $\mathcal{Q}$  closely reflects the entire dataset distribution by reducing the discrepancy between any sample from  $\mathcal{D}$  and the coreset  $\mathcal{Q}$ .

#### Sample Difficulty

It is evident that the prediction difficulty for the model varies across different samples, meaning not all samples contribute equally to improving the model’s instruction-following capability. Following the process of minimizing cross-entropy loss for each token during language model training, the sample-wise cross-entropy loss  $\mathcal{L}^{(i)}$  directly quantifies the model’s fit and prediction difficulty for any given sample  $z^{(i)}$ :

$$\mathcal{L}_t^{(i)} = -\log(p(y_t^{(i)} | x^{(i)}, y_{<t}^{(i)})), \quad (3)$$

$$\mathcal{L}^{(i)} = \frac{1}{|y^{(i)}|} \sum_t \mathcal{L}_t^{(i)}, \quad (4)$$

where  $y_{<t}^{(i)}$  denotes the token sequence preceding the  $t$ -th token,  $\mathcal{L}_t^{(i)}$  represents the loss of the  $t$ -th token given instruction  $x^{(i)}$  and  $y_{<t}^{(i)}$ , and  $|y^{(i)}|$  is the total number of tokens in  $y^{(i)}$ .

However, improving an LLM’s capability to follow instructions does not precisely align with minimizing loss on a limited amount of labeled instruction tuning data. As illustrated in Figure 2, due to the auto-regressive mechanism in current language models, inherent generation diversity caused by the preceding context commonly occurs. We term this phenomenon context-oriented generation diversity (CoGD). A more intuitive understanding is that, for open-ended text generation, a single instruction can correspond to many valid responses. During pre-training, the model learns the distribution of generation diversity from vast data, which provides a broad coverage of this diversity. In contrast, during instruction tuning, where the amount of training data is considerably reduced, overlooking CoGD in data selection risks overfitting. Specifically, selecting samples with high loss caused by CoGD may lead the LLM to over-fit particular samples or tokens from distributions with greater generation diversity, thereby undermining its instruction-following capability.

To mitigate the interference of CoGD on the difficulty measurement, we introduce the uncertainty-based prediction difficulty (UPD) to measure the model’s prediction difficulty

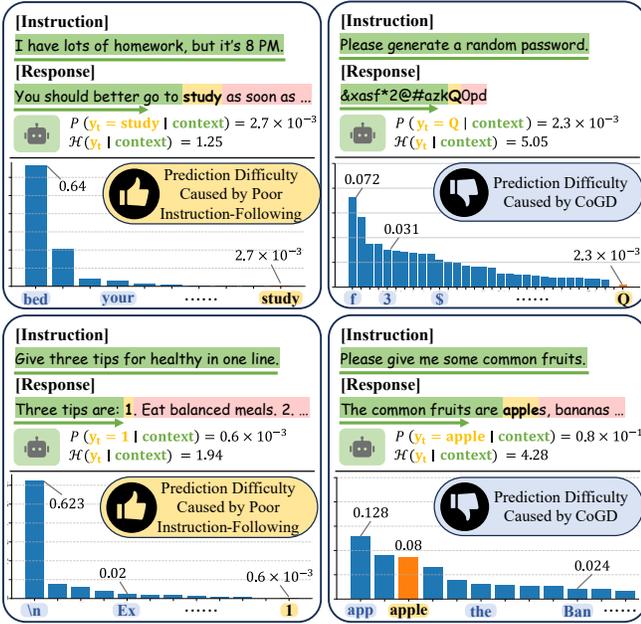


Figure 2: Analysis of top token probabilities reveals two causes for tokens that are difficult to predict: context-oriented diversity and poor instruction-following. We should prioritize enhancing poor instruction-following rather than fitting difficult tokens caused by CoGD, as this risks reducing generation diversity and overfitting.

for samples by accounting for generation diversity. Since it is impractical to consider all potential responses to a given instruction at the sentence level, we quantify this by calculating the Shannon entropy [Shannon, 1948] of token-wise predictions output from the LLM. This uncertainty reflects the generation diversity based on the current context and the pre-trained knowledge in LLM:

$$\mathcal{H}_t^{(i)} = - \sum_j p(w_j | x^{(i)}, y_{<t}^{(i)}) \log(p(w_j | x^{(i)}, y_{<t}^{(i)})), \quad (5)$$

where  $\mathcal{H}_t^{(i)}$  is the entropy of the  $t$ -th token prediction given  $x^{(i)}$  and  $y_{<t}^{(i)}$ , and  $w_j$  denotes each token in the vocabulary.

Despite the high loss, tokens with higher uncertainty indicate greater context-oriented generation diversity. In data selection, we reduce the attention given to such tokens to mitigate the risk of over-fitting during instruction tuning. The UPD score for tokens and samples is computed as follows:

$$\text{UPD}_t^{(i)} = \sigma(\mathcal{L}_t^{(i)}) \cdot \max\left(1 - \frac{\mathcal{H}_t^{(i)}}{[\log(v_{\text{size}})]^\beta}, 0\right), \quad (6)$$

$$d_2(z^{(i)}) = \text{UPD}(z^{(i)}) = \frac{1}{|y^{(i)}|} \sum_t \text{UPD}_t^{(i)}, \quad (7)$$

where  $\sigma(u) = 2\left(\frac{1}{1+e^{-u/\alpha}} - \frac{1}{2}\right) \in [0, 1]$  represents the sigmoid transformation mapping the loss score to  $[0, 1]$ , and  $v_{\text{size}}$  is the size of the token vocabulary. Samples with a smaller UPD score typically exhibit either lower loss, indicating easier prediction, or higher uncertainty, reflecting greater generation diversity. These samples should be assigned lower weights to be considered and selected in the coreset.

### Algorithm 1 The $D_3$ Data Selection and Fine-tuning.

- 1: **Input:** Instruction tuning data pool  $\mathcal{D}$ , Pre-trained LLM  $F$ , Teacher LLM  $F^*$ , Selection budget  $k$ , Warm-up budget  $k_0$ , Total rounds  $R$ , Fine-tuning Algorithm  $\mathcal{A}_{sft}$
- 2: **Procedure:**
- 3: Warm-up set  $\mathcal{Q}_0 \leftarrow \text{sample}(\mathcal{D}, k_0)$   $\triangleright$  Warm-up step
- 4: Warm-up LLM  $F_0 \leftarrow \mathcal{A}_{sft}(F, \mathcal{Q}_0)$
- 5: **for** round  $r$  in  $\{1 \cdots R\}$  **do**  $\triangleright D_3$  scoring step
- 6:   **for**  $z^{(i)} \in \mathcal{D}$  **do**
- 7:     Calc. the sample embedding  $\psi_{F_{r-1}}(z^{(i)})$  for the diversity function  $d_1(z^{(i)}, \cdot)$  in Eq.(1).
- 8:     Calc. the difficulty score  $d_2(z^{(i)})$  by Eq.(6), (7).
- 9:     Calc. the dependability score  $d_3(z^{(i)})$  by Eq.(8).
- 10:   **end for**
- 11:   Initialize  $\mathcal{Q}_r \leftarrow \emptyset$   $\triangleright D_3$  selection step
- 12:   **while**  $|\mathcal{Q}_r|/|\mathcal{D}| < k/R$  **do**
- 13:      $\mathcal{Q} = \mathcal{Q}_1 \cup \cdots \cup \mathcal{Q}_{r-1} \cup \mathcal{Q}_r$
- 14:      $z^* \leftarrow \arg \max_z \{d_1(z, \mathcal{Q}) \cdot d_2(z) \cdot d_3(z)\}$
- 15:      $\mathcal{Q}_r \leftarrow \mathcal{Q}_r \cup \{z^*\}$
- 16:   **end while**
- 17:    $F_r \leftarrow \mathcal{A}_{sft}(F_{r-1}, \mathcal{Q}_r)$   $\triangleright$  Instruction tuning
- 18: **end for**

It is worth noting that some uncertainty-based machine learning methods, such as in active learning [Shao *et al.*, 2022a; Shao *et al.*, 2022b], often prioritize samples with higher uncertainty, as these samples are less familiar to the model and more informative, which contrasts sharply with the motivation behind UPD. Unlike training small models from scratch, LLMs are typically pre-trained on vast corpora, enabling them to acquire extensive knowledge and text generation capabilities [Zhao *et al.*, 2024]. Instruction tuning, which involves relatively much less training data, primarily serves to equip the models with instruction-following capabilities rather than imparting new, low-level knowledge. Therefore, the priorities during instruction tuning should be correcting the model's confidently incorrect prediction patterns rather than reducing the context-oriented generation diversity. From this perspective, we observe that samples with higher UPD scores are primarily due to the model's weak instruction-following rather than factors related to CoGD. This suggests that the model is poorly aligned with these instruction samples, and such samples should be prioritized for instruction tuning to enhance the model's performance.

### Sample Dependability

The difficulty score based on UPD helps identify samples with greater difficulty, which is less likely to be caused by CoGD. However, if a sample itself contains factual errors, extraneous symbols, or lacks fluency, it can undermine both the data selection and instruction tuning. To tackle this, we integrate an external teacher model to assess the dependability of samples through prompt-based evaluation. This helps filter out unreliable samples, reducing the risk of incorporating erroneous supervision during instruction tuning. By leveraging the teacher's instruction-following capability, we obtain a soft dependability score  $d_3$  within  $(0, 1)$  for each sample:

$$d_3(z^{(i)}) = \frac{\exp(F^*(w_{\text{pos}} | T(z^{(i)})))}{\exp(F^*(w_{\text{pos}} | T(z^{(i)}))) + \exp(F^*(w_{\text{neg}} | T(z^{(i)})))}, \quad (8)$$

		Winning Score ( <i>budget k vs. Full</i> )						Leaderboard	
Dataset ( <i>budget k</i> )	Selection Strategy	WizardLM (218)	Sinstruct (252)	Vicuna (80)	Koala (180)	LIMA (300)	Average (1030)	AlpacaEval ( <i>budget k vs. Full</i> )	
Alpaca (5%)	PPL	0.20	0.30	0.08	0.16	0.15	0.19	4.34	28.00
	RAND	0.99	0.93	1.09	0.98	0.95	0.97	28.29	
	IFD	1.15	1.01	<b>1.38</b>	1.13	1.28	1.17	36.50	
	D <sub>3</sub>	<b>1.22</b>	<b>1.16</b>	1.21	<b>1.33</b>	<b>1.37</b>	<b>1.27</b>	<b>42.01</b>	
AlpacaGPT4 (5%)	PPL	0.29	0.33	0.11	0.36	0.20	0.27	20.92	65.68
	RAND	0.85	0.86	0.83	0.92	0.98	0.90	60.93	
	IFD	1.10	1.08	1.00	1.11	1.15	1.10	64.53	
	D <sub>3</sub>	<b>1.13</b>	<b>1.19</b>	<b>1.07</b>	<b>1.24</b>	<b>1.22</b>	<b>1.19</b>	<b>71.26</b>	
WizardLM (5%)	PPL	0.40	0.52	0.60	0.45	0.48	0.48	28.49	58.42
	RAND	0.92	0.81	1.08	0.79	0.94	0.89	54.48	
	IFD	0.93	0.93	<b>1.28</b>	0.99	1.11	1.02	54.94	
	D <sub>3</sub>	<b>0.99</b>	<b>1.14</b>	1.21	<b>1.07</b>	<b>1.17</b>	<b>1.11</b>	<b>63.09</b>	

Table 1: Performance comparisons of various data selection strategies across different datasets. The numbers in parentheses denote the sample size of the dataset. Higher winning scores and leaderboard metrics indicate better performance, with the best results highlighted in **bold**.

where  $F^*$  represents the teacher model,  $w_{\text{pos}}$  and  $w_{\text{neg}}$  correspond to the positive and negative tokens, respectively.  $T(z^{(i)})$  is the evaluation prompt generated based on sample  $z^{(i)}$ , and  $F^*(w|T)$  refers to the logit of the next token  $w$  given the inputs  $T(z^{(i)})$ . Details are provided in the appendix.

### 3.4 D<sub>3</sub> Selection Objective

Finally, under a specific selection budget  $k$ , we select a data coreset that captures the diversity of the overall distribution while ensuring its samples exhibit great difficulty and dependability. We formalize this multi-criteria sample selection problem as the D<sub>3</sub> weighted coreset objective, aiming to solve for the coreset  $\mathcal{Q}$  which jointly optimizes the three aspects of selection criteria under the budget constraint:

$$\begin{aligned} \mathcal{Q} &= \arg \min_{\mathcal{Q}} \max_{z \in \mathcal{D}} d_1(z, \mathcal{Q}) d_2(z) d_3(z) \quad s.t. \frac{|\mathcal{Q}|}{|\mathcal{D}|} = k \\ &= \arg \min_{\mathcal{Q}} \max_{z \in \mathcal{D}} \min_{z' \in \mathcal{Q}} d_2(z) d_3(z) \mathbf{d}_{\psi_F}(z, z'), \end{aligned} \quad (9)$$

where  $\mathbf{d}_{\psi_F}$  represents the cosine distance between the embeddings of samples  $\psi_F(\cdot)$  produced by the LLM  $F$ . We tackle this NP-hard combinatorial optimization problem with the greedy iterative approach. Specifically, we initialize  $\mathcal{Q}$  as an empty set and randomly select a sample to add to  $\mathcal{Q}$ . In each iteration, we compute the weighted nearest distance between the remaining samples and  $\mathcal{Q}$ , selecting the sample with the largest distance to add to  $\mathcal{Q}$ . We repeat this process until the desired number of samples is obtained.

We conduct multiple rounds of scoring and selection steps. In each round  $r$ , we select the coreset  $\mathcal{Q}_r$  based on the scoring results and then fine-tune the model to obtain  $F_r$ . Subsequently, leveraging the sample-level real-time feedback from  $F_r$ , we re-score and re-select the samples for subsequent fine-tuning, which enables the adaptive refinement of the selection focus. The process of D<sub>3</sub> is summarized in Algorithm 1.

		Winning Score ( <i>budget 10% vs. Full</i> )				
Strategy	Satis.	Pres.	Conv.	Spoke.	Overall	
PPL	0.87	0.84	0.81	0.89	0.82	
RAND	0.90	0.91	0.92	0.89	0.93	
IFD	0.77	0.74	0.76	0.82	0.71	
D <sub>3</sub>	<b>1.07</b>	<b>1.07</b>	<b>1.04</b>	<b>1.06</b>	<b>1.11</b>	

Table 2: Pairwise comparisons on real-world Taobao Live dataset across five evaluation dimensions of interest: Satisfaction, Presentation Skill, Conversion, Spoken Fluency, and Overall Performance.

## 4 Empirical Study

### 4.1 Experimental Setup

**Training Datasets.** We employ three commonly used instruction tuning datasets in the following experiments: 1) The Alpaca dataset [Taori *et al.*, 2023], generated through the self-instruct method [Wang *et al.*, 2023] and Davinci003 model, consists of 52,002 samples. Given that the data generation relies on Davinci003, the overall quality of the dataset is moderate, akin to datasets collected in real-world tasks that may include low-quality or erroneous samples. 2) The AlpacaGPT4 dataset [Peng *et al.*, 2023] is a refined version of the Alpaca dataset generated using GPT-4, resulting in significantly higher quality. 3) The WizardLM dataset [Xu *et al.*, 2023], created by the Evol-Instruct method with ChatGPT, consists of 70,000 high-quality instruction samples in total.

**Test Datasets.** We evaluate the effectiveness of different data selection strategies using five commonly used public test datasets: WizardLM [Xu *et al.*, 2023], Self-instruct [Wang *et al.*, 2023], Vicuna [Chiang *et al.*, 2023], Koala [Vu *et al.*, 2023], and LIMA [Zhou *et al.*, 2023]. These test datasets consist of approximately 1,000 instruction-based queries in total, curated across diverse tasks and domains. They provide a comprehensive evaluation of the models’ capabilities to follow instructions in real-world tasks.

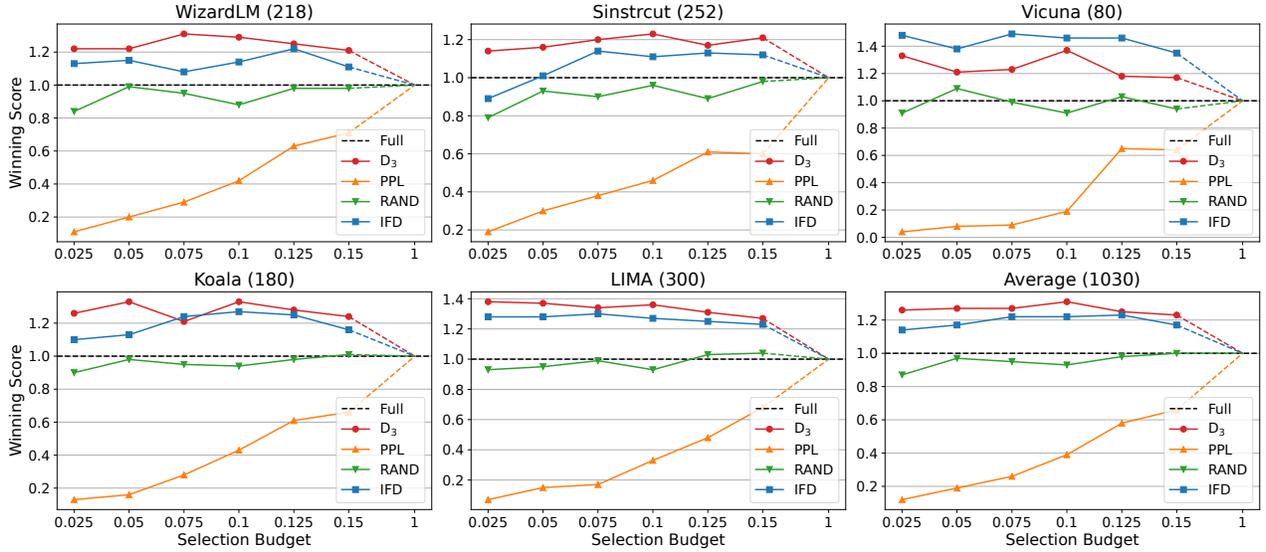


Figure 3: Performance variation across different selection budgets on Alpaca is presented. Detailed results can be found in the appendix.

**Real-world Dataset.** We further evaluate our method in a real-world setting within Taobao Live, where the language model is fine-tuned using collected live-streaming data to generate high-quality sales-pitch scripts tailored to product information, thereby empowering streamers to improve their performance. We utilize 64,881 collected samples for training, with a random selection of 300 reserved for testing. A pairwise evaluation of generated scripts is conducted across five evaluation dimensions of interest in the task context.

**Candidate Methods.** We compare our proposed  $D_3$  method with the following data selection strategies:

- **Full:** No data selection strategy. The entire instruction tuning dataset is utilized to fine-tune the LLMs.
- **RAND:** A subset of training data is randomly sampled from the original instruction tuning dataset.
- **PPL:** The samples with the highest perplexity are selected, corresponding to those with the highest loss.
- **IFD** [Li *et al.*, 2024b]: The IFD score, calculated as the ratio of loss to conditional loss, identifies samples with weaker instruction-following. The subset of samples with the highest IFD scores is then selected.

To ensure fairness and adherence to the settings and procedures of candidate methods, we use the pre-trained Llama2-7B as the initial model. All methods first undergo two training epochs on a warm-up set, which consists of 1% of the data randomly sampled from the data pool. Subsequently, each strategy selects a data subset of equal size for instruction tuning based on a given selection budget. Samples selected are trained for three epochs to obtain the fine-tuned LLMs.

## 4.2 Evaluation and Benchmark

**Pairwise Evaluation.** Evaluating the instruction-following capabilities of LLMs on open-ended natural language generation tasks remains a significant challenge. Building on

recent advancements in LLM evaluation, such as LLM-as-a-judge [Zheng *et al.*, 2023; Li *et al.*, 2025], a more capable LLM is employed as a judge to assess response quality. We generate two responses for each instruction sample in the test set  $\mathcal{D}_t$ : one from the model fine-tuned on a small subset selected and the other from a model fine-tuned on the full dataset. The judge LLM assigns preference scores, and based on these scores, we classify the outcome as a win, tie, or loss for each sample. Each response pair is evaluated twice in different orders to mitigate positional bias. The final winning score is computed from outcomes of the entire test set. Pairwise evaluations are performed across all five test datasets.

$$\text{winning\_score}(\mathcal{D}_t) = \frac{\text{num}(\text{wins}) - \text{num}(\text{loses})}{\text{num}(\mathcal{D}_t)} + 1 \quad (10)$$

**Benchmark.** We also evaluate models on the widely recognized AlpacaEval Leaderboard [Li *et al.*, 2023], which provides automated assessments using the AlpacaFarm dataset [Dubois *et al.*, 2023]. This evaluation compares the model outputs with those of Davinci003 by ChatGPT. Additional details on the evaluation are provided in the appendix.

## 4.3 Main Results

We present the primary comparison results in Table 1. For all three instruction tuning datasets, we allocate a  $k = 5\%$  selection budget for sample-efficient instruction tuning. To ensure a fair comparison, all strategies select data to reach the budget in one go, except for the Full strategy, which uses all the data. Specifically, we set the total round  $R = 1$  in  $D_3$ . Our proposed  $D_3$  method demonstrates superior results by effectively selecting a small, yet highly valuable subset of data, enhancing the model’s instruction-following capabilities on both medium and high-quality instruction tuning datasets. It outperforms the Full strategy in both performance and efficiency. In contrast, using perplexity, or loss, as a measure of sample value for data selection guidance fails to account for

	Winning Score (5% vs. Full)		
Strategy	Alpaca	AlpacaGPT4	WizardLM
$D_3$ -w/o1	0.38	0.83	0.85
$D_3$ -w/o2	1.08	1.03	1.01
$D_3$ -w/o3	1.12	1.12	1.06
$D_3$	<b>1.27</b>	<b>1.19</b>	<b>1.11</b>

Table 3: Comparisons of  $D_3$  with three ablation strategies. Scores represent the average winning scores across all test sets.

diversity and difficulty, resulting in considerably poorer performance. The RAND strategy, which can be seen as a simple approach that focuses solely on diversity to capture the overall data distribution by random sampling, achieves performance slightly weaker than Full. However, it falls short of outperforming Full due to the lack of a more in-depth sample-level evaluation. The IFD method provides a deeper analysis of sample-wise instruction-following difficulty, helping the model enhance performance with less data. However, as it relies solely on a single metric for data selection, it leaves room for further improvement. Besides, the results reported by AlpacaEval show that it is slightly less competitive than the Full strategy on two high-quality instruction tuning datasets. The performance comparisons on the real-world dataset are presented in Table 2, under a fixed selection budget  $k = 10\%$ , further demonstrating the superiority of our method.

#### 4.4 Performance with Varied Budgets

We systematically examine the impact of selection budgets on various strategies by incrementally increasing the budget. In line with the setup in the main experiment, all strategies select the required data in a single round. Figure 3 presents the winning score on the Alpaca dataset. As the budget increases, the performance of more potent strategies exhibits a trend of initially improving and then declining. Conversely, weaker strategies gradually approach the performance of the Full strategy. This is because selecting more valuable data positively contributes to model instruction tuning. However, once the effective data reaches saturation, incorporating additional data of slightly lower quality or with redundancy leads to diminishing returns, ultimately causing the performance to converge with that of the Full strategy. Specifically,  $D_3$  achieves superior performance with the smallest budget, highlighting the effectiveness of our strategy in selecting the most valuable coreset data for sample-efficient instruction tuning. Similarly, the results on the other two high-quality datasets, as reported in the appendix, also empirically validate this. In contrast, other strategies display weaker performance trajectories or slower growth, which further underscores the superiority and effectiveness of our approach.

#### 4.5 Ablation Study

We design three ablation strategies by individually removing diversity (w/o1), difficulty (w/o2) and dependability (w/o3) to validate the effectiveness of the three data selection criteria in  $D_3$ . Table 3 shows the performance of the three ablation strategies. Given that Alpaca is a medium-quality dataset,

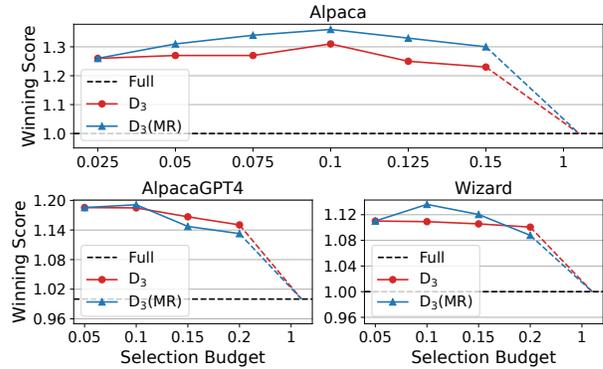


Figure 4: The average performance comparisons of  $D_3$  and  $D_3(MR)$ .

neglecting diversity causes the model to overly focus on certain challenging and highly dependable samples, resulting in a significant performance drop. In the case of the other two high-quality datasets, most samples have dependability scores close to 1, making  $D_3$ -w/o2, to some extent, a strategy that only considers diversity. As a result, its performance is similar to but slightly better than the RAND strategy shown in the main results. The performance impact of removing dependability is more pronounced for Alpaca due to the quality difference of datasets. These validate the necessity of comprehensively considering all three criteria in data selection.

#### 4.6 Selection with Multi-Round

In this section, we present the performance of the  $D_3$  with multi-round selection. The budget increment is 2.5% for the Alpaca dataset and 5% for the other two datasets. Figure 4 compares the performance of multi-round selection  $D_3(MR)$  with single-round selection  $D_3$ . As observed, on the Alpaca dataset with medium quality,  $D_3(MR)$  could adaptively refine the selection focus during the instruction tuning process by leveraging real-time model feedback, thereby better identifying the most valuable subset of data to enhance performance significantly. On datasets with inherently high-quality samples, although the selected small amount of data is already sufficient, multi-round selection still provides some benefits in rapidly improving performance. Note that since  $D_3(MR)$  involves re-scoring in each round, employing multi-round selection necessitates a trade-off between performance gains and the additional computational overhead.

### 5 Conclusion

This paper investigates data selection for sample-efficient LLM instruction tuning. We first establish data selection criteria based on three distinct aspects of diversity, difficulty, and dependability, offering the systematic guidance for data selection. We then propose a novel method  $D_3$  to identify valuable subsets from large datasets by scoring and jointly optimizing three criteria. Our results show that instruction-following capabilities can be achieved with significantly smaller datasets, improving the effectiveness and efficiency of instruction tuning. One direction for future work is to explore token-wise weights on loss based on selection during instruction tuning.

## Acknowledgments

This research was supported by the Key Program of Jiangsu Science Foundation (BK20243012), the Jiangsu Science Foundation (BG2024036), the Fundamental Research Funds for the Central Universities (022114380023), and Alibaba Group through Alibaba Innovative Research Program.

## Contribution Statement

Jia Zhang and Yao Liu contribute equally as co-first authors. Yi Liu and Lan-Zhe Guo serve as co-corresponding authors.

## References

- [Bai *et al.*, 2022] 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. Training a helpful and harmless assistant with reinforcement learning from human feedback. arXiv:2204.05862, 2022.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, and Amanda Askell. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- [Cao *et al.*, 2024] Yihan Cao, Yanbin Kang, Chi Wang, and Lichao Sun. Instruction mining: Instruction data selection for tuning large language models. arXiv:2307.06290, 2024.
- [Chiang *et al.*, 2023] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality, 2023.
- [Du *et al.*, 2023] Qianlong Du, Chengqing Zong, and Jiajun Zhang. MoDS: Model-oriented data selection for instruction tuning. arXiv:2311.15653, 2023.
- [Dubois *et al.*, 2023] Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. AlpacaFarm: A simulation framework for methods that learn from human feedback, 2023.
- [Ge *et al.*, 2024] Yuan Ge, Yilun Liu, Chi Hu, Weibin Meng, Shimin Tao, Xiaofeng Zhao, Mahong Xia, Zhang Li, Boxing Chen, Hao Yang, Bei Li, Tong Xiao, and Jingbo Zhu. Clustering and ranking: Diversity-preserved instruction selection through expert-aligned quality estimation. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 464–478, 2024.
- [He *et al.*, 2024] Yexiao He, Ziyao Wang, Zheyu Shen, Guoheng Sun, Yucong Dai, Yongkai Wu, Hongyi Wang, and Ang Li. Shed: Shapley-based automated dataset refinement for instruction fine-tuning. arXiv:2405.00705, 2024.
- [Honovich *et al.*, 2022] Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. Unnatural instructions: Tuning language models with (almost) no human labor. arXiv:2212.09689, 2022.
- [Khashabi *et al.*, 2020] Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. Unifiedqa: Crossing format boundaries with a single qa system. arXiv:2005.00700, 2020.
- [Kung *et al.*, 2023] Po-Nien Kung, Fan Yin, Di Wu, Kai-Wei Chang, and Nanyun Peng. Active instruction tuning: Improving cross-task generalization by training on prompt sensitive tasks. arXiv:2311.00288, 2023.
- [Li *et al.*, 2023] Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. AlpacaEval: An automatic evaluator of instruction-following models, 2023.
- [Li *et al.*, 2024a] Ming Li, Yong Zhang, Shwai He, Zhitao Li, Hongyu Zhao, Jianzong Wang, Ning Cheng, and Tianyi Zhou. Superfiltering: Weak-to-strong data filtering for fast instruction-tuning. arXiv:2402.00530, 2024.
- [Li *et al.*, 2024b] Ming Li, Yong Zhang, Zhitao Li, Jiu-hai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. From quantity to quality: Boosting llm performance with self-guided data selection for instruction tuning. arXiv:2308.12032, 2024.
- [Li *et al.*, 2025] Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhattacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, Kai Shu, Lu Cheng, and Huan Liu. From generation to judgment: Opportunities and challenges of llm-as-a-judge. arXiv:2411.16594, 2025.
- [Long *et al.*, 2024] Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. On llms-driven synthetic data generation, curation, and evaluation: A survey. arXiv:2406.15126, 2024.
- [Naveed *et al.*, 2024] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language models. arXiv:2307.06435, 2024.
- [Ouyang *et al.*, 2022] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, and Alex Ray. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [Peng *et al.*, 2023] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. arXiv:2304.03277, 2023.

- [Rafailov *et al.*, 2024] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36, 2024.
- [Sener and Savarese, 2018] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. arXiv:1708.00489, 2018.
- [Shannon, 1948] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [Shao *et al.*, 2022a] Jie-Jing Shao, Lan-Zhe Guo, Xiao-Wen Yang, and Yu-Feng Li. Log: Active model adaptation for label-efficient ood generalization. *Advances in Neural Information Processing Systems*, 35:11023–11034, 2022.
- [Shao *et al.*, 2022b] Jie-Jing Shao, Yunlu Xu, Zhanzhan Cheng, and Yu-Feng Li. Active model adaptation under unknown shift. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1558–1566, 2022.
- [Taori *et al.*, 2023] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- [Touvron *et al.*, 2023] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruiti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Es-iobu, 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. Llama 2: Open foundation and fine-tuned chat models. arXiv:2307.09288, 2023.
- [Vu *et al.*, 2023] Thuy-Trang Vu, Xuanli He, Gholamreza Haffari, and Ehsan Shareghi. Koala: An index for quantifying overlaps with pre-training corpora. arXiv:2303.14770, 2023.
- [Wang *et al.*, 2022] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Maitreya Patel, Kuntal Kumar Pal, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Shailaja Keyur Sampat, Savan Doshi, Siddhartha Mishra, Sujay Reddy, Sumanta Patro, Tanay Dixit, Xudong Shen, Chitta Baral, Yejin Choi, Noah A. Smith, Hannaneh Hajishirzi, and Daniel Khashabi. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. arXiv:2204.07705, 2022.
- [Wang *et al.*, 2023] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. arXiv:2212.10560, 2023.
- [Wang *et al.*, 2024a] Ke Wang, Jiahui Zhu, Minjie Ren, Zeming Liu, Shiwei Li, Zongye Zhang, Chenkai Zhang, Xiaoyu Wu, Qiqi Zhan, Qingjie Liu, and Yunhong Wang. A survey on data synthesis and augmentation for large language models. arXiv:2410.12896, 2024.
- [Wang *et al.*, 2024b] Zhichao Wang, Bin Bi, Shiva Kumar Pentylala, Kiran Ramnath, Sougata Chaudhuri, Shubham Mehrotra, Zixu, Zhu, Xiang-Bo Mao, Sitaram Asur, Na, and Cheng. A comprehensive survey of llm alignment techniques: RLHF, RLAIFF, PPO, DPO and more. arXiv:2407.16216, 2024.
- [Xia *et al.*, 2024] Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. Less: Selecting influential data for targeted instruction tuning. arXiv:2402.04333, 2024.
- [Xu *et al.*, 2023] Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. Wizardlm: Empowering large language models to follow complex instructions. arXiv:2304.12244, 2023.
- [Ye *et al.*, 2021] Qinyuan Ye, Bill Yuchen Lin, and Xiang Ren. Crossfit: A few-shot learning challenge for cross-task generalization in nlp. arXiv:2104.08835, 2021.
- [Zhao *et al.*, 2024] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. A survey of large language models. arXiv:2303.18223, 2024.
- [Zheng *et al.*, 2023] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, and Eric Xing. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.
- [Zhou *et al.*, 2023] 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. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36:55006–55021, 2023.