

# THE BEST OF BOTH WORLDS: BRIDGING QUALITY AND DIVERSITY IN DATA SELECTION WITH BIPARTITE GRAPH

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The performance of large language models (LLMs) in natural language processing (NLP) tasks is significantly influenced by the quality and diversity of data used for supervised fine-tuning (SFT). Current data selection methods often focus solely on quality or diversity, leading to underperforming models due to suboptimal training data. In this paper, we introduce GRAPHFILTER, a novel method that represents the dataset as a bipartite graph, linking sentences to their constituent n-grams. This representation effectively captures the relationships between sentences and linguistic patterns, facilitating the selection of sentences that enhance n-gram diversity. To balance quality and diversity during selection, we propose a priority function that combines the quality metric with the diversity metric in a multiplicative manner. GRAPHFILTER iteratively selects high-priority sentences, updates the bipartite graph by removing covered n-grams, and re-calculates priorities to reflect the evolving data landscape. We conduct extensive experiments using three model backbones across six widely used benchmarks. The results demonstrate that GRAPHFILTER outperforms all nine baseline approaches, achieving superior model performance and computational efficiency. Our analyses validate the effectiveness of our design choices, examine the subsets selected by GRAPHFILTER and other methods, highlight the importance of instruction diversity, and explore the role of quality and diversity in relation to subset sizes. GRAPHFILTER establishes a new foundation for effective data selection strategies, encouraging further research in data selection for LLMs.

## 1 INTRODUCTION

Large language models (LLMs) have significantly advanced the field of natural language processing (NLP), enabling models to generate coherent and contextually relevant text across a variety of tasks (Ouyang et al., 2022; Sanh et al., 2022; OpenAI, 2023; Anil et al., 2023b; Touvron et al., 2023a;b; Anil et al., 2023a; Mesnard et al., 2024; Yang et al., 2024). Central to the success of these models is the quality and diversity of the data used during supervised fine-tuning (SFT). Fine-tuning on high-quality data ensures that the model learns accurate language patterns and responds appropriately to inputs (Wang et al., 2023; Zhou et al., 2023), while diversity in the data allows the model to generalize across different contexts and topics (Abbas et al., 2023; Maharana et al., 2024). However, the vastness of available SFT data presents a challenge: selecting a subset of data that balances both quality and diversity to optimize model performance.

Recent methods for data selection often prioritize either quality or diversity, rarely achieving an optimal balance of both. Approaches that focus exclusively on quality may overlook the variety of language patterns necessary for effective generalization (Marion et al., 2023; Ankner et al., 2024). Conversely, methods emphasizing diversity might include lower-quality data, which could negatively impact model performance (Abbas et al., 2023; Lu et al., 2024). This focus can result in models that either overfit to specific data patterns or underperform due to the inclusion of irrelevant or poor-quality data. Hence, it is crucial to develop a data selection strategy that simultaneously maximizes both data quality and diversity for effective supervised fine-tuning.

In response to this challenge, we propose a novel method, GRAPHFILTER, as shown in Figure 1, which models the dataset as a bipartite graph to capture the relationships between sentences and their constituent n-grams. In this model, sentences and n-grams are distinct sets of nodes, with edges indicating n-grams’ presence in sentences, providing a comprehensive overview of n-gram coverage. This structure allows us to prioritize sentences that contribute unique n-grams, enhancing the diversity of the selected subset. To ensure high-quality and diverse sentence selection, we use a priority function that evaluates sentences on these dimensions. For quality, we use the SUPERFILTER, a metric that measures the informativeness of a response by comparing its perplexity when conditioned on the instruction with its standalone perplexity, favoring more relevant responses. For diversity, we calculate Term Frequency-Inverse Document Frequency (TF-IDF) scores for n-grams within sentences. By summing the TF-IDF scores of all n-grams present in a sentence, we estimate the sentence’s contribution to covering significant and less frequent linguistic patterns in the dataset. The priority function combines these two measures multiplicatively, assigning higher priority to sentences that are both informative (high-quality) and contribute substantially to n-gram diversity. During the selection process, GRAPHFILTER iteratively selects sentences with the highest priority scores, updates the bipartite graph by removing the covered n-grams, and re-calculates the priorities based on the modified graph.

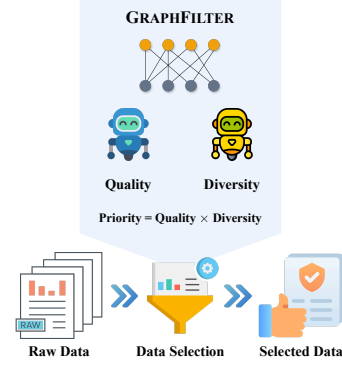


Figure 1: GRAPHFILTER.

To demonstrate the effectiveness of GRAPHFILTER, we conducted extensive experiments, comparing GRAPHFILTER against nine baseline approaches using three model backbones: GEMMA-2-2B (Team et al., 2024), MISTRAL-7B-v0.3 (Jiang et al., 2023), and LLAMA-3-8B (Dubey et al., 2024). The evaluation was performed across six widely-used benchmarks: MMLU (Hendrycks et al., 2021), ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), GSM8K (Cobbe et al., 2021), AlpacaEval-2.0 (Dubois et al., 2024), and MT-Bench (Zheng et al., 2023). Our empirical results indicate that GRAPHFILTER significantly outperforms recent state-of-the-art baselines and achieves notably better computational efficiency. Specifically, in terms of overall performance, GRAPHFILTER outperforms the baselines by up to +2.37 for GEMMA-2-2B, +3.02 for MISTRAL-7B-v0.3, and +3.38 for LLAMA-3-8B, and is up to 61× faster than these baselines without requiring GPUs for computation. Furthermore, we performed an in-depth analysis to validate the effectiveness of our design choices in GRAPHFILTER, examine the characteristics of the selected subsets and the importance of instruction diversity, and investigate the significance of quality and diversity under various data scales.

In summary, the contributions of this work are threefold:

- We introduce a novel bipartite graph model for datasets, named GRAPHFILTER, which effectively captures the complex relationships between sentences and n-grams. This model enables efficient data selection for supervised fine-tuning (see Section 3.2).
- We propose a priority function that seamlessly integrates quality and diversity metrics for re-ranking the data. This ensures that the selected data maximizes n-gram coverage while maintaining high quality (see Section 3.3).
- Through experiments, we demonstrate that our method, GRAPHFILTER, surpasses existing data selection strategies, achieving significantly better computational efficiency (see Section 4). Additionally, our detailed analyses provide valuable insights into the design choices of GRAPHFILTER, the characteristics of the selected subset, and the importance of quality and diversity in relation to the subset sizes (see Section 5).

## 2 RELATED WORK

**Data Engineering for Large Language Models** The success of recent large language models (LLMs) largely relies on the data used during their training process (Zha et al., 2023). State-of-the-art LLMs are generally trained on vast corpora (OpenAI, 2023; Team et al., 2024; Dubey et al., 2024). A significant area of research focuses on curating high-quality corpora for pre-training these

models (Raffel et al., 2020; Computer, 2023; Soldaini et al., 2024; Penedo et al., 2024). Furthermore, Wang et al. (2023) demonstrate that LLMs are capable of synthesizing high-quality datasets for supervised fine-tuning, which leads to a surge of research on dataset synthesis (Xu et al., 2023; Li et al., 2023; Gunasekar et al., 2023; Ding et al., 2023; Cui et al., 2023; Wu et al., 2024; Chen et al., 2024a; Xu et al., 2024). These research efforts facilitate the synthesis of large-scale datasets containing billions of tokens for various purposes, resulting in a significant demand for selecting valuable subsets.

**Data Selection** Data selection strategies aim to identify the most informative data subsets for training or fine-tuning models by considering quality and diversity. Quality-focused approaches prioritize metrics like complexity, difficulty, or informativeness (Marion et al., 2023; Chen et al., 2024b; Liu et al., 2024; Li et al., 2024b;a), but may neglect the range of language patterns needed for generalization. Conversely, diversity-focused methods capture a broad spectrum of linguistic patterns and contexts, potentially incorporating lower-quality data that could impair model performance (Abbas et al., 2023; Lu et al., 2024).

**Ours** To overcome limitations in current data selection methods, we propose GRAPHFILTER, a novel approach that represents the dataset as a bipartite graph of sentences and their n-grams. By balancing quality and diversity with a priority function, our method improves model performance across various downstream tasks.

### 3 METHODOLOGY

In this section, we first introduce the data selection problem for supervised fine-tuning in Section 3.1. Subsequently, we describe the modeling of the dataset as a bipartite graph in Section 3.2. Finally, we explain the re-ranking of the graph nodes using a priority function that integrates quality and diversity metrics during the data selection process in Section 3.3.

#### 3.1 DATA SELECTION PROBLEM

The data selection problem involves the challenge of identifying and selecting the most relevant and informative subset of supervised instances from a larger dataset to fine-tune large language models (LLMs). Formally, let  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$  be the supervised fine-tuning (SFT) dataset, where  $x_i$  represents the instruction and  $y_i$  its corresponding response for the  $i$ -th training instance. Our aim is to select a subset  $\mathcal{S}_\pi$  of size  $k$  from  $\mathcal{D}$ , utilizing the data selection strategy  $\pi$ , where  $k$  is the *data selection budget*. The objective is to determine the optimal data selection strategy  $\pi^*$  that is capable of selecting a subset  $\mathcal{S}_\pi$  maximizing the performance of the fine-tuned LLM  $f_\theta$  on the downstream tasks  $\mathcal{D}_{\text{lst}}$ . Therefore, the data selection problem can be formally formulated as:

$$\pi^* = \arg \max_{\pi} \mathcal{R}(f_\theta; \mathcal{D}_{\text{lst}}), \text{ subject to } |\mathcal{S}_\pi| = k, \text{ where } \theta = \text{FineTune}(\mathcal{F}, \mathcal{S}_\pi), \quad (1)$$

where  $\mathcal{S}_\pi$  is the subset of the training data selected by the strategy  $\pi$ ,  $\theta = \text{FineTune}(\mathcal{F}, \mathcal{S}_\pi)$  denotes the parameters of the model backbone  $\mathcal{F}$  after fine-tuning on the selected data subset  $\mathcal{S}_\pi$ ,  $f_\theta$  is the fine-tuned model with parameters  $\theta$ , and  $\mathcal{R}(f_\theta; \mathcal{D}_{\text{lst}})$  is the performance metric (e.g., accuracy) of the fine-tuned model  $f_\theta$  evaluated on the downstream tasks  $\mathcal{D}_{\text{lst}}$ .

#### 3.2 GRAPHFILTER: MODELING DATASETS AS BIPARTITE GRAPHS

In our approach, we model the dataset as a bipartite graph to effectively represent the relationships between sentences and their constituent n-grams. A bipartite graph is a special type of graph whose vertices can be divided into two disjoint and independent sets such that every edge connects a vertex from one set to a vertex from the other set. Formally, a bipartite graph  $\mathcal{G} = (\mathcal{U}, \mathcal{V}, \mathcal{E})$  consists of *sentence nodes* ( $\mathcal{U} = \{u_i\}_{i=1}^N$ ), *n-gram nodes* ( $\mathcal{V} = \{v_j\}_{j=1}^M$ ), and *edges* ( $\mathcal{E} \subseteq \mathcal{U} \times \mathcal{V}$ ). This structure allows us to capture the occurrence of n-grams within sentences, providing a foundation for selecting sentences that maximize n-gram coverage while adhering to specific priorities. We introduce the details of the priority for re-ranking the sentences in Section 3.3.

**Algorithm 1:** GRAPHFILTER

---

**Input :**  $\mathcal{U} = \{u_i\}_{i=1}^N$ , the set of sentence nodes;  $\mathcal{V} = \{v_j\}_{j=1}^M$ , the set of n-gram nodes;  
 $\mathcal{E} \subseteq \mathcal{U} \times \mathcal{V}$ , the set of edges between sentence nodes and n-gram nodes;  $k$ , the data  
selection budget;  $\phi(u)$ , the priority function for each  $u \in \mathcal{U}$ ;  
**Output:** The selected subset  $\mathcal{S}$ ;

```

1  $\mathcal{S} = \emptyset$ ;
2 while  $|\mathcal{S}| < k \wedge \mathcal{U} \neq \emptyset$  do
3   // Select the sentence with the highest priority
4    $u^* \leftarrow \arg \max_{u \in \mathcal{U}} \phi(u)$ ;
5   // Find n-gram nodes connected to  $u^*$ 
6    $\mathcal{V}_{u^*} \leftarrow \{v \in \mathcal{V} \mid (u^*, v) \in \mathcal{E}\}$ ;
7   // Add  $u^*$  to the selected set
8    $\mathcal{S} \leftarrow \mathcal{S} \cup \{u^*\}$ ;
9   // Remove  $u^*$  from the remaining sentences
10   $\mathcal{U} \leftarrow \mathcal{U} \setminus \{u^*\}$ ;
11  // Remove edges connected to  $u^*$ 
12   $\mathcal{E} \leftarrow \mathcal{E} \setminus \{(u^*, v) \mid v \in \mathcal{V}_{u^*}\}$ ;
13  // Remove  $\mathcal{V}_{u^*}$  and edges connected to it
14  foreach  $v \in \mathcal{V}_{u^*}$  do
15     $\mathcal{E} \leftarrow \mathcal{E} \setminus \{(u, v) \mid u \in \mathcal{U}\}$ ;
16     $\mathcal{V} \leftarrow \mathcal{V} \setminus \{v\}$ ;
17  end
18 end

```

---

Our objective is to select a subset of sentences, denoted as  $\mathcal{S}$ , from the entire dataset, constrained by a data selection budget  $k$ . The aim is to maximize the coverage of unique n-grams while aligning with a priority function  $\phi(u)$  for each sentence  $u \in \mathcal{U}$ . As illustrated in Algorithm 1, our method, referred to as GRAPHFILTER, operates iteratively by updating the graph structure to reflect the n-gram coverage as sentences are selected. The process begins with an empty set of selected sentences,  $\mathcal{S} = \emptyset$ , and a bipartite graph  $\mathcal{G}$  that includes sentence nodes, n-gram nodes, and connecting edges. In each iteration, we select the sentence  $u^* \in \mathcal{U}$  that has the highest priority score  $\phi(u^*)$ , add  $u^*$  to  $\mathcal{S}$ , and then remove  $u^*$  from the set of remaining sentences  $\mathcal{U}$ . Next, we identify the n-grams covered by  $u^*$ , denoted as  $\mathcal{V}_{u^*}$ . We then remove all edges that connect  $u^*$  to the n-gram nodes in  $\mathcal{V}_{u^*}$ . Subsequently,  $\mathcal{V}_{u^*}$  and all edges between its n-grams and other sentences are eliminated from the graph. Note that the priority of each sentence  $u \in \mathcal{U}$  is computed based on the most recent graph  $\mathcal{G}$  during each iteration.

Moreover, we present a minimalist example in Figure 2. Initially, the bipartite graph is displayed in Figure 2a. In Figure 2b, the sentence node  $u_1$  is selected as  $u^*$  and is highlighted in yellow, along with its associated n-gram nodes,  $\mathcal{V}_{u_1}$ , which are highlighted in red. Figure 2c demonstrates the removal of edges connected to  $u_1$  and  $\mathcal{V}_{u_1}$ , as indicated by dashed lines. Finally, Figure 2d illustrates the removal of isolated nodes, shown in white. The next selected sentence node is  $u_4$ . In this example, GRAPHFILTER can cover all the n-grams by selecting only  $u_1$  and  $u_4$ .

Our problem formulation is related to the classical *set cover NP-hard problem* (Garey & Johnson, 1979). In the set cover problem, given a universe of elements and a collection of sets whose union comprises the universe, the objective is to identify the smallest number of sets whose union still contains all elements in the universe. Similarly, in a special case of our problem where the priority function assigns the same score to all sentences (i.e.,  $\phi(u) = 1$  for all  $u \in \mathcal{U}$ ), and the goal is to find the minimal set of sentences that cover *all* n-grams, our task becomes analogous to the set cover problem. In this scenario, the *greedy approach* used in Algorithm 1 can be shown to have an approximation factor of  $H(r)$  (Vazirani, 2001), where  $r$  is the maximum degree of the sentence nodes in the graph (the largest number of n-grams contained in any sentence), and  $H(r) = \sum_{k=1}^r \frac{1}{k}$  is the  $r$ -th harmonic number. This relationship highlights the theoretical foundations of our method and provides insight into its performance guarantees in this special case.

By modeling the dataset as a bipartite graph and employing an iterative selection algorithm, GRAPHFILTER effectively selects a subset of sentences that maximizes n-gram coverage while adhering to

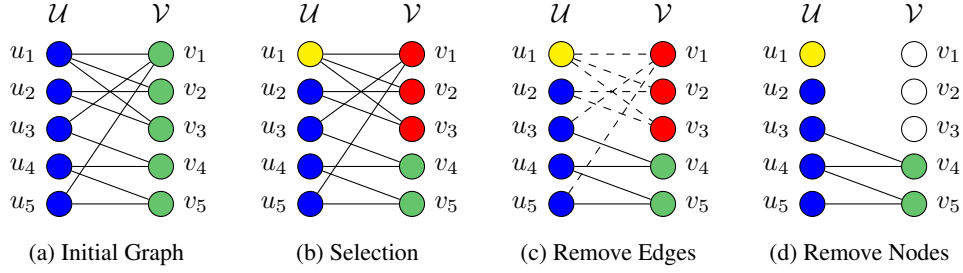


Figure 2: An example of a single iteration of GRAPHFILTER without the priority function. In this case, the degree of a sentence node serves as the priority score. Sentence nodes are in blue and n-gram nodes in green. The selected sentence node is yellow, while connected n-gram nodes are red. Removed n-gram nodes are white, with removed edges as dashed lines. Node  $u_1$  is selected in the current iteration, and  $u_4$  will be the next.

specified priorities. Each SFT training instance comprises instructions and responses. In this work, we apply GRAPHFILTER solely to the instructions of the SFT data.

**Implementation** In a brute-force implementation, the computational complexity of our algorithm is  $\mathcal{O}(N)$  per iteration. This complexity results from the need to perform operations such as selecting the highest-priority sentence and removing edges, which involve scanning the sets of sentences ( $\mathcal{U}$ ), n-grams ( $\mathcal{V}$ ), and edges ( $\mathcal{E}$ ). These sets are not optimized for efficient access or modification. To enhance computational efficiency, we employ a max-heap (or priority queue) to select the highest-priority sentence, allowing this selection to be performed in  $\mathcal{O}(\log N)$  time per iteration. This reduces the selection complexity from  $\mathcal{O}(N)$  to  $\mathcal{O}(\log N)$ . Additionally, the max-heap data structure facilitates the localization of priority updates to affected nodes, eliminating the need to enumerate all nodes and edges.

### 3.3 BALANCING QUALITY AND DIVERSITY WITH PRIORITY FUNCTION

As illustrated in Algorithm 1, GRAPHFILTER naturally selects a subset with maximal n-gram coverage, emphasizing data diversity. However, the quality of the data is equally important for effective language model training. To balance both quality and diversity in our selection process, we define a priority function  $\phi(u)$  for each sentence node  $u \in \mathcal{U}$ , which is used to re-rank the sentence nodes during selection.

**SUPERFILTER for Quality** For quality, we employ the SUPERFILTER as the quality measure (Li et al., 2024a;b). The SUPERFILTER metric evaluates the informativeness of a response by comparing the perplexity of the response conditioned on the instruction with the perplexity of the response alone. Formally, for a given sentence node  $u$  associated with the instruction-response pair  $(x, y)$ , the quality priority metric is defined as:

$$\text{QUALITY}(u) = \text{SUPERFILTER}(x, y) = \frac{\text{PPL}(y | x)}{\text{PPL}(y)}, \quad (2)$$

$$\text{where } \text{PPL}(\mathbf{w}) = \exp \left( -\frac{1}{T} \sum_{t=1}^T \log P(w_t | \mathbf{w}_{<t}) \right),$$

where  $\text{PPL}(\mathbf{w})$  is the perplexity of the sentence  $\mathbf{w}$  with a length of  $T$ ,  $\text{PPL}(y)$  is the perplexity of the response  $y$ , and  $\text{PPL}(y | x)$  is the perplexity of the response  $y$  conditioned on the instruction  $x$ . A higher SUPERFILTER value indicates that the response is more relevant and informative given the instruction, thus reflecting higher quality. It is important to note that quality metrics, such as SUPERFILTER, can be pre-computed prior to the selection process.

**TF-IDF for Diversity** For diversity, we use the Term Frequency-Inverse Document Frequency (TF-IDF) as a measure of the significance of each n-gram within the dataset. The TF-IDF score of an n-gram  $v$  is calculated as  $\text{TF-IDF}(v) = \text{TF}(v) \times \text{IDF}(v)$ , where  $\text{TF}(v)$  (Term Frequency) is

the number of times n-gram  $v$  appears in the corpus, and  $\text{IDF}(v)$  (Inverse Document Frequency) is defined as  $\text{IDF}(v) = \log\left(\frac{N}{d_v}\right)$ , with  $N$  being the total number of sentences in the corpus, and  $d_v$  being the number of sentences containing n-gram  $v$ . Furthermore, we compute the sum of TF-IDF scores of all n-grams (of varying lengths) present in the sentence:

$$\text{DIVERSITY}(u) = \sum_{v \in \mathcal{V}_u} \text{TF-IDF}(v), \quad (3)$$

where  $\mathcal{V}_u$  is the set of n-grams connected to sentence  $u$  in the graph  $\mathcal{G}$ . In our work,  $\mathcal{V}_u$  includes unigrams ( $n = 1$ ), bigrams ( $n = 2$ ), and trigrams ( $n = 3$ ) present in sentence  $u$ , capturing both word-level and phrase-level features.

**Combined Priority Function** To effectively prioritize sentences based on both quality and diversity, we combine the QUALITY score and the DIVERSITY score for the sentence node  $u$  into a single priority function:

$$\phi(u) = \text{QUALITY}(u) \times \text{DIVERSITY}(u). \quad (4)$$

This function assigns higher priority to sentences that are both high-quality and contribute significantly to n-gram diversity. By integrating both quality and diversity into the priority function, our selection algorithm can effectively choose a subset that not only covers a wide range of linguistic patterns but also maintains a high standard of data quality.

## 4 EXPERIMENTS

In this section, we initially outline our experimental setup in Section 4.1, followed by a presentation of our main results in Section 4.2.

### 4.1 EXPERIMENTAL SETUP

**Training Dataset** Xu et al. (2024) utilize state-of-the-art open-source large language models (LLMs) to create a high-quality dataset collection known as *Magpie*. In our research, we employ the *Magpie* dataset, which is generated by LLAMA-3-70B-INSTRUCT and comprises 300K training instances.<sup>1</sup> For this study, we choose a subset of 10K training instances using various selection methods from the entire dataset, unless otherwise stated.

**Baselines** We compare our approach, GRAPHFILTER, with a diverse array of baseline methods:

- **Heuristic:** (1) RANDOM randomly selects a subset from the entire dataset; (2) LONGEST sorts the training instances in descending order based on the length of the instructions;
- **Quality-based:** (3) PERPLEXITY utilizes perplexity values, where larger values typically indicate higher difficulty and quality of training instances; (4) ARMORM is the state-of-the-art open-sourced reward model presented by Wang et al. (2024);<sup>2</sup> (5) ALPAGASUS demonstrates that state-of-the-art LLMs can be directly prompted for estimating data quality (Chen et al., 2024b); (6) DEITA leverages CHATGPT to synthesize a quality estimation dataset and fine-tune LLMs for data quality estimation (Liu et al., 2024); (7) SUPERFILTER indicates the Instruction-Following Difficulty (IFD) metric computed by smaller language models. Li et al. (2024b) introduce this method, while Li et al. (2024a) demonstrate that IFD scores from smaller models are as accurate as those from larger models;
- **Diversity-based:** (8) KMEANS clusters the training instances with the state-of-the-art sentence embedding model and selects the training instances that are closest to their respective cluster centroids (Arthur & Vassilvitskii, 2007); (9) INSTAG is designed for analyzing the SFT dataset by tagging the topics of training instances, and can be used for selecting the subset with the most diverse topics from the entire dataset (Lu et al., 2024).

We present more details of these baseline approaches in Section A.1. To demonstrate the effectiveness and generality of GRAPHFILTER, we conduct experiments on three diverse model back-

<sup>1</sup><https://huggingface.co/datasets/Magpie-Align/Magpie-Pro-300K-Filtered>

<sup>2</sup><https://huggingface.co/RLHFlow/ArmORM-Llama3-8B-v0.1>

Table 1: Main results given by GEMMA-2-2B, MISTRAL-7B-v0.3, and LLAMA-3-8B on the standardized benchmarks and LLM-as-a-Judge benchmarks. The datasets HS, G8K, and AE-2 correspond to HellaSwag, GSM8K, and AlpacaEval-2.0, respectively. The best results are highlighted in **bold**, and the second-best results are highlighted in underline.

	Standardized					LLM-as-a-Judge					$\mu_{ALL}$	
	MMLU	ARC	HS	G8K	$\mu_{BENCH}$	AE-2		MT-Bench		$\mu_{LLM}$		
	Acc	Acc	Acc	Acc		LC	WR	$\mu_{MT}$	1st			2nd
GEMMA-2-2B												
RANDOM	25.25	47.52	58.27	9.10	35.03	10.77	13.73	4.73	5.44	4.01	29.01	33.03
LONGEST	25.50	47.06	56.43	8.79	34.45	10.40	13.10	4.79	5.54	4.05	29.17	32.69
PERPLEXITY	23.34	47.58	59.04	6.48	34.11	12.19	14.76	4.98	5.75	4.21	31.00	33.07
ARMORM	25.42	<b>48.06</b>	56.19	<u>10.62</u>	35.07	<b>13.40</b>	<b>16.39</b>	4.84	5.55	4.14	30.92	33.69
ALPAGASUS	26.56	47.18	58.69	<u>10.57</u>	35.75	13.12	15.76	4.89	5.68	4.11	<u>31.03</u>	34.18
DEITA	28.72	47.51	58.35	10.16	36.18	12.99	15.86	4.82	5.62	4.01	<u>30.57</u>	34.31
SUPERFILTER	<u>28.82</u>	47.20	59.18	9.33	36.13	12.87	15.55	4.88	5.49	<u>4.26</u>	30.81	<u>34.36</u>
KMEANS	28.39	46.96	56.59	10.31	35.56	12.19	14.76	4.98	5.74	4.23	31.00	34.04
INSTAG	27.60	47.75	<b>59.98</b>	9.86	<u>36.29</u>	12.75	15.47	4.79	5.45	4.13	30.31	34.30
GRAPHFILTER	<b>29.06</b>	<u>47.92</u>	<u>59.38</u>	<b>10.71</b>	<b>36.77</b>	<u>13.14</u>	<u>15.99</u>	<b>5.01</b>	<b>5.77</b>	<b>4.25</b>	<b>31.64</b>	<b>35.06</b>
MISTRAL-7B-v0.3												
RANDOM	25.50	52.17	67.44	9.17	38.57	14.76	17.41	5.03	5.93	4.13	32.51	36.55
LONGEST	25.17	52.11	67.32	10.30	38.73	13.67	16.14	4.96	6.00	3.91	31.62	36.36
PERPLEXITY	30.64	52.42	<u>69.31</u>	4.62	39.25	13.60	16.18	4.98	6.01	3.95	31.70	36.73
ARMORM	28.84	50.85	68.85	9.63	39.54	<b>15.56</b>	<b>18.89</b>	5.13	5.93	4.34	<u>33.43</u>	37.51
ALPAGASUS	28.67	51.92	68.61	9.48	39.67	14.67	18.14	5.21	<u>6.13</u>	4.30	33.40	37.58
DEITA	29.86	50.82	67.99	10.60	39.82	14.08	16.49	5.03	5.93	4.13	32.18	37.27
SUPERFILTER	<b>33.59</b>	<u>52.45</u>	68.56	9.93	<u>41.13</u>	13.59	16.75	<u>5.23</u>	6.01	<u>4.44</u>	32.92	<u>38.40</u>
KMEANS	28.77	50.58	67.81	11.55	39.68	13.98	16.83	5.11	5.93	4.29	32.52	37.29
INSTAG	28.29	50.99	67.44	<b>12.59</b>	39.82	14.55	17.36	5.11	5.86	4.36	32.84	37.50
GRAPHFILTER	<u>33.24</u>	<b>52.48</b>	<b>69.69</b>	<u>11.92</u>	<b>41.83</b>	<u>15.16</u>	<u>18.85</u>	<b>5.38</b>	<b>6.23</b>	<b>4.54</b>	<b>34.49</b>	<b>39.38</b>
LLAMA-3-8B												
RANDOM	49.55	52.00	67.30	22.14	47.75	22.17	25.05	5.99	6.95	5.03	41.04	45.51
LONGEST	44.52	50.56	67.99	24.56	46.91	20.17	22.67	5.97	6.82	5.13	39.96	44.59
PERPLEXITY	51.08	52.31	<b>68.74</b>	20.96	48.27	20.38	22.87	6.02	7.02	5.01	40.28	45.61
ARMORM	47.84	52.24	68.11	24.64	48.21	<b>23.45</b>	<u>26.60</u>	<u>6.19</u>	7.14	5.24	<u>42.66</u>	46.36
ALPAGASUS	49.90	51.63	68.40	25.89	48.96	22.90	25.94	6.09	7.05	5.13	41.90	46.60
DEITA	48.49	52.40	<u>68.46</u>	25.78	48.78	22.23	24.42	6.12	7.12	5.11	41.70	46.42
SUPERFILTER	50.16	51.10	67.70	<u>27.45</u>	49.10	22.54	24.68	6.13	<b>7.23</b>	5.03	41.91	46.70
KMEANS	51.98	51.35	67.15	25.12	48.90	22.06	24.80	6.14	7.03	<u>5.25</u>	41.72	46.51
INSTAG	<u>53.16</u>	<u>52.85</u>	67.86	25.85	<u>49.93</u>	22.10	24.64	6.13	7.05	5.21	41.72	<u>47.19</u>
GRAPHFILTER	<b>53.73</b>	<b>52.92</b>	67.76	<b>27.81</b>	<b>50.55</b>	<u>22.95</u>	<b>26.71</b>	<b>6.26</b>	<u>7.21</u>	<b>5.31</b>	<b>42.79</b>	<b>47.97</b>

bones, including GEMMA-2-2B (Team et al., 2024),<sup>3</sup> MISTRAL-7B-v0.3 (Jiang et al., 2023),<sup>4</sup> and LLAMA-3-8B (Dubey et al., 2024).<sup>5</sup> The optimization details are in Section A.2.

**Evaluation** We conduct evaluations on six popular benchmarks, categorized into two groups:

- **Standardized:** We assess the LLMs using LM-EVALUATION-HARNESS (Gao et al., 2024) on four standardized benchmarks: MMLU (Hendrycks et al., 2021), ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), and GSM8K (Cobbe et al., 2021). The model performance on these benchmarks is measured by accuracy. We use the macro-average accuracy across four benchmarks as the overall performance of this group, denoted as  $\mu_{BENCH}$ .
- **LLM-as-a-Judge:** We evaluate LLMs using two benchmarks: AlpacaEval-2.0 (Dubois et al., 2024) and MT-Bench (Zheng et al., 2023), with GPT-4O-2024-05-13 as

<sup>3</sup><https://huggingface.co/google/gemma-2-2b>

<sup>4</sup><https://huggingface.co/mistralai/Mistral-7B-v0.3>

<sup>5</sup><https://huggingface.co/meta-llama/Meta-Llama-3-8B>

the judge. For AlpacaEval-2.0, GPT-4-1106-PREVIEW generates reference answers, and we report both the length-controlled win rate (LC) and the original win rate (WR). For MT-Bench, performance is denoted as  $\mu_{\text{MT}}$ , the macro-average across all categories. Overall performance of this group,  $\mu_{\text{LLM}}$ , is the macro-average of LC and  $\mu_{\text{MT}}$ .

We define overall model performance,  $\mu_{\text{ALL}}$ , as the macro-average of results from four standardized benchmarks, LC, and  $\mu_{\text{MT}}$ . In calculating  $\mu_{\text{ALL}}$  and  $\mu_{\text{LLM}}$ ,  $\mu_{\text{MT}}$  is scaled by  $10\times$  to align with a range of 1 to 100, matching other benchmarks. Further evaluation details are in Section A.3.

## 4.2 MAIN RESULTS

**GRAPHFILTER surpasses all baseline approaches.** As shown in Table 1, GRAPHFILTER consistently outperforms all baseline approaches across the three model backbones on both standardized benchmarks and LLM-as-a-Judge benchmarks. It achieves either the best or second-best results on most individual benchmarks. Specifically, in terms of  $\mu_{\text{ALL}}$ , GRAPHFILTER outperforms the baselines by up to +2.37 for GEMMA-2-2B, +3.02 for MISTRAL-7B-V0.3, and +3.38 for LLAMA-3-8B, compared to LONGEST. These results demonstrate the superiority of GRAPHFILTER which effectively combines the quality and diversity in data selection.

**Quality-based data selection approaches appear to exhibit biases towards specific benchmarks.** Quality-based approaches often use neural models to estimate the quality of each training instance. However, these models display biases that can significantly affect downstream performance. As demonstrated in Table 1, models fine-tuned on subsets chosen by ARMORM perform well on AlpacaEval-2.0 but poorly on other benchmarks. Furthermore, the PERPLEXITY-selected subset consistently results in the worst performance on GSM8K, highlighting the risks of depending solely on neural models for selecting high-quality data.

**GRAPHFILTER is highly efficient, with its variant running quickly on a CPU.** Recent baselines typically rely on neural models for quality estimation, which generally require a GPU. We compare the runtimes of various baselines on a system equipped with an A100 80G GPU and 20 CPU cores, as shown in Table 2. As elaborated in Section 3.3, GRAPHFILTER defaults to using a quality estimation model for  $\text{QUALITY}(u)$ . When utilizing SUPERFILTER, GRAPHFILTER completes its tasks in 2.48 hours, highlighting its efficiency. Notably, without using the priority function  $\phi(u)$  for re-ranking, GRAPHFILTER becomes even faster, taking only 0.53 hours on a CPU. This is up to  $61\times$  faster than other baselines, compared to the 32.34 hours by ALPAGASUS.

Table 2: Runtime (in hours) for data selection approaches when selecting 10K training instances.  $\dagger$  indicate the CPU-only method.

	Runtime (hrs)
PERPLEXITY	0.92
ARMORM	5.93
ALPAGASUS	32.34
DEITA	22.65
SUPERFILTER	1.95
KMEANS	2.26
INSTAG	25.48
GRAPHFILTER	2.48
w/o priority $\phi(u)$	0.53 $^\dagger$

## 5 ANALYSIS

In this section, we perform an ablation study for GRAPHFILTER (Section 5.1), analyze the selected subsets (Section 5.2), highlight instruction diversity (Section 5.3), and examine the interplay between quality and diversity concerning subset sizes (Section 5.4).

### 5.1 ABLATION STUDY

**Combining n-grams captures features at different levels.** We examine the effectiveness of n-gram combinations, which are designed to capture both word-level and phrase-level features. The results are presented in Table 3. Our observations suggest that the variant of GRAPHFILTER, which integrates unigrams ( $n = 1$ ), bigrams ( $n = 2$ ), and trigrams ( $n = 3$ ), significantly outperforms other variations that do not incorporate n-gram combinations. Different n-grams capture features at varying levels, and merging them can effectively consolidate this information.




Table 3: Ablation study for n-gram combination of GRAPHFILTER with LLAMA-3-8B. ✓ indicates that various n-grams are used.

N-gram			$\mu_{\text{BENCH}}$	$\mu_{\text{LLM}}$	$\mu_{\text{ALL}}$
Unigram	Bigram	Trigram			
✓	✓	✓	50.55	42.79	47.97
✓			49.02	41.41	46.48
	✓		49.09	41.70	46.63
		✓	49.84	41.78	47.15

Table 4: Ablation study for  $\text{QUALITY}(u)$  and  $\text{DIVERSITY}(u)$  in the priority function of GRAPHFILTER with LLAMA-3-8B. ✗ indicates the component is not used.

$\text{QUAL}(u)$	$\text{DIV}(u)$	$\mu_{\text{BENCH}}$	$\mu_{\text{LLM}}$	$\mu_{\text{ALL}}$
SUPERFILTER	TF-IDF	50.55	42.79	47.97
PERPLEXITY	TF-IDF	49.21	40.85	46.43
✗	TF-IDF	48.94	41.87	46.58
SUPERFILTER	✗	49.52	41.28	46.78
✗	✗	48.27	40.28	45.61



(a) Quality-Diversity relationship (b) GRAPHFILTER vs. ARMORM (c) GRAPHFILTER vs. INSTAG

Figure 3: Figure 3a displays the quality-diversity relationships of subsets selected by different methods, with  $\uparrow$  indicating a preference for higher values. Figure 3b shows the semantic diversity in a t-SNE plot of subsets from GRAPHFILTER and ARMORM, where green rectangles indicate data points chosen by GRAPHFILTER but not by ARMORM. Figure 3c depicts the semantic diversity in a t-SNE plot comparing subsets from GRAPHFILTER and INSTAG.

**Both  $\text{QUALITY}(u)$  and  $\text{DIVERSITY}(u)$  in priority function enhance the data selection.** We provide empirical evidence in Table 4 showcasing the effectiveness of our proposed priority function. By incorporating the  $\text{QUALITY}(u)$  metric (using SUPERFILTER) and the  $\text{DIVERSITY}(u)$  metric (using TF-IDF) into GRAPHFILTER, we achieve superior performance across all evaluation metrics. This demonstrates that our combined priority function significantly enhances the model’s ability to select high-quality and diverse training data. Omitting either the quality metric (✗ + TF-IDF) or the diversity metric (SUPERFILTER + ✗) results in noticeable performance declines. Furthermore, replacing the SUPERFILTER metric with PERPLEXITY as the quality measure leads to reduced performance, highlighting the importance of using optimal metrics. These findings support our decision to integrate quality and diversity in the priority function.

## 5.2 WHAT DATA ARE SELECTED BY GRAPHFILTER?

**GRAPHFILTER effectively balances quality and diversity in its selected datasets.** In this section, we analyze the subsets selected by GRAPHFILTER and other methods, with results shown in Figure 3. To confirm that GRAPHFILTER maintains quality and diversity, we measure lexical diversity using the MTLT metric (McCarthy & Jarvis, 2010) and assess data quality with the advanced reward model, SKYWORKRM (Liu & Zeng, 2024).<sup>6</sup> As depicted in Figure 3a, GRAPHFILTER achieves the highest lexical diversity and ranks second in data quality. We also visualize GRAPHFILTER instructions compared with ARMORM and INSTAG using the BGE-LARGE-EN-V1.5 model.<sup>7</sup> It is evident that GRAPHFILTER selects instructions not chosen by ARMORM, shown by green rectangles in Figure 3b. Furthermore, Figure 3c illustrates that GRAPHFILTER and INSTAG exhibit similar semantic diversity. These results suggest that GRAPHFILTER not only selects high-quality data but also maximizes dataset diversity.

<sup>6</sup><https://huggingface.co/Skywork/Skywork-Reward-Llama-3.1-8B>

<sup>7</sup><https://huggingface.co/BAAI/bge-large-en-v1.5>

Table 5: Applying GRAPHFILTER to instructions and responses with LLAMA-3-8B. The ✓ indicates that GRAPHFILTER is applied. Lexical diversity is measured by MTLT (McCarthy & Jarvis, 2010), and quality is assessed using ARMORM, scaled by  $100\times$ .

	Content Type		Benchmarks			Lexical Diversity		Quality
	Inst.	Resp.	$\mu_{\text{BENCH}}$	$\mu_{\text{LLM}}$	$\mu_{\text{ALL}}$	Inst.	Resp.	
GRAPHFILTER	✓		50.55	42.79	47.97	102.43	71.74	81.54
		✓	47.16	39.71	44.68	90.22	73.57	81.52
	✓	✓	48.03	41.20	45.76	90.13	72.60	81.52

### 5.3 THE DIVERSITY OF INSTRUCTION AND RESPONSE: WHICH MATTERS MORE?

**Prioritizing instruction diversity most effectively improves model performance.** Each SFT training instance comprises an instruction and its response. This study evaluates the impact of applying GRAPHFILTER to instructions, responses, or both on model performance. As shown in Table 5, applying GRAPHFILTER only to instructions produces the best benchmark results, greatly improving lexical diversity in instructions with minimal effect on response diversity compared to other methods. Notably, all three variations maintain similar quality with different downstream outcomes, underscoring the importance of instruction diversity.

### 5.4 QUALITY AND DIVERSITY: WHICH SHOULD BE PRIORITIZED?

After showcasing GRAPHFILTER’s superiority in previous sections, an open question remains: *When should diversity be prioritized over quality, and vice versa?*

**The priority of quality and diversity varies with data selection budgets, and GRAPHFILTER excels at balancing these two factors effectively.** We hypothesize that the data selection budget plays a crucial role in determining the priority between quality and diversity and present the results in Figure 4. Our results indicate that the effectiveness of quality-based and diversity-based strategies is budget-dependent. Specifically, the quality-based SUPERFILTER excels with smaller budgets (1K and 5K instances), but its advantage diminishes as the budget increases. This suggests that quality-based methods with neural models may exhibit biases toward certain linguistic patterns, which limits model generalization when the budget is sufficiently large. Conversely, the diversity-based INSTAG performs poorly with small budgets but surpasses SUPERFILTER with larger ones. This observation demonstrates that diversity-based methods are more prone to introducing low-quality data with smaller budgets. Notably, GRAPHFILTER consistently achieves significant performance gains compared to RANDOM across all budget levels. These findings show that the data selection budget influences the effectiveness of different approaches, and GRAPHFILTER successfully integrates both quality and diversity.




Figure 4: Performance gap ( $\Delta_{\text{ALL}}$ ) with respect to  $\mu_{\text{ALL}}$ , comparing SUPERFILTER, INSTAG, and GRAPHFILTER against RANDOM, across various data selection budgets.

## 6 CONCLUSION

In this work, we introduce GRAPHFILTER, a novel method for data selection that models the dataset as a bipartite graph linking sentences to their constituent n-grams. To balance quality and diversity, we use a priority function that combines a quality metric with a diversity metric, allowing us to select subsets that enhance n-gram diversity and maintain high response quality. Our extensive experiments demonstrate GRAPHFILTER’s effectiveness across three model backbones and six benchmark datasets. Compared to nine baseline methods, GRAPHFILTER consistently delivers superior model performance and computational efficiency. Our analyses validate our design choices, assess the subsets chosen by GRAPHFILTER and other methods, highlight the importance of instruction diversity, and examine the role of quality and diversity relative to subset sizes.

## 7 ETHICS STATEMENT

In this work, we present GRAPHFILTER, a data selection method for supervised fine-tuning of large language models (LLMs). We acknowledge the ethical considerations related to data usage, potential biases, and the societal impact of LLMs. All datasets utilized in our experiments are publicly available and have been used extensively in prior research. We have adhered to all applicable licenses and terms of use for these datasets. However, we recognize that biases present in the training data can be propagated or even amplified by LLMs. To mitigate this risk, we recommend that practitioners applying GRAPHFILTER conduct thorough analyses of the selected data subsets to identify and address potential biases. Furthermore, while our goal is to enhance model performance and computational efficiency, we are aware that improved models could be misused in ways that are harmful or unethical. We advocate for the responsible deployment of LLMs and encourage users to follow ethical guidelines and best practices to prevent misuse.

## 8 REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our results presented in this paper. To facilitate replication and verification by the research community, we provide comprehensive details of our proposed method, GRAPHFILTER, in the main paper.

All hyperparameters, training configurations, and implementation specifics are thoroughly documented. For our experimental evaluations, we use publicly available datasets and benchmarks, namely MMLU (Hendrycks et al., 2021), ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), GSM8K (Cobbe et al., 2021), AlpacaEval-2.0 (Dubois et al., 2024), and MT-Bench (Zheng et al., 2023). All experiments were conducted using standard computational resources without the need for specialized hardware. Details about the computational setup and resource requirements are outlined in this work. By ensuring that all components of our work are transparently documented and accessible, we aim to facilitate reproducibility and encourage further exploration of our method by the research community.

We will release the source code of GRAPHFILTER and all scripts used for data selection, model training, and evaluation upon acceptance of this paper.

## REFERENCES

- Amro Abbas, Kushal Tirumala, Daniel Simig, Surya Ganguli, and Ari S. Morcos. Semdedup: Data-efficient learning at web-scale through semantic deduplication. *CoRR*, abs/2303.09540, 2023. doi: 10.48550/ARXIV.2303.09540. URL <https://doi.org/10.48550/arXiv.2303.09540>.
- Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Slav Petrov, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy P. Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul Ronald Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka, Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, and et al. Gemini: A family of highly capable multimodal models. *CoRR*, abs/2312.11805, 2023a. doi: 10.48550/ARXIV.2312.11805. URL <https://doi.org/10.48550/arXiv.2312.11805>.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernández Ábrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan A. Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa

- Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vladimir Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, and et al. Palm 2 technical report. *CoRR*, abs/2305.10403, 2023b. doi: 10.48550/ARXIV.2305.10403. URL <https://doi.org/10.48550/arXiv.2305.10403>.
- Zachary Ankner, Cody Blakeney, Kartik Sreenivasan, Max Marion, Matthew L. Leavitt, and Mansheej Paul. Perplexed by perplexity: Perplexity-based data pruning with small reference models. *CoRR*, abs/2405.20541, 2024. doi: 10.48550/ARXIV.2405.20541. URL <https://doi.org/10.48550/arXiv.2405.20541>.
- David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In Nikhil Bansal, Kirk Pruhs, and Clifford Stein (eds.), *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pp. 1027–1035. SIAM, 2007. URL <http://dl.acm.org/citation.cfm?id=1283383.1283494>.
- Jiuhai Chen, Rifaa Qadri, Yuxin Wen, Neel Jain, John Kirchenbauer, Tianyi Zhou, and Tom Goldstein. Genqa: Generating millions of instructions from a handful of prompts. *CoRR*, abs/2406.10323, 2024a. doi: 10.48550/ARXIV.2406.10323. URL <https://doi.org/10.48550/arXiv.2406.10323>.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. Alpapasus: Training a better alpaca with fewer data. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024b. URL <https://openreview.net/forum?id=FdVXgSJhvvz>.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the AI2 reasoning challenge. *CoRR*, abs/1803.05457, 2018. URL <http://arxiv.org/abs/1803.05457>.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021. URL <https://arxiv.org/abs/2110.14168>.
- Together Computer. Redpajama: an open dataset for training large language models, 2023. URL <https://github.com/togethercomputer/RedPajama-Data>.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. Ultrafeedback: Boosting language models with high-quality feedback. *CoRR*, abs/2310.01377, 2023. doi: 10.48550/ARXIV.2310.01377. URL <https://doi.org/10.48550/arXiv.2310.01377>.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. Enhancing chat language models by scaling high-quality instructional conversations. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 3029–3051, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.183. URL <https://aclanthology.org/2023.emnlp-main.183>.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lomakin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael

- Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, and Kevin Stone. The llama 3 herd of models. *CoRR*, abs/2407.21783, 2024. doi: 10.48550/ARXIV.2407.21783. URL <https://doi.org/10.48550/arXiv.2407.21783>.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B. Hashimoto. Length-controlled alpacaeval: A simple way to debias automatic evaluators. *CoRR*, abs/2404.04475, 2024. doi: 10.48550/ARXIV.2404.04475. URL <https://doi.org/10.48550/arXiv.2404.04475>.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, 07 2024. URL <https://zenodo.org/records/12608602>.
- M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. Textbooks are all you need. *CoRR*, abs/2306.11644, 2023. doi: 10.48550/ARXIV.2306.11644. URL <https://doi.org/10.48550/arXiv.2306.11644>.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. Mistral 7b. *CoRR*, abs/2310.06825, 2023. doi: 10.48550/ARXIV.2310.06825. URL <https://doi.org/10.48550/arXiv.2310.06825>.
- Haonan Li, Fajri Koto, Minghao Wu, Alham Fikri Aji, and Timothy Baldwin. Bactrian-x : A multilingual replicable instruction-following model with low-rank adaptation. *CoRR*, abs/2305.15011, 2023. doi: 10.48550/ARXIV.2305.15011. URL <https://doi.org/10.48550/arXiv.2305.15011>.
- 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. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14255–14273, Bangkok, Thailand, August 2024a. Association for Computational Linguistics. URL <https://aclanthology.org/2024.acl-long.769>.
- Ming Li, Yong Zhang, Zhitao Li, Jiuhai 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. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 7602–7635, Mexico City, Mexico, June 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.421. URL <https://aclanthology.org/2024.naacl-long.421>.

- Chris Yuhao Liu and Liang Zeng. Skywork reward model series. <https://huggingface.co/Skywork>, September 2024. URL <https://huggingface.co/Skywork>.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for alignment? A comprehensive study of automatic data selection in instruction tuning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=BTKAeLqLMw>.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. #instag: Instruction tagging for analyzing supervised fine-tuning of large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=pszewhybU9>.
- Adyasha Maharana, Prateek Yadav, and Mohit Bansal. D2 pruning: Message passing for balancing diversity & difficulty in data pruning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=thbtoAkCe9>.
- Max Marion, Ahmet Üstün, Luiza Pozzobon, Alex Wang, Marzieh Fadaee, and Sara Hooker. When less is more: Investigating data pruning for pretraining llms at scale. *CoRR*, abs/2309.04564, 2023. doi: 10.48550/ARXIV.2309.04564. URL <https://doi.org/10.48550/arXiv.2309.04564>.
- Philip M McCarthy and Scott Jarvis. Mtl-d, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392, 2010.
- Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, and et al. Gemma: Open models based on gemini research and technology. *CoRR*, abs/2403.08295, 2024. doi: 10.48550/ARXIV.2403.08295. URL <https://doi.org/10.48550/arXiv.2403.08295>.
- OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023. doi: 10.48550/arXiv.2303.08774. URL <https://doi.org/10.48550/arXiv.2303.08774>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022. URL [http://papers.nips.cc/paper\\_files/paper/2022/hash/blefde53be364a73914f58805a001731-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/blefde53be364a73914f58805a001731-Abstract-Conference.html).
- Guilherme Penedo, Hynek Kydlíček, Loubna Ben Allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro von Werra, and Thomas Wolf. The fineweb datasets: Decanting the web for the finest text data at scale. *CoRR*, abs/2406.17557, 2024. doi: 10.48550/ARXIV.2406.17557. URL <https://doi.org/10.48550/arXiv.2406.17557>.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020. URL <http://jmlr.org/papers/v21/20-074.html>.

- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=9Vrb9D0WI4>.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Author, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Jha, Sachin Kumar, Li Lucy, Xinxu Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma Strubell, Nishant Subramani, Oyvind Tafjord, Evan Walsh, Luke Zettlemoyer, Noah Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, and Kyle Lo. Dolma: an open corpus of three trillion tokens for language model pretraining research. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15725–15788, Bangkok, Thailand, August 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.acl-long.840>.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhu-patiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Fer-ret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Char-line Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchi-son, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Wein-berger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshv, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mo-hamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leti-cia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Mar-tins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khat-wani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Os-car Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Ko-cisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. Gemma 2: Improving open language models at a practical size, 2024. URL <https://arxiv.org/abs/2408.00118>.

- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023a. doi: 10.48550/ARXIV.2302.13971. URL <https://doi.org/10.48550/arXiv.2302.13971>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023b. doi: 10.48550/ARXIV.2307.09288. URL <https://doi.org/10.48550/arXiv.2307.09288>.
- Vijay V. Vazirani. *Approximation Algorithms*. Springer-Verlag, Berlin, Germany, 2001.
- Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. Interpretable preferences via multi-objective reward modeling and mixture-of-experts. *CoRR*, abs/2406.12845, 2024. doi: 10.48550/ARXIV.2406.12845. URL <https://doi.org/10.48550/arXiv.2406.12845>.
- 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. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 13484–13508, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.754. URL <https://aclanthology.org/2023.acl-long.754>.
- Minghao Wu, Abdul Waheed, Chiyu Zhang, Muhammad Abdul-Mageed, and Alham Fikri Aji. LaMini-LM: A diverse herd of distilled models from large-scale instructions. In Yvette Graham and Matthew Purver (eds.), *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 944–964, St. Julian’s, Malta, March 2024. Association for Computational Linguistics. URL <https://aclanthology.org/2024.eacl-long.57>.
- 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. *CoRR*, abs/2304.12244, 2023. doi: 10.48550/ARXIV.2304.12244. URL <https://doi.org/10.48550/arXiv.2304.12244>.
- Zhangchen Xu, Fengqing Jiang, Luyao Niu, Yuntian Deng, Radha Poovendran, Yejin Choi, and Bill Yuchen Lin. Magpie: Alignment data synthesis from scratch by prompting aligned llms with nothing. *CoRR*, abs/2406.08464, 2024. doi: 10.48550/ARXIV.2406.08464. URL <https://doi.org/10.48550/arXiv.2406.08464>.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yaqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report. *CoRR*, abs/2407.10671, 2024.



doi: 10.48550/ARXIV.2407.10671. URL <https://doi.org/10.48550/arXiv.2407.10671>.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. HellaSwag: Can a machine really finish your sentence? In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1472. URL <https://aclanthology.org/P19-1472>.

Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. Data-centric artificial intelligence: A survey. *CoRR*, abs/2303.10158, 2023. doi: 10.48550/ARXIV.2303.10158. URL <https://doi.org/10.48550/arXiv.2303.10158>.

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. Judging llm-as-a-judge with mt-bench and chatbot arena. *CoRR*, abs/2306.05685, 2023. doi: 10.48550/arXiv.2306.05685. URL <https://doi.org/10.48550/arXiv.2306.05685>.

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. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. URL [http://papers.nips.cc/paper\\_files/paper/2023/hash/ac662d74829e4407celd126477f4a03a-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/ac662d74829e4407celd126477f4a03a-Abstract-Conference.html).

## A EXPERIMENTAL SETUP

### A.1 BASELINES

In this work, we compare GRAPHFILTER against following baselines:

- **RANDOM** selects a random subset of size  $k$  from the entire dataset, where  $k$  is the designated data selection budget.
- **LONGEST** chooses the top- $k$  instances from the entire dataset, ranking them in descending order based on the number of words in each instruction.
- **PERPLEXITY** selects the top- $k$  instances from the entire dataset, sorted in descending order according to the perplexity values of the instructions. For the perplexity computation in this work, we utilize GPT2 (Radford et al., 2019).<sup>8</sup>
- **ARMORM** represents one of the state-of-the-art reward models (Wang et al., 2024). It evaluates multiple rewards from diverse perspectives and integrates these rewards using a gating network.
- **ALPAGASUS** employs GPT-3.5-TURBO to assess data quality (Chen et al., 2024b). Given the improved model performance and limited budget, we substitute GEMMA-2-27B-IT in this work, using the prompt illustrated in Figure 5. GEMMA-2-27B-IT is the state-of-the-art open large language model (LLM) and significantly surpasses GPT-3.5-TURBO according to the Chatbot Arena Leaderboard.<sup>9</sup>
- **DEITA** utilizes CHATGPT to create a quality estimation dataset and fine-tune large language models (LLMs) for evaluating data quality (Liu et al., 2024). We employ the official codes and models provided by Liu et al. (2024) for data selection.<sup>10</sup>
- **SUPERFILTER** refers to the Instruction-Following Difficulty (IFD) metric, which is calculated using smaller language models. Introduced by Li et al. (2024b), this method is shown by Li et al. (2024a) to provide IFD scores from smaller models that are as reliable as those from larger models. In this study, GPT2 is used for computing these scores (Radford et al., 2019).
- **KMEANS** involves clustering training instances using a state-of-the-art sentence embedding model and selecting instances that are nearest to their respective cluster centroids (Arthur & Vassilvitskii, 2007). In this work, we begin by sampling 50K instances from the entire dataset and encoding their instructions into sentence embeddings using the BGE-LARGE-EN-V1.5 model.<sup>11</sup> These embeddings are used for training the KMEANS model with 10K clusters. Once the KMEANS model is established, we cluster the sentence embeddings of instructions for the entire dataset and select the instances closest to each cluster centroid.
- **INSTAG** is designed to analyze the SFT dataset by tagging the topics of training instances. It can be used to select a subset with the most diverse topics from the entire dataset (Lu et al., 2024). We utilize the official codes and models released by Lu et al. (2024) for data selection.<sup>12</sup>

### A.2 OPTIMIZATION

**Hyperparameters** In this study, all experiments utilize the same set of hyperparameters. Specifically, we employ a batch size of 64, a learning rate of  $2 \times 10^{-5}$ , a warmup ratio of 0.05, and a linear learning rate schedule. All the experiments run for 3 epochs.

**Computation Infrastructure** For this study, all methods are trained using two A100 80GB GPUs, which are interconnected via PCIe.

<sup>8</sup><https://huggingface.co/openai-community/gpt2>

<sup>9</sup><https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

<sup>10</sup><https://github.com/hkust-nlp/deita>

<sup>11</sup><https://huggingface.co/BAAI/bge-large-en-v1.5>

<sup>12</sup><https://github.com/OFA-Sys/InsTag>

```

### System:
We would like to request your feedback on the performance of AI assistant in response to the instruction and
↳ the given input displayed following.

###Instruction:
{instruction}

### Input:
{input}

### Response:
{output}

### USER:
Please rate according to the accuracy of the response to the instruction and the input. Each assistant
↳ receives a score on a scale of 0 to 5, where a higher score indicates higher level of the accuracy.
↳ Please first output a single line containing value indicating the scores. In the subsequent line, please
↳ provide a comprehensive explanation of your evaluation, avoiding any potential bias.

```

Figure 5: The prompt used for ALPAGASUS annotation.

### A.3 EVALUATION

In this work, we evaluate the approaches on six widely used benchmarks:

- **MMLU** (Hendrycks et al., 2021) is a benchmark designed to assess knowledge acquired during pretraining, by evaluating models exclusively in zero-shot and few-shot settings. It covers 57 subjects across STEM, the humanities, social sciences, and more, totaling approximately 14,000 test examples.
- **ARC** (Clark et al., 2018) is a multiple-choice question-answering dataset containing questions from science exams for grades 3 to 9, amounting to approximately 4,000 test examples.
- **HellaSwag** (Zellers et al., 2019) is a challenging dataset for evaluating commonsense natural language inference, which is particularly difficult for state-of-the-art models, though its questions are trivial for humans. It contains approximately 10,000 test examples.
- **GSM8K** (Cobbe et al., 2021) comprises a collection of diverse grade school math word problems created by human problem writers, containing approximately 1,000 test examples.
- **AlpacaEval-2.0** (Dubois et al., 2024) is an automated tool for evaluating instruction-following language models. Its test set consists of 805 instructions generated by large language models (LLMs). Models are evaluated based on the winning rate against a reference answer, judged by a state-of-the-art LLM, such as GPT-4. AlpacaEval-2.0 is an upgraded version of the original AlpacaEval, featuring reduced length bias for a fairer evaluation of responses of varying lengths.
- **MT-Bench** (Zheng et al., 2023) is a multi-turn test set containing 80 questions that cover 8 aspects: writing, roleplay, reasoning, math, coding, extraction, STEM, and humanities. A state-of-the-art LLM, such as GPT-4, is used to score model outputs on a scale from 1 to 10.