

# A Survey on Data Selection for LLM Instruction Tuning

BOLIN ZHANG\*, Harbin Institute of Technology, China  
JIAHAO WANG\*, Institute of Automation, Chinese Academy of Sciences, China  
QIANLONG DU, Institute of Automation, Chinese Academy of Sciences, China  
JIAJUN ZHANG<sup>†</sup>, Institute of Automation, Chinese Academy of Sciences, China  
ZHIYING TU, Harbin Institute of Technology (Weihai), China  
DIANHUI CHU, Harbin Institute of Technology (Weihai), China

Instruction tuning is a vital step of training large language models (LLMs), so how to enhance the effect of instruction tuning has received increased attention. Existing works indicate that the quality of the dataset is more crucial than the quantity during instruction tuning of LLMs. Therefore, recently a lot of studies focus on exploring the methods of selecting high-quality subset from instruction datasets, aiming to reduce training costs and enhance the instruction-following capabilities of LLMs. This paper presents a comprehensive survey on data selection for LLM instruction tuning. Firstly, we introduce the widely used instruction datasets. Then, we propose a new taxonomy of the data selection methods and provide a detailed introduction of recent advances, and the evaluation strategies and results of data selection methods are also elaborated in detail. Finally, we emphasize the open challenges and present new frontiers of this task.

**JAIR Associate Editor:** Huang Xuanjing

## JAIR Reference Format:

Bolin Zhang, Jiahao Wang, Qianlong Du, Jiajun Zhang, Zhiying Tu, and Dianhui Chu. 2025. A Survey on Data Selection for LLM Instruction Tuning. *Journal of Artificial Intelligence Research* 83, Article 32 (August 2025), 14 pages. DOI: [10.1613/jair.1.17625](https://doi.org/10.1613/jair.1.17625)

## 1 Introduction

Large language models (LLMs), such as PaLM [5], GPT-4 [17], and LLaMA [26], have demonstrated impressive capabilities across a wide range of natural language understanding and generation tasks. Central to their success is instruction tuning—a process that aligns model outputs with human intent by fine-tuning pretrained models on curated (instruction, response) pairs [18, 34]. This technique not only improves the controllability and safety of LLMs but also enables them to adapt to specific domains with minimal computational overhead.

In the existing works [9, 28, 15], instruction tuning mainly focuses on constructing large-scale instruction-following datasets. These works can be categorized into two types: (1) converting annotated NLP datasets into instruction-style formats using templates (e.g., P3 [22] and FLAN [28]), and (2) automatically generating instruction-response pairs using powerful models like GPT-3.5 (e.g., Self-Instruct [27]). While these works have

\*Equal contribution.

<sup>†</sup>Corresponding author.

---

Authors' Contact Information: Bolin Zhang, ORCID: [0000-0002-6315-1856](https://orcid.org/0000-0002-6315-1856), [brolin@hit.edu.cn](mailto:brolin@hit.edu.cn), Harbin Institute of Technology, Harbin, Heilongjiang, China; Jiahao Wang, ORCID: [0009-0001-3280-851X](https://orcid.org/0009-0001-3280-851X), [jiahaowang0917@gmail.com](mailto:jiahaowang0917@gmail.com), Institute of Automation, Chinese Academy of Sciences, Beijing, , China; Qianlong Du, ORCID: [0000-0002-7187-2530](https://orcid.org/0000-0002-7187-2530), [qianlong.du@nlpr.ia.ac.cn](mailto:qianlong.du@nlpr.ia.ac.cn), Institute of Automation, Chinese Academy of Sciences, Beijing, , China; Jiajun Zhang, ORCID: [0000-0001-5293-7434](https://orcid.org/0000-0001-5293-7434), [jjzhang@nlpr.ia.ac.cn](mailto:jjzhang@nlpr.ia.ac.cn), Institute of Automation, Chinese Academy of Sciences, Beijing, , China; Zhiying Tu, ORCID: [0000-0001-8800-4513](https://orcid.org/0000-0001-8800-4513), [tzy\\_hit@hit.edu.cn](mailto:tzy_hit@hit.edu.cn), Harbin Institute of Technology (Weihai), , Shandong, China; Dianhui Chu, ORCID: [0000-0003-2973-7252](https://orcid.org/0000-0003-2973-7252), [chudh@hit.edu.cn](mailto:chudh@hit.edu.cn), Harbin Institute of Technology (Weihai), , Shandong, China.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.17625](https://doi.org/10.1613/jair.1.17625)

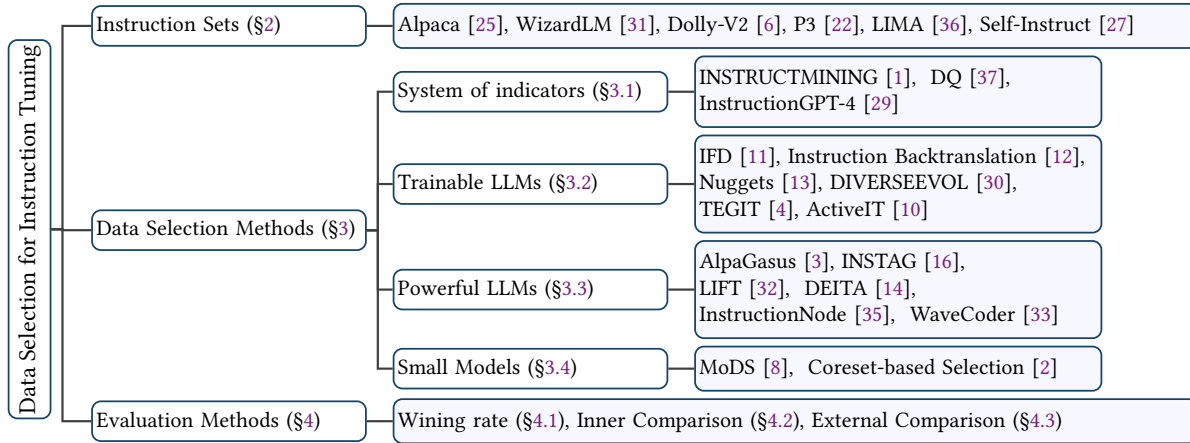


Fig. 1. Overview of Data Selection for LLM Instruction Tuning.

greatly expanded the availability of instruction data, they often suffer from issues related to data redundancy, limited diversity, and inconsistent quality. Moreover, these works primarily rely on a large volume of data, which can lead to a significant waste of computational resources and electricity.

In contrast to the previous approaches primarily focused on large-scale data construction [25, 31, 6], LIMA [36] proposes an approach centered on careful data selection. Their findings demonstrated that a model fine-tuned on merely 1,000 carefully curated instruction examples could achieve performance comparable, and in some cases superior, to models trained on vastly larger, often automatically generated, datasets. This seminal work underscored a critical insight: the quality of instruction data can significantly outweigh its quantity in determining the efficacy of instruction tuning. Consequently, this discovery has prompted a shift in research focus, moving from the large-scale generation of instruction data towards developing better methods for constructing and selecting high-quality, impactful, and efficient datasets.

In these data selection approaches, manual instruction data selection often involves high costs and introduces human bias. As a result, automated methods are crucial for instruction data selection. These methods have shown promising results. For instance, the IFD[11] method significantly outperforms the Alpaca model by using only about 5% of the Alpaca dataset, and it also surpasses the WizardLM model by about 10%. Using high-quality subsets for fine-tuning not only enhances the instruction-following skills of LLMs but also significantly reduces computational costs and time.

However, the data selection task for LLM instruction tuning is challenging due to the complex factors and multidimensional considerations involved. For example, it is difficult to assess the quality of individual instructions and ensure the overall diversity of the selected data. Another challenge is reducing costs and improving the efficiency of the selection process. In light of these factors, various data selection methods have been developed. Some methods use a system of indicators to evaluate individual data points, while others rely on trainable LLMs or powerful external LLMs. These exploit the LLMs' own capabilities to select instructions. Additionally, methods that use smaller models and design a comprehensive process to achieve a balanced effect across all aspects are noteworthy.

Despite these promising advances, the field remains fragmented, with no unified understanding of the principles, methods, and trade-offs involved in instruction data selection. In this survey, we aim to fill this gap by providing a systematic overview of instruction data selection for LLMs. We categorize existing methods, analyze their

core ideas and assumptions, and highlight open challenges such as the lack of standardized evaluation protocols, inefficiencies in processing large-scale instruction data, and the heavy dependence on powerful foundation models for selection. By synthesizing recent progress and identifying future directions, we hope to establish a foundation for further research and practical development in instruction tuning.

As illustrated by Fig. 1, this paper provides a comprehensive review of existing works on data selection methods for the instruction tuning of LLM. To facilitate the community, we maintain a paper list<sup>1</sup>, collect commonly instruction sets for data selection. Section 2 describes the mainstream datasets with different source and construction methods used in instruction tuning and Section 3 describes four types data selection methods in details: indicators sets, trainable LLMs, powerful LLMs and small models. Section 4 presents the evaluation methods and shows the results of different instruction selection methods. Section 5 summarizes the paper and emphasize the open challenges and future direction of the instruction selection.

## 2 Instruction Datasets

Various instruction tuning datasets (e.g. Self-Instruct and Alpaca), generated by LLMs, offer a wealth of samples without human labor, but their data quality depends on the performance of LLMs and is uncertain. Conversely, the manually curated datasets (e.g. LIMA and Dolly) obtain higher quality through meticulous human selection but are potentially influenced by human biases. Alternative dataset construction methods, like prompt mapping and evol-instruct, aim to enhance dataset quality and diversity but introduce new challenges in quality assurance. This variability in dataset construction and sourcing significantly affects data quality, highlighting the importance of careful data selection for the instructional tuning of LLM. This section describes the scale and construction procedures of several commonly instruction tuning datasets:

Self-instruct [27] consists of 52,000 training and 252 testing instructions. Initial instructions from seed tasks were selected, categorized, and diversified using InstructGPT [19] to generate inputs and outputs through output-first or input-first strategies. Post-processing refined the dataset for uniqueness and relevance, providing a versatile resource for natural language processing applications.

Alpaca [25] is a dataset of 52,002 instruction-response pairs generated using the Self-Instruct framework [27] with the text-davinci-003 model. It is designed to fine-tune LLaMA for improved instruction-following capabilities.

WizardLM [31] is a dataset containing 250,000 instruction-response pairs generated via evolutionary algorithms. Two types of algorithms (depth evolution and breadth evolution) are used to enhance the complexity and scope of basic instructions, generating more complex and diverse high-quality instruction data by ChatGPT.

LIMA [36] is a dataset consisting of 1,000 training samples, 300 test samples, and 50 development samples. In addition to the manually authored samples, samples collected from Q&A websites are strictly selected by human. Despite its modest scale, LIMA stands out for its meticulous compilation and design. LLMs fine-tuned on LIMA show remarkable ability in following instructions and adapting to unseen tasks.

DollyV2 [6] is a dataset of 15,000 human-authored instruction-response pairs covering tasks such as brainstorming, classification, question answering, and summarization. The data was created by employees using only Wikipedia as a source, with restrictions against using web content or generative AI tools when crafting responses.

P3 [22] is a large-scale prompt collection that integrates 170 NLP datasets with 2,052 handcrafted prompts, enabling unified training and evaluation across diverse language tasks. These prompts, also known as task templates, convert conventional NLP tasks like question answering or text classification into natural language input-output pairs. The P3 dataset itself is formed by randomly selecting prompts from PromptSource and organizing data into triplets of inputs, answer choices, and targets.

<sup>1</sup><https://github.com/Bolin97/awesome-instruction-selector>

### 3 Data Selection Methods

Formally, define an instruction dataset  $X$  with a size of  $n$ , where  $X = \{x_1, x_2, \dots, x_n\}$ , and each  $x_i$  represents an instance of instruction fine-tuning data. To employ a specific instruction data selection method  $\pi$  from  $X$  and choose a subset  $S_\pi^{(m)}$  with a size of  $m$ , we then use a predefined evaluation metric  $Q$  to assess the quality of  $S_\pi^{(m)}$ . The obtained subset quality, as measured by the evaluation metric, can gauge the effectiveness of the chosen instruction data selection method. The process of designing the selection method can be seen as:

$$S_\pi^{(m)} = \pi(X) \quad (1)$$

$$\pi^* = \arg \max_{\pi} Q(S_\pi^{(m)}) \quad (2)$$

The classification of instruction data selection methods is based on what scoring rules the method uses and what model base it employs. The methods can be divided into the following four categories: methods based on system of indicators, trainable LLMs, powerful LLMs like ChatGPT and small models.

#### 3.1 Methods Based on a System of Indicators

Methods based on a system of indicators aim to define a comprehensive set of evaluation metrics  $I_1, I_2, \dots, I_n$ , each governed by a specific computational formula. These metrics are designed to assess different aspects of data quality or relevance. In certain cases, the metrics may incorporate deep learning models for feature extraction, which function as learned indicators.

For a given data instance  $x_j$ , each metric  $I_i$  produces a score  $\text{score}_{ij} = I_i(x_j)$ . These individual metric scores are then aggregated using a combination function  $G(\cdot)$ , resulting in an overall score:

$$\text{score}_j = G(I_1(x_j), I_2(x_j), \dots, I_n(x_j)) \quad (3)$$

Once the scoring function  $G$  is defined, the system can be applied to the entire dataset  $X$  to evaluate each data instance. Based on a predefined threshold  $\tau$ , a high-quality subset  $S_\pi$  can be selected as follows:

$$S_\pi = \{x \in X \mid G(x) > \tau\} \quad (4)$$

This indicator-based framework provides a structured and scalable approach to data selection, enabling the integration of both heuristic and learned metrics.

INSTRUCTMINING [1] is a linear rule-based method for assessing the quality of instructional data. This approach initially identifies key natural language metrics, such as instruction length, perplexity, reward scores, KNN-i[21, 7], among others. These metrics are subsequently utilized to develop a linear equation. To explore the correlation between data quality and these metrics, and to define the parameters of equation, comprehensive fine-tuning experiments are conducted. Diverse quality datasets are partitioned into subsets, combined, and then employed in the fine-tuning of large-scale models. Quality labels for each subset are derived by evaluating performance of the model on a test set. The least squares method is applied to these experimental outcomes to estimate the parameters in INSTRUCTMINING. This involves fitting a linear equation to the loss observed when evaluating the model against the test set. Once the parameters have been determined, this formula can be utilized to calculate the quality of instructions, thereby facilitating data selection.

InstructionGPT-4 [29] is a data selection method designed for fine-tuning multimodal large language models. It surpasses MiniGPT-4 [38] in various assessments using less data. The first step, metrics like the CLIP Score[20], instruction length, and others are used. Visual and textual data are encoded into vectors and then subjected to dimensional reduction, which is regarded as special metrics. These metrics are combined into a vector  $e$ . The vector is then fed to a trainable data selector, such as a multi-layer perceptron or a self-attention network. This

approach, similar to the method [1] of calculating quality labels, employs clustering algorithms to segment datasets. Quality labels are assigned post fine-tuning and evaluation of each subset.

DQ [37] is an innovative data compression approach originally developed for large-scale computer vision datasets, and later adapted for use in the LLM domain. The approach involves several key steps: Initially, a gain function  $P(x)$  is defined as:

$$P(x) = \sum_{p \in S} \|f(p) - f(x)\|_2^2 - \sum_{p \in D \setminus S} \|f(p) - f(x)\|_2^2, \quad (5)$$

incorporating a feature function  $f(\cdot)$  – analogous to a metric – and using the current subset  $S$  and the entire dataset  $D$ . This gain function essentially forms a metric set function  $G(\cdot)$ . The dataset is then iteratively partitioned into non-overlapping subsets, guided by the gain function to maximize the defined gain. Subsequently, a representative sample from each subset is uniformly selected, ensuring coverage of the entire dataset while optimizing for maximum data diversity. This approach prioritizes maintaining the overall diversity of the dataset.

### 3.2 Methods Based on Trainable LLMs

This section introduces methods that leverage trainable large language models (LLMs) to develop scoring functions for instruction data selection. In these methods, the LLM itself is fine-tuned or adapted to act as a data selector, assigning a score to each instruction instance:

$$\text{score}_i = \text{LLM}_{\text{trainable}}(x_i) \quad (6)$$

Unlike indicator-based methods, which rely on predefined or hand-crafted metrics, trainable LLM-based methods learn to evaluate instruction quality directly from data. These methods not only focus on scoring individual instructions but also emphasize the alignment between the data selection process and the downstream performance of the fine-tuned model. Specific methods and frameworks that implement this paradigm are discussed in the following sections.

Instruction following difficulty (IFD) [11] represents a trainable data selection approach grounded in the hypothesis that large language models (LLMs) can first be trained on a carefully curated subset of instruction data to acquire basic instruction-following capabilities. This foundational skill then enables the model to better estimate instruction difficulty and overall data quality.

The method begins by fine-tuning an LLM on a small, clustered subset of the instruction dataset. It then introduces a novel metric, also named IFD, which quantifies the challenge of responding to a given instruction. Specifically, it compares the model's ability to generate the correct answer both with and without the given instruction. The *Conditioned Answer Score* is computed as:

$$s_{\theta}(A|Q) = \frac{1}{N} \sum_{i=1}^n \log P(w_i | Q, w_1, \dots, w_{i-1}), \quad (7)$$

where  $Q$  is the instruction and  $A$  is the target answer consisting of tokens  $w_1, \dots, w_n$ . This measures the model's likelihood of generating  $A$  when conditioned on  $Q$ .

The final IFD score is defined as the ratio between the conditioned and unconditioned scores:

$$r_{\theta}(Q, A) = \frac{s_{\theta}(A|Q)}{s_{\theta}(A)}, \quad (8)$$

where  $s_{\theta}(A)$  is the score of generating  $A$  without conditioning on  $Q$ . Higher IFD scores indicate samples where the instruction significantly influences the output, thereby reflecting greater instruction-following difficulty. By

applying a threshold on the IFD score, the method selects informative instructions to further refine the model in subsequent fine-tuning stages.

Instruction Backtranslation [12] is a method for generating and filtering instructions. Starting with a baseline instruction-following model and a web corpus, the model generates instructions for each web document, forming a dataset. Then the model is fine-tuned with seed instructions to gain basic capabilities. And it autonomously scores each instruction, with those exceeding a set threshold forming a high-score subset for further fine-tuning. This iterative process enhances instruction generation and filtering efficiency.

Nuggets [13] is a framework that implements a dual-phase approach. Initially, it employs a variety of predefined tasks to evaluate the proficiency of a LLM across multiple scenarios, a process known as zero-shot scoring. Subsequently, each entry in the instruction dataset is utilized as a unique prompt in a one-shot manner. These prompts are presented before the predefined tasks, and the performance of LLM is reassessed, a step termed one-shot scoring. This approach harnesses the disparity between one-shot and zero-shot scores to calculate a definitive 'gold score' for each instruction. Upon acquiring gold scores for all instructions, those constituting the highest-scoring subset are chosen as the 'gold subset'. This subset is then directly employed in the fine-tuning of the model. This methodology exploits the inherent contextual learning capabilities of extensive models.

DIVERSEEVOL [30] is an innovative iterative data selection mechanism. It leverages large-scale models like LLaMa to generate embedding vectors for instructional data. The mechanism employs the k-center-greedy algorithm[23] to facilitate diversity in selecting a subset of data for fine-tuning the LLaMa model. This procedure is repetitively applied, progressively enlarging the chosen subset, culminating in the creation of a high-quality instructional dataset.

TEGIT [4] is a novel method for generating high-quality instruction fine-tuning data. Particularly notable is their method of filtering instruction data. Utilizing ChatGPT, a small corpus of documentation is transformed into a format suitable for instruction data, forming a meta-dataset. This dataset then serves to train two Llama2 models - one functioning as a task generator and the other as a task discriminator. The role of generator is to devise tasks from provided texts, while the discriminator evaluates these tasks, ensuring their quality.

Active Instruction Tuning [10] is a method that focuses on task sensitivity selection to enhance large model fine-tuning with fewer tasks while improving out-of-task generalization. This method introduces the concept of Prompt Uncertainty, which is determined by generating k perturbed instructions through the random deletion of words from an original instruction. The output of LLM probability deviation is then averaged for these k perturbed instructions. Tasks that exhibit higher Prompt Uncertainty are prioritized for instruction fine-tuning, using the extent of Prompt Uncertainty as a measure of task uncertainty.

### 3.3 Methods Based on Powerful LLMs like ChatGPT

This section introduces the methods of using powerful LLMs such as GPT-4 and ChatGPT as data selectors. The approaches primarily involve designing prompt templates and leveraging the capabilities of LLMs to assess the quality of instruction data.

$$S_{\pi} = \{x | \text{ChatGPT}(\text{score} | \text{prompt}, x), \text{score} > \tau\} \quad (9)$$

ALPAGASUS [3] is an innovative data filtering methodology aimed at improving efficiency and accuracy in Instruction-Following Task (IFT) data curation. This approach utilizes a well-designed prompt applied to ChatGPT, to assess the quality of each data tuple, comprising instruction, input, and response. The methodology focuses on excluding tuples that fall below a predefined quality threshold. When this filtering process was applied to a substantial dataset, it was observed that a considerable portion of the data was compromised by quality issues. Notably, the application of the LLM-based filtering process led to the development of a model that surpassed

the performance of the original model, which was trained using the unfiltered dataset with instruction-based fine-tuning.

INSTAG [16] is an automated instruction tagging method that employs ChatGPT to generate detailed, open-ended labels for instructions, enhancing the diversity and complexity of the resulting subset. The process involves annotating data with labels that reflect the semantic and intent of each instruction and normalizing them for a label-based selection approach. This approach follows a complexity-first diversification sampling strategy. It starts by sorting queries in descending order of label count and then iteratively adds queries to a subset based on the uniqueness of their labels, until the subset reaches a desired size  $N$ . The result is a curated set of queries with a higher number of labels, indicating greater complexity and diversity.

To enhance the distribution and quality of datasets, [32] propose the method LIFT to reduce the redundancy of samples. It involves two stages: expanding dataset distribution and curating the diversity and quality of the dataset. Initially, ChatGPT enhances the data by generating diverse instructions and vectorizing them. A subset is then selected based on row variances. Secondly, ChatGPT scores instructions for accuracy, interpretability, clarity, difficulty, and length. The initial subset is re-selected based on these scores.

DEITA [14] integrates a multifaceted approach for selecting instruction data, focusing on complexity, quality, and diversity. Complexity describes factors such as the length, difficulty, and complexity of the instruction. And quality captures the accuracy of the output. Utilizing the WizardLM technique, ChatGPT is employed to augment instructions, which are then evaluated for complexity and quality. These evaluations involve scoring instructions based on complexity, using a specially trained complexity scorer, and assessing the output quality. Each instruction in the dataset is assigned complexity scores ( $c$ ) and quality scores ( $q$ ), and a comprehensive score is calculated by multiplying these two metrics. The dataset is then organized according to these comprehensive scores and vectorized for further analysis. To ensure diversity, a subset is created by adding samples that exceed a set distance threshold ( $\tau$ ) from their nearest neighbor in the subset. This process continues until the subset reaches a predetermined size.

InstructionNode [35] is a tree-instruct method, which improves the quality of instruction by enhance the complexity of it. DEITA [14] uses this method to measure instruction complexity and sets thresholds for instruction filtering. Tree-instruct employs GPT-4 to generate semantic parse trees for instruction data, using the number of nodes in the tree as a measure of complexity. The complexity is enhanced by adding nodes to the tree, and then GPT-4 is used to convert the new tree back into natural language, resulting in new high-quality instructions.

WaveCoder [33] is a code-focused LLM enhanced by an instruction improvement technique. Its training incorporates generated data, with the data filtering phase being particularly significant. Following data generation, a LLM-based discriminator, leveraging GPT-4, evaluates instruction data against established criteria divided into subtopics. This approach enables finer control over the filtering process, effectively eliminating low-quality instruction instances.

### 3.4 Methods Based on Small Models

This section introduces methods involving the use of external small models as scorers, or the transformation of instructions into embedding vectors using small models followed by further processing. Generally, these methods are quite comprehensive. The scoring filtering or embedding generation by Small Models often serves only as a part of the entire methodological process.

MoDS [8] focuses on instruction selection through three criteria: quality (instructional data fidelity), coverage (variety of instruction types), and necessity ( impact of instruction on LLM fine-tuning). The process unfolds in four key steps. Firstly, employing a reward model to evaluate the instruction dataset quality, selecting a subset denoted as  $D_h$ , comprising instructions that surpass a predefined quality threshold. Secondly, utilizing the k-center-greedy algorithm [23] to identify seed instructions, which ensures a diverse and representative

sample of the instruction dataset. Thirdly, fine-tuning a pre-trained LLM with the seed instructions. Subsequently, this refined model is applied to  $D_h$  to generate a new dataset,  $D_{inference}$ . This dataset is then evaluated using the reward model, which identifies instructions critical for learning of the LLM by focusing on those with lower scores. A threshold is established to select Augmented Instruction Data, which is specifically tailored to enhance performance of the model. Finally, the seed instructions with the Augmented Instruction Data to form a high-quality instruction subset that is poised to facilitate effective fine-tuning of the LLM.

The work [2] proposes a coreset-based and task-related data selection method: firstly, obtain sentence embeddings of samples by pre-trained language models (like BERT), then select the center point by applied unsupervised clustering on these embeddings, and finally use the KCenterGreedy [24] algorithm to retrieve core samples from the given dataset. This method effectively reduces the required training data volume, maintaining or potentially enhancing model performance.

### 3.5 Comparative Analysis of Data Selection Methods

In this subsection, we provide an in-depth comparison of the existing data selection methods across several key dimensions: model dependence, scalability, domain adaptability, interpretability, and selection effectiveness. The comparison results are shown in Table 1.

Table 1. Comparison of Instruction Data Selection Methods

Dimensions	Methods in Sec.3.1	Methods in Sec.3.2	Methods in Sec.3.3	Methods in Sec.3.4
Model Dependence	Model-agnostic and heuristic	Requires supervised model training	Depends on external black-box APIs (e.g., GPT-4)	Uses small pre-trained models (e.g., BERT, TinyLM)
Scalability	High	Low	Low	Moderate
Domain Adaptability	Low	High	Medium	Low to Medium
Interpretability	High	Low	Low	Moderate
Selection Effectiveness	Variable	High	High	Medium

The first dimension, model dependence, captures how much each method relies on external or trainable models. Indicator-based approaches are generally model-agnostic, relying on heuristic or statistical features derived directly from the data, without needing to query or train large models. In contrast, trainable LLM-based methods require supervised training of a dedicated selector or reward model, often built on large-scale LLMs, making them heavily model-dependent. Methods that use powerful LLMs like GPT-4 operate in a black-box manner, scoring or filtering instructions via external APIs, which introduces reliance on proprietary infrastructure. Small model approaches strike a balance by leveraging small-scale pretrained models (e.g., BERT or TinyLM), reducing resource demands while maintaining some model-based reasoning.

Indicator-based methods clearly excel in scalability: they are fast, parallelizable, and suitable for large-scale filtering without model inference. Trainable methods, by contrast, face serious scalability challenges due to the high computational cost of model training and inference. Black-box LLM-based methods suffer from scalability bottlenecks due to API latency, usage limits, and costs, making them impractical for processing massive datasets unless sufficient resources are available. Small model methods offer a middle ground—they are not as scalable as pure heuristics but can be applied to larger datasets with reasonable efficiency.

Domain adaptability varies notably among these methods. Trainable methods exhibit the highest adaptability, as their learning process can directly align with new domains through supervised fine-tuning on relevant data. Powerful LLMs also adapt reasonably well in few-shot or zero-shot scenarios, depending on the prompt design and instruction diversity. Indicator-based methods, however, tend to be less adaptive since their features are usually handcrafted and not easily transferable across domains. Small models may offer limited adaptability depending on their pretraining corpus and whether further domain-specific fine-tuning is conducted.

When it comes to interpretability, heuristic indicator methods clearly lead, offering intuitive, transparent criteria for data selection (e.g., length, diversity, perplexity). This is important for tasks requiring human oversight or explainable decisions. In contrast, both trainable and black-box LLM-based methods operate as opaque systems, making it difficult to trace why certain instances are chosen. Small models offer moderate interpretability—they can be analyzed via model outputs and saliency techniques but are less transparent than heuristics.

Finally, the selection effectiveness dimension compares how well each method identifies high-quality instruction data. Both trainable and powerful LLM-based methods tend to produce strong results, as they explicitly optimize or evaluate instruction usefulness using powerful language models. Small models offer moderate performance—better than pure heuristics but generally not on par with full-scale LLMs. Indicator-based methods exhibit variable performance, with effectiveness largely dependent on the specific metrics used and the task domain.

In summary, no single method dominates across all dimensions. Practitioners must balance trade-offs among efficiency, generalizability, and interpretability when choosing a data selection strategy, or consider hybrid designs that combine the strengths of multiple approaches.

#### 4 Evaluation methods and Result Analysis

To measure the effectiveness of different instruction data selection methods, several evaluation metrics are proposed and can be divided into three categories: winning rate, inner comparison, and external comparison. In this section, we first introduce each distinct evaluation and then compare different data selection methods across multiple benchmarks under these evaluations to analyze which methods achieve better performance.

Table 2. Performances of different selection methods on winning rate.

Selection methods	Model pairs for comparison	Training set	Win rate across various benchmarks	
			Vicuna/Koala/WizardLM/Self-inst/LIMA	total WS
IFD [11]	Llama-7B(5%), Llama-7B(full)	alpaca	1.125/0.97/1.077/1/1.1	1.04
	Llama-7B(10%), Llama-7B(full)		1.037/1.055/1.114/1.123/1.103	1.097
	Llama-7B(15%), Llama-7B(full)		1/1.038/1.114/1.027/1.09	1.064
	Llama-7B(10%), Llama-7B(full)	WizardLM	1.1625/1.1278/1.1147/1.0278/1.1067	1.0971
	Llama2-7B(5%), Llama2-7B(full)	alpaca	1.5875/1.4889/1.4266/1.2937/1.4733	1.4311
Random Sampling	Llama-7B(5%), Llama-7B(full)	alpaca	-	0.9
InstructionGPT4 [29]	miniGPT4(6%), miniGPT4(full)	cc_sbu_align	-	1.167
Alpagsaus [3]	Llama-7B(9k), Llama-7B(full)	alpaca	1.2125/1.0222/1.0596/1.0556/-	1.0658
	Llama-7B(9k), Llama-7B(3k)		1.1/1.183/1.082/1.17/-	1.074
	Llama-7B(9k), Llama-7B(6k)		1.05/1.072/1.05/1.087/-	1.082
	Llama-13B(9k), Llama-13B(full)		1.2125/1.0167/1.133/1.0198/-	1.074
MoDS [8]	Llama2-7B(2k), Llama2-7B(full)	alpaca	1.7125/1.5111/1.4725/1.369/1.4933	1.4786
InstructionMining [1]	Llama-7B(2k), Llama-7B(full)	dolly	-	1.088

'-' denotes values not reported in the original paper, and model(x) denotes tuning the model on the given training set using x samples or x% of samples.

#### 4.1 Wining Rate

To evaluate the effectiveness of the dataset selection methods, [3] define the winning rate of the LLM fine-tuned with less samples (i.e., LLM-sub) compared to the LLM fine-tuned with full samples (i.e., base LLM) and is calculated as:

$$\text{Win Rate} = \frac{\text{Num}(\text{win}) - \text{Num}(\text{lose})}{\text{Num}(\text{all})} + 1, \quad (10)$$

where Num(win) represents the number of wining cases, Num(lose) represents the number of losing cases, Num(all) represents the number of all cases in the testing benchmarks, and 1 indicates equal performance, values above 1 indicate improved performance, and values below 1 suggest degraded performance.

LLM-sub represents the LLM fine-tuned on the subset that filtered from the training set by the selection methods while the base LLM often involves 2 types: i) fine-tuned on the full training set and ii) fine-tuned on the same-scale subset filtered by the regular selections (e.g. random sampling and instruction length). The outputs of LLM-sub and the base LLM are rated by the judge on a scale from 1 to 10, and usually GPT4 is employed as the judge. To address the positional bias of the judge, IFD [11] send the outputs of these two LLMs to the judge twice with different orders. According to this work, the wining cases refer to LLM-sub outperforming the base LLM in both times or winning in one and tying in the other. The losing cases refer to LLM-sub lagging behind the base LLM in both times or tying in one and losing in the other. The wining rates of different selection methods on the testing benchmarks are summarized in Table 2.

Table 3. Inner comparisons of LLM tuned on subsets with itself tuned on full sets.

Selection methods	Training set(samples)	Base model	Win rate across various benchmarks
activeIT [10]	selfinstruct(2k)	Llama-7B	total WS 1.107
	selfinstruct(full)		1.293
DQ [37]	alpaca(20%)	Llama-7B	BBH/DROP/MMLU/Human-Eval/Avg 32.7/26.7/39.8/9.2/27.1
	alpaca(full)		32.9/26.3/41.6/10/27.7
DIVERSEEVOL [30]	Dolly(1k)	Llama-7B	Vicuna RS/Vicuna WTR/Koala RS/ Koala WTR 79.69/20/62.29/6.67
	Dolly(full)		73.84/5/57.9/3.33
	Selfins(1k)		79.16/7.5/66.95/6.11
	Selfins(full)		73.03/2.5/69.5/3.89
LIFT [32]	Platypus(15k random)		HellaSwag/ARC/TruthfulQA/MMLU 0.82/0.607/0.438/0.625
	Platypus(15k)		0.844/0.643/0.49/0.645
coreset [2]	P3(0.5)	Galactica-1.3B	RTE/CB/ANLI R1/ANLI R2/ANLI R3 74.73/73.21/49.6/41.9/43.75
	P3(full)		76.17/75/44/35.7/39.42

dataset(x) denotes tuning the base model on the x samples of the given set.

## 4.2 Inner Comparison

To evaluate the effectiveness of the dataset selection methods simply and directly, LLM-sub is compared with the same LLM, but it is fine-tuned on either the full training set or on a same-scale subset filtered by the regular selections. We refer to this evaluation method as inner comparison, because it only compares the LLM fine-tuned on a subset to itself. The inner comparison performances of different selection methods on the testing benchmarks are summarized in Table 3.

Inner comparison reveals the trade-off between dataset size and quality. Although models fine-tuned on the full training set generally perform well, carefully selected subset methods can achieve comparable or even better performance with fewer samples, demonstrating improved efficiency. However, the performance gains vary across different methods and tasks, highlighting the importance of adjusting selection strategies based on the characteristics of downstream tasks and the diversity of instructions.

## 4.3 External Comparison

Another simple and straightforward type of evaluation method is external comparison, which compares LLM-sub with the external LLMs (i.e. differ from the model of LLM-sub) on different testing benchmarks. The external comparison performances of different selection methods on the testing benchmarks are summarized in Table 4.

Table 4. External comparisons of LLM tuned on subsets with the other LLM.

Methods	Training set(samples)	Base model	MT-bench	AlpacaEval
gpt-4	✗	✗	8.99	95.28
gpt-3.5-turbo	✗	✗	7.94	91.36
alpaca-13B	alpaca	Llama-13B	4.53	-
NUGGETS[13]	alpaca(7.5k)	Llama-7B	5.34	-
TAGLM-13B-v1.0 [16]		Llama-13B	6.44±0.04	72.8
TAGLM-13B-v2.0 [16]		Llama2-13B	6.55±0.02	-
Instruction Length	mixture(6k)		5.89	-
Random Sampling		Llama-13B	5.84	-
IFD[11]			5.91	-
Instruction Node[35]			5.65	-
DEITA [14]	mixture(10k)	Llama2-13B	6.79	81.09

'mixture' denotes a combined dataset of WizardLM(Alpaca), WizardLM(ShareGPT), UltraChat, and ShareGPT.

External comparison helps reveal the generalization ability of data selection methods across different base models and architectures. For example, as shown in Table 4, IFD[11] demonstrates robust performance when transferring from Llama-7B to Llama2-7B, indicating that the selected data subsets contain general instruction information beneficial to multiple models. However, this evaluation also exposes some limitations: certain methods that perform well in inner comparison show diminished advantages in external comparison, possibly due to overfitting to the training distribution of specific models. For instance, although Alpagausaus [3] performs excellently in inner comparison, its performance is unstable when compared against multiple external LLMs.

## 4.4 Result Analysis

**The proposed selection methods outperform the regular selections, which proves the significance of the data selection in instruction tuning.** As shown in Tabel 2 and 3, TAGLM-13B-v1.0 and IFD outperform

the regular selections on MT-benc based on instruction length and random sampling when tuning Llama-13B on the mixture dataset, and Alpagsaus and IFD outperform the random sampling on total WS when tuning Llama-7B on alpaca.

**The more advanced LLM (Llama2-7B) achieves higher performance than the standard LLM (Llama-7B)**, both of which are fine-tuned on the same subset. Table 4 shows that TAGLM-13B-v2.0 outperforms TAGLM-13B-v1.0, and Table 2 shows that Llama2-7B(5%) outperforms Llama-7B(5%) when employing IFD selection method. These improvements are attributed to the inherent sophistication of the advanced LLMs, leading to higher learning efficiency on the subset.

**Instruction tuning the specific LLM on a larger subset does not necessarily guarantee performance improvement.** This may be related to the inherent features of the selection methods. As shown in Table 2, the performance of Llama-7B does not improve with the increase of the subset size when applying the method IFD in the training set alpaca. However, the performance of Llama-7B improves with the increase of the subset size when applying the method Alpagsaus in the same training set.

The above findings indicate that no single method holds an absolute advantage across all tasks. Instead, hybrid approaches—such as using metrics for preliminary filtering followed by detailed evaluation by large language models (LLMs)—may represent a more promising direction. Future evaluations should further take into account the characteristics of downstream tasks and the types of instructions, rather than relying solely on coarse-grained metrics like overall accuracy or win rate.

## 5 Discussion and Open Challenges

This paper provides a comprehensive overview of methods and challenges in instruction data selection for LLM tuning, emphasizing the crucial role of high-quality data during the fine-tuning stage. We review existing instruction datasets and their construction approaches, highlighting common issues such as imbalanced data distribution and inconsistent quality. On this basis, we categorize four representative types of data selection methods. Indicator-based methods can effectively assess sample-level quality but often overlook dataset-level characteristics. Trainable LLM-based methods adapt to model-specific preferences through scoring functions. Methods based on powerful external LLMs (e.g., GPT-4) offer superior general scoring capacity. Lastly, small model-based methods often adopt modular pipelines to approximate high-cost selections efficiently.

We also introduce three evaluation paradigms: win rate (against reference data), inner comparison (vs. full/small subsets with the same model), and external comparison (vs. other LLMs). Although existing methods show promising results, several challenges remain:

**Lack of unified evaluation standards.** As detailed in Section 4, different methods adopt different evaluation schemes and benchmarks, making it difficult to compare across techniques fairly. Future research should aim to develop standardized, automatic evaluation protocols that comprehensively assess selection quality under consistent conditions.

**High computational cost and dependence on large models.** When filtering hundreds of thousands of instructions, methods using large LLMs or APIs (e.g., GPT-4) can be prohibitively expensive and slow. There is a pressing need for efficient selection approaches that reduce reliance on external models, for example, by leveraging distilled or lightweight selectors that mimic the behavior of larger models.

**Limited coverage across domains and languages.** Most current selection frameworks are designed for English and general-purpose tasks. Extending them to other languages (e.g., Chinese, Arabic) and specialized domains (e.g., biomedical, legal) remains an open challenge.

To concretize the future directions above, we propose several specific and actionable research paths:

First, addressing the **evaluation inconsistency**, future work should build a standardized benchmark suite for instruction selection. This includes diverse downstream tasks (reasoning, summarization, code generation),

a variety of model architectures (e.g., LLaMA, Mistral), and multilingual datasets. In addition, we advocate for *task-conditioned evaluation*, where selection quality is judged by its impact on real-world applications rather than just coarse-grained win rates.

Second, to improve **efficiency and scalability**, a promising strategy is to develop *distillable selector models*—compact models trained to emulate powerful LLMs’ selection behavior. Another approach is *retrieval-augmented pruning*, in which fast heuristics narrow down candidate data, followed by LLM-based fine filtering only on top-ranked items, significantly reducing total cost.

Third, to enhance **domain and language adaptability**, future methods can integrate *multilingual encoders* (e.g., XLM-R) and incorporate *domain-specific priors*, such as metadata or instruction patterns. Modular design would allow swapping components to suit medical, legal, or programming domains, supporting cross-lingual and cross-domain generalization.

Lastly, we strongly encourage the **open-sourcing of reproducible selection pipelines**, including code, pretrained selector models, scoring results, and evaluation benchmarks. This will ensure fair comparisons and promote rapid iteration across the community.

**In summary**, instruction data selection is a critical but underexplored component in the instruction tuning pipeline. We hope this survey and analysis serve as a solid foundation to guide further research and practical improvements in making LLMs more instruction-following and efficient.

## Acknowledgments

This work is supported by National Key R&D Proram of China (No.2022ZD0160602 and No.2022YFF0902703), the National Natural Science Foundation of China (No.62472121), and the Special Funding Program of Shandong Taishan Scholars Project.

## References

- [1] Y. Cao, Y. Kang, and L. Sun. 2023. Instruction mining: high-quality instruction data selection for large language models. *CoRR*, abs/2307.06290.
- [2] H. Chen, Y. Zhang, Q. Zhang, H. Yang, X. Hu, X. Ma, Y. Yanggong, and J. Zhao. 2023. Maybe only 0.5% data is needed: A preliminary exploration of low training data instruction tuning. *CoRR*, abs/2305.09246.
- [3] L. Chen et al. 2024. Alpagasus: training a better alpaca with fewer data. In *The Twelfth International Conference on Learning Representations*.
- [4] Y. Chen, H. Jiang, X. Huang, S. Shi, and G. Qi. 2023. Tegit: generating high-quality instruction-tuning data with text-grounded task design. *CoRR*, abs/2309.05447.
- [5] A. Chowdhery et al. 2023. Palm: scaling language modeling with pathways. *Journal of Machine Learning Research*, 24, 240, 1–113.
- [6] M. Conover et al. 2023. Free dolly: introducing the world’s first truly open instruction-tuned llm. (2023).
- [7] W. Dong, M. Charikar, and K. Li. 2011. Efficient k-nearest neighbor graph construction for generic similarity measures. In *Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011*. S. Srinivasan, K. Ramamritham, A. Kumar, M. P. Ravindra, E. Bertino, and R. Kumar, (Eds.) ACM, 577–586.
- [8] Q. Du, C. Zong, and J. Zhang. 2023. Mods: model-oriented data selection for instruction tuning. *CoRR*, abs/2311.15653.
- [9] S. Ghosh, C. K. R. Evuru, S. Kumar, S. Ramaneswaran, D. Aneja, Z. Jin, R. Duraiswami, and D. Manocha. 2024. A closer look at the limitations of instruction tuning. In *International Conference on Machine Learning*. PMLR, 15559–15589.
- [10] P. Kung, F. Yin, D. Wu, K. Chang, and N. Peng. 2023. Active instruction tuning: improving cross-task generalization by training on prompt sensitive tasks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*. H. Bouamor, J. Pino, and K. Bali, (Eds.) Association for Computational Linguistics, 1813–1829.
- [11] M. Li, Y. Zhang, Z. Li, J. Chen, L. Chen, N. Cheng, J. Wang, T. Zhou, and J. Xiao. 2024. From quantity to quality: boosting llm performance with self-guided data selection for instruction tuning. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 7595–7628.
- [12] X. Li, P. Yu, C. Zhou, T. Schick, O. Levy, L. Zettlemoyer, J. E. Weston, and M. Lewis. 2024. Self-alignment with instruction backtranslation. In *The Twelfth International Conference on Learning Representations*.
- [13] Y. Li et al. 2024. One-shot learning as instruction data prospector for large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4586–4601.

- [14] W. Liu, W. Zeng, K. He, Y. Jiang, and J. He. 2024. What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning. In *The Twelfth International Conference on Learning Representations*.
- [15] S. Longpre et al. 2023. The flan collection: designing data and methods for effective instruction tuning. In *International Conference on Machine Learning*. PMLR, 22631–22648.
- [16] K. Lu, H. Yuan, Z. Yuan, R. Lin, J. Lin, C. Tan, C. Zhou, and J. Zhou. 2024. # instag: instruction tagging for analyzing supervised fine-tuning of large language models. In *The Twelfth International Conference on Learning Representations*.
- [17] OpenAI. 2023. Gpt-4 technical report. *CoRR*, abs/2303.08774.
- [18] L. Ouyang et al. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [19] L. Ouyang et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35, 27730–27744.
- [20] A. Radford et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event* (Proceedings of Machine Learning Research). M. Meila and T. Zhang, (Eds.) Vol. 139. PMLR, 8748–8763.
- [21] N. Reimers and I. Gurevych. 2019. Sentence-bert: sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992.
- [22] V. Sanh et al. 2022. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*.
- [23] O. Sener and S. Savarese. 2018. Active learning for convolutional neural networks: A core-set approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- [24] O. Sener and S. Savarese. 2018. Active learning for convolutional neural networks: A core-set approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- [25] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto. 2023. Stanford alpaca: an instruction-following llm model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca). (2023).
- [26] H. Touvron et al. 2023. Llama: open and efficient foundation language models. *CoRR*, abs/2302.13971.
- [27] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi. 2023. Self-instruct: aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*. A. Rogers, J. L. Boyd-Graber, and N. Okazaki, (Eds.) Association for Computational Linguistics, 13484–13508.
- [28] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le. 2021. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.
- [29] L. Wei, Z. Jiang, W. Huang, and L. Sun. 2023. Instructiongpt-4: A 200-instruction paradigm for fine-tuning minigpt-4. *CoRR*, abs/2308.12067.
- [30] S. Wu, K. Lu, B. Xu, J. Lin, Q. Su, and C. Zhou. 2023. Self-evolved diverse data sampling for efficient instruction tuning. *CoRR*, abs/2311.08182.
- [31] C. Xu, Q. Sun, K. Zheng, X. Geng, P. Zhao, J. Feng, C. Tao, and D. Jiang. 2023. Wizardlm: empowering large language models to follow complex instructions. *CoRR*, abs/2304.12244.
- [32] Y. Xu, Y. Yao, Y. Huang, M. Qi, M. Wang, B. Gu, and N. Sundaresan. 2023. Rethinking the instruction quality: LIFT is what you need. *CoRR*, abs/2312.11508.
- [33] Z. Yu, X. Zhang, N. Shang, Y. Huang, C. Xu, Y. Zhao, W. Hu, and Q. Yin. 2023. Wavecoder: widespread and versatile enhanced instruction tuning with refined data generation. *CoRR*, abs/2312.14187. arXiv: 2312.14187. doi: 10.48550/ARXIV.2312.14187.
- [34] S. Zhang et al. 2023. Instruction tuning for large language models: A survey. *CoRR*, abs/2308.10792.
- [35] Y. Zhao, B. Yu, B. Hui, H. Yu, F. Huang, Y. Li, and N. L. Zhang. 2023. A preliminary study of the intrinsic relationship between complexity and alignment. (2023). arXiv: 2308.05696 [cs.CL].
- [36] C. Zhou et al. 2023. LIMA: less is more for alignment. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [37] D. Zhou, K. Wang, J. Gu, X. Peng, D. Lian, Y. Zhang, Y. You, and J. Feng. 2023. Dataset quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17205–17216.
- [38] D. Zhu, J. Chen, X. Shen, X. Li, and M. Elhoseiny. 2024. Minigpt-4: enhancing vision-language understanding with advanced large language models. In *The Twelfth International Conference on Learning Representations*.

Received 14 November 2024; accepted 12 April 2025