

# Call for Rigor in Reporting Quality of Instruction Tuning Data

Hyeonseok Moon, Jaehyung Seo, Heuseok Lim<sup>†</sup>

<sup>†</sup>Department of Computer Science and Engineering, Korea University  
{glee889, seojae777, limhseok}@korea.ac.kr

## Abstract

Instruction tuning is crucial for adapting large language models (LLMs) to align with user intentions. Numerous studies emphasize the significance of the quality of instruction tuning (IT) data, revealing a strong correlation between IT data quality and the alignment performance of LLMs. In these studies, the quality of IT data is typically assessed by evaluating the performance of LLMs trained with that data. However, we identified a prevalent issue in such practice: hyperparameters for training models are often selected arbitrarily without adequate justification. We observed significant variations in hyperparameters applied across different studies, even when training the same model with the same data. In this study, we demonstrate the potential problems arising from this practice and emphasize the need for careful consideration in verifying data quality. Through our experiments on the quality of LIMA data and a selected set of 1,000 Alpaca data points, we demonstrate that arbitrary hyperparameter decisions can make any arbitrary conclusion.

## 1 Introduction

Instruction Tuning (IT) is a widely adopted strategy for enabling a human-interactive use of the knowledge embedded in large language models (LLMs) (Cao et al., 2023; Wang et al., 2024). By training with datasets composed of instruction-response pairs, LLM can attain the ability to generate appropriate responses to given instructions (Dubois et al., 2023; Zheng et al., 2023; Xu et al., 2023; Conover et al., 2023).

In implementing IT, data quality is considered a critical factor (Zhou et al., 2023a; Wang et al., 2024; Zhao et al., 2024b; Lu et al., 2024). Several studies have proven that selectively using high-quality IT data for training leads to better alignment performance than using the entire dataset (Liu et al., 2024b; Chen et al., 2024; Zhao et al., 2024a; Mekala et al., 2024).

Traditionally, the quality of IT data is measured by evaluating the performance of models trained on it (Liu et al., 2024b; Chen et al., 2024; Zhao et al., 2024a; Xia et al., 2024a). This approach stems from the consensus that *data is deemed good if it produces a good model*. Consequently, most studies on data quality establish a training configuration for models that represent data quality. Then, the performance of the trained model is regarded as the data quality. (Zhou et al., 2023a; Zhao et al., 2024a; Xia et al., 2024b; Du et al., 2023; Zhou et al., 2023b).

However, we observed that these studies often lack justification for the hyperparameter settings used in model training. Table 1 presents the diverse hyperparameter configurations utilized in previous research implementing IT with a sampled 1K general domain IT dataset. We discovered that the configurations may vary across studies, even when training the same model with identical data sizes.

In this study, we question whether reaching coherent conclusions under varying settings is possible. Specifically, we emphasize that conclusions regarding data quality can easily be altered based on arbitrarily chosen hyperparameter settings. For instance, even if one might report that dataset A is superior to dataset B, another could claim that B is better by training models under different settings, even with the same dataset, model, and test settings. This variability poses a risk of causing significant confusion.

As a representative case, we consider two general-domain IT datasets: LIMA (Zhou et al., 2023a) and sampled 1K dataset from Alpaca (Alpaca-longest (Zhao et al., 2024a)). In (Zhao et al., 2024a), it was reported that a model trained on Alpaca-longest outperformed a model trained on LIMA. However, our experiments contrarily demonstrate that LIMA can also be regarded as better than Alpaca-longest, depending on the selected training setting. Given the current research trend

Paper	Epochs	LR	LR Scheduler	Batch	Data Pool
<i>Training Llama-2-7B with sampled 1K general domain IT data</i>					
Ghosh et al. (2024)	3	5e-5	-	32	Lima
Raghavendra et al. (2024)	3	1e-5	-	8	Dolly
Yu et al. (2024)	3	1e-5	-	64	Alpaca / WizardLM
Du et al. (2023)	3	2e-5	Cosine	128	Alpaca+HC3+WizardLM+Dolly+Self-Instruct+Lima
Li et al. (2024)	3	2e-5	-	128	Alpaca / WizardLM
Liu et al. (2024a)	3	2e-5	Cosine	128	Alpaca-gpt4 / Lima
Mekala et al. (2024)	3	2e-5	Cosine	128	Alpaca / Dolly
Kong et al. (2024)	8	1e-5	Cosine	64	Lima
Zhao et al. (2024a)	15	1e-5	Linear	128	Alpaca / WizardLM / Lima
Zhou et al. (2023a)	15 (ES)	1e-5	Linear	64	Lima
<i>Training Llama-2-13B with 1K general domain IT data</i>					
Ghosh et al. (2024)	3	5e-5	-	32	Lima
Zhao et al. (2024b)	10	1e-4	-	16	Alpaca-gpt4
Liu et al. (2024a)	3	2e-5	Cosine	128	Alpaca-GPT4 / Lima
Mekala et al. (2024)	3	2e-5	Cosine	128	Alpaca / Dolly
Zhao et al. (2024a)	15	1e-5	Linear	128	Alpaca / WizardLM / Lima
<i>Training Mistral 7B with 1K general domain IT data</i>					
Kong et al. (2024)	4	1e-5	Cosine	64	Lima
Zhao et al. (2024a)	15	2e-6	Linear	128	Alpaca / WizardLM / Lima
Ghosh et al. (2024)	3	5e-5	-	32	Lima
Yu et al. (2024)	3	1e-5	-	64	Alpaca / WizardLM
Yin et al. (2024)	4	4e-6	-	128	WizardLM / UltraChat / ShareGPT

Table 1: Hyperparameters reported by previous studies, adopted to train LLMs with 1K general domain IT data. The data pool details the sources from which the 1K data samples were drawn. Detailed descriptions of these data pools are provided in the Table 4. The '+' symbol indicates experiments where samples were drawn from a combined data mix of all mentioned datasets. The '/' symbol reports studies that sampled individually from each data pool.

of arbitrarily determining hyperparameters for validation models, this confusion can be identified as a severe yet persistent problem.

Through our experiments, we emphasize the necessity of rigor in reporting data quality. Furthermore, our discussion suggests the importance of identifying (at least) locally optimal hyperparameters and reporting data quality under these settings.

## 2 Related Works

Previous research has widely acknowledged the importance of data quality in performing IT. Chen et al. (2024) proposed that training LLMs with a small, carefully selected subset of high-quality data can significantly improve alignment performance within the vast IT data pool. Furthermore, Zhou et al. (2023a) even suggested that carefully curated high-quality 1,000 data points are sufficient to attain alignment performance for LLMs. Motivated by these findings, numerous studies are exploring various methodologies focused on selecting high-quality instruction tuning data (Wang et al., 2024; Chen et al., 2024; Zhao et al., 2024a; Xia et al., 2024b; Lu et al., 2024; Liu et al., 2024b).

However, most studies lack justification for the selected hyperparameter setting to train verification models. Consequently, the training setups become diversified even when using the same LLM and

dataset. We argue that the importance of selecting appropriate hyperparameters has long been emphasized (Yu and Zhu, 2020; McCandlish et al., 2018; Halfon et al., 2024). The community widely recognizes that optimal hyperparameters are often specific to particular LLMs and datasets, and reported performance may vary based on the experimental setup (Van Rijn and Hutter, 2018; Jin, 2022; Gkouti et al., 2024; Bi et al., 2024). However, we find that research on data quality frequently reports performance without adequately considering these factors.

In this study, we highlight the potential confusion that can result from neglecting these considerations and demonstrate the necessity of a rigorous experimental setup to report data quality.

## 3 Experimental Setting

### 3.1 Exam-taker Dataset

In our experiments, we adopt two general domain IT datasets, each comprising 1,000 samples, as our exam-taker datasets. By comparing the quality of these two datasets, we examine how the judgment on the exam-taker datasets varies with different arbitrarily chosen hyperparameter settings.

**LIMA (Zhou et al., 2023a)** LIMA is a high-quality dataset comprising 1,000 IT data points,




Figure 1: The performance comparison between the two models trained with LIMA and Alpaca-Longest. We train Llama-2-7B model with each dataset, We evaluate the data quality when training each dataset with the Llama-2-7B model. is depicted on the Y-axis represents the hyperparameter settings used in each experiment. We bolded the settings that consistently demonstrated conclusive results across all three evaluation datasets.

carefully curated by human efforts with an emphasis on quality and diversity.

**Alpaca-Longest (Zhao et al., 2024a)** Zhao et al. (2024a) selected the 1,000 entries with the longest token lengths from the Alpaca dataset (Taori et al., 2023). This approach proved more effective than training on the entire Alpaca dataset and significantly outperformed other baselines such as Alpargus (Chen et al., 2024). According to the original paper, training with this data resulted in higher alignment performance than LIMA.

### 3.2 Experimental Model

The quality of the Exam-taker Dataset is determined by the performance of the experimental model trained on it. We conduct experiments using the Llama-2-7B model (Touvron et al., 2023) and the Mistral-7B-v0.3 model (Jiang et al., 2023). The main paper reports the results for the Llama-2-7B model, and Appendix A includes the results for the Mistral-7B model.

### 3.3 Experimental Setting

This study focuses on four commonly reported hyperparameters: learning rate, learning rate scheduler, batch size, and number of epochs. We report the experimental results obtained from varying these parameters. We conduct comparative experiments for each setting by choosing two prevalent

yet distinct values. While numerous other potential variations exist, such as weight decay and dropout, we leave these for future exploration. Apart from the hyperparameters under investigation, detailed experimental settings are provided in Appendix C.

### 3.4 Test Dataset

To evaluate the performance of the trained model, we use three LLM alignment benchmarks: Koala (Geng et al., 2023), MT-Bench (Zheng et al., 2023), and Self-Instruct (Wang et al., 2023). These benchmarks serve as an instruction-following evaluation tool, assessing LLMs by evaluating the quality of text generated in response to given instructions. We employ GPT-4o (Hurst et al., 2024)<sup>1</sup> as a judge to compare the performance of experimental models for each benchmark. The judge prompts used in the experiments are detailed in Appendix D.

## 4 Experimental Results

### 4.1 LIMA vs Alpaca-Longest

Figure 1 presents the experimental results based on hyperparameter variations. Our results show that if we choose specific settings (*e.g.*, Settings 4, 5, 10, 12, 13), we can report that Alpaca-longest exhibits superior data quality compared to LIMA. At the same time, if we choose other configurations

<sup>1</sup><https://openai.com/index/hello-gpt-4o/>

Dataset	LIMA			Alpaca-Longest		
Comparison with Setting 1	Setting 1 Wins	Tie	Setting $x$ Wins	Setting 1 Wins	Tie	Setting $x$ Wins
vs Setting 2	112	10	58	121	7	52
vs Setting 3	101	7	72	97	10	73
vs Setting 4	142	7	31	144	10	26
vs Setting 5	98	9	73	102	6	72
vs Setting 6	114	6	60	109	7	64
vs Setting 7	83	8	89	73	7	100
vs Setting 8	121	8	51	136	9	35
vs Setting 9	86	10	84	84	14	82
vs Setting 10	116	9	55	130	9	41
vs Setting 11	101	5	74	96	9	75
vs Setting 12	146	7	27	145	9	26
vs Setting 13	96	7	77	102	8	70
vs Setting 14	124	8	48	114	5	61
vs Setting 15	82	7	91	78	6	96
vs Setting 16	109	10	61	145	5	30

Table 2: We report the performance of the Llama-2-7B model, trained under each setting, as evaluated on the Koala dataset. Details for each setting are presented in the Figure 1.

(*e.g.*, Settings 8, 16), we can report that LIMA still demonstrates higher data quality.

Considering that authors have determined such hyperparameters arbitrarily, this represents a significant concern. We view that the ability to alter reported conclusions based on subjective decisions can severely undermine the reliability of scientific discussions.

## 4.2 Among the Same Dataset

Then, which setting should we choose to report? Considering that the primary goal of the IT dataset is to construct high-performance models, it would be reasonable and practical to report results based on the best achievable performance with the given data (Koehn et al., 2018, 2020; Budach et al., 2022; Van Rijn and Hutter, 2018).

In this section, we identify the optimal settings among the configurations tested. We recognize that other configurations with better performance may have been overlooked. We focus on local optimality within our considered settings and discuss its implications. Figure 2 compares model performance across various hyperparameter settings, using **Setting1** as the baseline.

Our experiments reveal that **Setting7, 15** (2e-5 LR / 256 Batch / 15 Epochs) maximizes model performance within our study. Notably, we can find that such configurations are far beyond the widely chosen settings in existing research. As our brief survey in Table 1 indicates, most studies opt to train Llama-2-7B for only three epochs when using 1K IT datasets. However, our results show that this setup yielded significantly lower performance than training for 15 epochs under the same

conditions. This finding suggests that the reported performance in many studies may reflect the under-trained performance of models, which may fail to fully represent the potential of the exam-taker dataset.

## 5 Discussion

**We argue that it is inevitable to evaluate the downstream model performance.** We acknowledge that assessing data quality through the performance of a trained model can be ambiguous. However, we also argue that the quality of training data must inevitably be assessed through the model’s performance after training.

We would like to discuss how the quality of training data is generally acknowledged. The goal of constructing the training dataset is to develop a model that aligns with the intended purpose. Thus, in terms of training data, "good data" is defined as data that produces a "good model" (Chen et al., 2024; Koehn et al., 2020). This fundamentally differs from constructing benchmark datasets for evaluation. Since training data’s primary aim is to build a strong model, data that appear high-quality to humans (or any frontier LLMs) may offer little value if the trained model’s performance remains subpar (Liu et al., 2024b).

In this context, the quality of training data is fundamentally linked to the performance of the model trained on it. While there are various plausible methods to assess training data, these methods might remain indirect indicators, without validating with the performance of the trained model.

Consequently, most studies demonstrate the quality of the data under evaluation by training it onto one (possibly several) model and reporting their performance. We do not consider this approach erroneous; instead, we view it as a natural and inevitable choice. Our stance is that if model-based verification is unavoidable, a more thorough and rigorous training configuration would be essential to verify data quality. Our experiments demonstrate that hyperparameters can introduce unintended biases that skew the objective evaluation of the data quality.

**We argue that authors researching data quality have responsibility for such validation.** To address these ambiguities, we suggest selecting the hyperparameter setting that yields the highest performance within a given data and reporting the model’s performance under this setting. This ap-



proach seeks to evaluate the model by maximizing the data’s potential.

Given that hyperparameter search is being performed in relatively small-sized PLMs (e.g., BERT (Devlin et al., 2019), BART (Lewis et al., 2020)) (Latif and Kim, 2024; Ljubešić et al., 2024; Roele, 2021), we argue that it is challenging to justify its omission in LLMs other than its high cost. Even when researchers do not conduct their own hyperparameter search, there have been multiple attempts to use existing configurations (Zhou et al., 2023b). However, as shown in Table 1 with the example of Mistral, there appears to be no established standard configuration when tuning relatively recent LLMs.

Reporting the best performance would certainly require additional costs for experiments, but we believe this is a necessary sacrifice to strengthen scientific discourse. We argue that arbitrary conclusions stemming from arbitrary hyperparameter choices pose a greater risk than incurring additional costs. While a comprehensive hyperparameter search may not always be necessary, we claim that authors should clearly justify their chosen hyperparameters. Even if they do not report peak performance, employing the best settings from our paper or established training configurations (ex. LIMA configuration) would still be a rational approach.

## 6 Conclusion

In our examination of various studies addressing data quality, we observed a recurring issue where researchers often arbitrarily select hyperparameters when training models to verify data quality. Our experiments reveal that arbitrary hyperparameter choices can lead to arbitrary conclusions. Moreover, we found that hyperparameters chosen without justification often fail to achieve optimal performance on the exam-taker datasets, resulting in unreliable conclusions. To address this, we propose establishing a local hyperparameter pool and training models under locally optimal settings within this pool. While additional costs for hyperparameter validation are inevitable, we consider this a necessary sacrifice for attaining the reliability of scientific discourse. To ensure rigorous reporting and sustainable consensus, we urge careful attention.

## Limitation

The numerous hyperparameter settings we did not consider may remain a limitation of our study.

Within our budget constraints, we verified as many possibilities as possible. Fortunately, we found significant variations in model performance even within the four factors we examined, allowing us to draw generalized conclusions. Through our brief survey, we can also find that various other hyperparameter variants, such as weight decay, warmup steps, are also introduced without justification. Exploring additional possibilities to identify an optimal setup could be a meaningful area for future research.

In this study, we did not attempt hyperparameter optimization (HPO), as finding optimal values was outside the scope of our research. Applying HPO when reporting on data quality, could serve as an excellent direction for future research.

Although we conducted experiments using only two datasets, we do not see this as a limitation. We believe this setup clearly demonstrates the inherent ambiguity in reporting data quality. Possibly numerous other datasets can exist where hyperparameter settings could alter reporting conclusions. Instead of identifying additional dataset pairs, we consider it more valuable to focus future research on strategies to mitigate such ambiguities.

## Ethics Statement

We do not challenge the quality assessment results reported by the Zhao et al. (2024a), which proposed Alpaca-Longest and estimated that Alpaca-Longest is better than LIMA. Based on the results from our experiments using locally optimal settings within our pool, Alpaca-Longest can arguably be considered superior to Lima. This finding aligns with the results reported in the original paper. In all experiments, we exclusively used artifacts approved for research purposes.

## Acknowledges

This work was partly supported by ICT Creative Consilience Program through the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (IITP-2025-RS-2020-II201819, 25%), Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIT) (RS-2024-00398115, Research on the reliability and coherence of outcomes produced by Generative AI, 25%), Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Min-

istry of Education(NRF-2021R1A6A1A03045425, 25%), and Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2024 (Project Name : Development of generative AI-based publishing content analysis and sharing platform technology to respond to changes in the publishing environment. Project Number : : RS-2024-00442061, 25%)

## References

- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. 2024. Deepseek llm: Scaling open-source language models with longtermism. *CoRR*.
- Lukas Budach, Moritz Feuerpfeil, Nina Ihde, Andrea Nathansen, Nele Noack, Hendrik Patzlaff, Felix Naumann, and Hazar Harmouch. 2022. The effects of data quality on machine learning performance. *arXiv preprint arXiv:2207.14529*.
- Yihan Cao, Yanbin Kang, Chi Wang, and Lichao Sun. 2023. Instruction mining: Instruction data selection for tuning large language models. *arXiv preprint arXiv:2307.06290*.
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srinivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. 2024. [Alpagasus: Training a better alpaca with fewer data](#). In *The Twelfth International Conference on Learning Representations*.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. Free dolly: Introducing the world’s first truly open instruction-tuned llm. *Company Blog of Databricks*.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [Flashattention: Fast and memory-efficient exact attention with io-awareness](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 16344–16359. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. [Enhancing chat language models by scaling high-quality instructional conversations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3029–3051, Singapore. Association for Computational Linguistics.
- Qianlong Du, Chengqing Zong, and Jiajun Zhang. 2023. Mods: Model-oriented data selection for instruction tuning. *arXiv preprint arXiv:2311.15653*.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Xinyang Geng, Arnav Gudibande, Hao Liu, Eric Wallace, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. Koala: A dialogue model for academic research. *Blog post, April*, 1:6.
- Sreyan Ghosh, Chandra Kiran Reddy Evuru, Sonal Kumar, Ramaneswaran S, Deepali Aneja, Zeyu Jin, Ramani Duraiswami, and Dinesh Manocha. 2024. [A closer look at the limitations of instruction tuning](#). In *Forty-first International Conference on Machine Learning*.
- Nefeli Gkouti, Prodromos Malakasiotis, Stavros Toumpis, and Ion Androutsopoulos. 2024. [Should I try multiple optimizers when fine-tuning a pre-trained transformer for NLP tasks? should I tune their hyperparameters?](#) In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2555–2574, St. Julian’s, Malta. Association for Computational Linguistics.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. How close is chatgpt to human experts? comparison corpus, evaluation, and detection. *arXiv preprint arXiv:2301.07597*.
- Alon Halfon, Shai Gretz, Ofir Arviv, Artem Specator, Orith Toledo-Ronen, Yoav Katz, Liat Ein-Dor, Michal Shmueli-Scheuer, and Noam Slonim. 2024. Stay tuned: An empirical study of the impact of hyperparameters on llm tuning in real-world applications. *arXiv preprint arXiv:2407.18990*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

- Honghe Jin. 2022. Hyperparameter importance for machine learning algorithms. *arXiv preprint arXiv:2201.05132*.
- Dhiraj Kalamkar, Dheevatsa Mudigere, Naveen Mellem-pudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharma Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, et al. 2019. A study of bfloat16 for deep learning training. *arXiv preprint arXiv:1905.12322*.
- Philipp Koehn, Vishrav Chaudhary, Ahmed El-Kishky, Naman Goyal, Peng-Jen Chen, and Francisco Guzmán. 2020. Findings of the wmt 2020 shared task on parallel corpus filtering and alignment. In *Proceedings of the Fifth Conference on Machine Translation*, pages 726–742.
- Philipp Koehn, Huda Khayrallah, Kenneth Heafield, and Mikel L. Forcada. 2018. [Findings of the wmt 2018 shared task on parallel corpus filtering](#). In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, pages 739–752, Belgium, Brussels. Association for Computational Linguistics.
- Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Jiaming Zhou, and Haoqin Sun. 2024. [Self-prompt tuning: Enable autonomous role-playing in llms](#). *CoRR*, abs/2407.08995.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- Atif Latif and Jihie Kim. 2024. Evaluation and analysis of large language models for clinical text augmentation and generation. *IEEE Access*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Ming Li, Yong Zhang, Zhitao Li, Jiuhai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing 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)*, pages 7602–7635, Mexico City, Mexico. Association for Computational Linguistics.
- Liangxin Liu, Xuebo Liu, Derek F Wong, Dongfang Li, Ziyi Wang, Baotian Hu, and Min Zhang. 2024a. Selectit: Selective instruction tuning for large language models via uncertainty-aware self-reflection. *arXiv preprint arXiv:2402.16705*.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024b. [What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning](#). In *The Twelfth International Conference on Learning Representations*.
- Nikola Ljubešić, Vít Suchomel, Peter Rupnik, Taja Kuzman, and Rik van Noord. 2024. [Language models on a diet: Cost-efficient development of encoders for closely-related languages via additional pretraining](#). In *Proceedings of the 3rd Annual Meeting of the Special Interest Group on Under-resourced Languages @ LREC-COLING 2024*, pages 189–203, Torino, Italia. ELRA and ICCL.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2024. [#instag: Instruction tagging for analyzing supervised fine-tuning of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Sam McCandlish, Jared Kaplan, Dario Amodei, and OpenAI Dota Team. 2018. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*.
- Dheeraj Mekala, Alex Nguyen, and Jingbo Shang. 2024. Smaller language models are capable of selecting instruction-tuning training data for larger language models. *arXiv preprint arXiv:2402.10430*.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*.
- Mohit Raghavendra, Vaskar Nath, and Sean Hendryx. 2024. Revisiting the superficial alignment hypothesis. *arXiv preprint arXiv:2410.03717*.
- Cees Roele. 2021. [WVOQ at SemEval-2021 task 6: BART for span detection and classification](#). In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 270–274, Online. Association for Computational Linguistics.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deep-speed and megatron to train megatron-turing nl-g 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990*.
- Dusan Stosic and Paulius Micikevicius. 2021. Accelerating ai training with nvidia tf32 tensor cores. *NVIDIA Technical Blog*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).



- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Jan N Van Rijn and Frank Hutter. 2018. Hyperparameter importance across datasets. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2367–2376.
- Jiahao Wang, Bolin Zhang, Qianlong Du, Jiajun Zhang, and Dianhui Chu. 2024. A survey on data selection for llm instruction tuning. *arXiv preprint arXiv:2402.05123*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh 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)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. 2024a. [LESS: Selecting influential data for targeted instruction tuning](#). In *Forty-first International Conference on Machine Learning*.
- Tingyu Xia, Bowen Yu, Kai Dang, An Yang, Yuan Wu, Yuan Tian, Yi Chang, and Junyang Lin. 2024b. Rethinking data selection at scale: Random selection is almost all you need. *arXiv preprint arXiv:2410.09335*.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Mingjia Yin, Chuhan Wu, Yufei Wang, Hao Wang, Wei Guo, Yasheng Wang, Yong Liu, Ruiming Tang, Defu Lian, and Enhong Chen. 2024. [Entropy law: The story behind data compression and llm performance](#). *CoRR*, abs/2407.06645.
- Simon Yu, Liangyu Chen, Sara Ahmadian, and Marzieh Fadaee. 2024. Diversify and conquer: Diversity-centric data selection with iterative refinement. *arXiv preprint arXiv:2409.11378*.
- Tong Yu and Hong Zhu. 2020. Hyper-parameter optimization: A review of algorithms and applications. *arXiv preprint arXiv:2003.05689*.
- Hao Zhao, Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. 2024a. [Long is more for alignment: A simple but tough-to-beat baseline for instruction fine-tuning](#). In *Forty-first International Conference on Machine Learning*.
- Yingxiu Zhao, Bowen Yu, Binyuan Hui, Haiyang Yu, Minghao Li, Fei Huang, Nevin L. Zhang, and Yongbin Li. 2024b. [Tree-instruct: A preliminary study of the intrinsic relationship between complexity and alignment](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16776–16789, Torino, Italia. ELRA and ICCL.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-bench and chatbot arena](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- 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. 2023a. [Lima: Less is more for alignment](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 55006–55021. Curran Associates, Inc.
- Daquan Zhou, Kai Wang, Jianyang Gu, Xiangyu Peng, Dongze Lian, Yifan Zhang, Yang You, and Jiashi Feng. 2023b. Dataset quantization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17205–17216.



## A Experimental Results - Mistral

We conducted the same experiments described in Section 4 using the Mistral-7B model. The results are reported in Table 3 and Figure 2.

Dataset	LIMA			Alpaca-Longest		
Comparison with Setting 1	Setting 1 Wins	Tie	Setting $x$ Wins	Setting 1 Wins	Tie	Setting $x$ Wins
vs Setting 2	79	12	<b>89</b>	58	11	<b>111</b>
vs Setting 3	75	12	<b>93</b>	50	14	<b>116</b>
vs Setting 4	<b>100</b>	8	72	<b>107</b>	14	59
vs Setting 5	<b>116</b>	13	51	<b>126</b>	13	41
vs Setting 6	<b>97</b>	17	66	<b>95</b>	15	70
vs Setting 7	<b>93</b>	12	75	70	14	<b>96</b>
vs Setting 8	<b>125</b>	8	47	<b>130</b>	12	38
vs Setting 9	<b>87</b>	16	77	<b>82</b>	18	80
vs Setting 10	<b>83</b>	14	<b>83</b>	58	15	<b>107</b>
vs Setting 11	81	7	<b>92</b>	45	6	<b>129</b>
vs Setting 12	<b>102</b>	10	68	<b>96</b>	9	75
vs Setting 13	<b>124</b>	13	43	<b>129</b>	18	33
vs Setting 14	<b>96</b>	22	62	<b>87</b>	16	77
vs Setting 15	<b>88</b>	14	78	78	14	<b>88</b>
vs Setting 16	<b>109</b>	9	62	<b>118</b>	10	52

Table 3: We report the performance of the Mistral-7B model, trained under each setting, as evaluated on the Koala dataset. Details for each setting are presented in the Figure 1.

As shown in Table 3, models trained for 15 epochs generally outperformed those trained for only 3 epochs, even within the same settings. This finding suggests that commonly adopted hyperparameter settings in prior research may not be optimal and that reported performance might not fully exploit the data’s potential.

Figure 2 illustrates the potential conclusions we can draw from various settings using Mistral. There is still significant diversity between settings, supporting our earlier conclusions in Section 4. We demonstrate that merely using multiple models is insufficient to enhance robustness in data quality validation, emphasizing the necessity of hyperparameter generalization.

## B Dataset Details

Dataset	Paper / Description	Data Size
Alpaca	Taori et al. (2023)	52K
Alpaca-GPT4	Peng et al. (2023)	52K
Dolly	Conover et al. (2023)	15K
HC3	Guo et al. (2023)	24.3K
ShareGPT	Chiang et al. (2023)	52K
UltraChat	Ding et al. (2023)	200K
WizardLM	Xu et al. (2023)	700K

Table 4: We report only the data aimed at performing IT in a general domain, which are adopted to previous studies. Each dataset consists of a pair, featuring a human instruction and an appropriate response.

## C Experimental Details

We conducted experiments with a weight decay of 0.0, a warmup of 0.0, and a maximum length of 2,048, utilizing the HuggingFace trainer (Wolf et al., 2020). To enhance learning efficiency, we applied bf16 (Kalamkar et al., 2019) and tf32 (Stosic and Micikevicius, 2021) strategy. All training was performed using FlashAttention-2 (Dao et al., 2022) and DeepSpeed Stage 2 (Smith et al., 2022). For inference, we employed vllm (Kwon et al., 2023). Our setup included four RTX-A6000 GPUs with 48GB each for model training and inference. The original batch size per GPU was set to 2, and we used gradient accumulation to increase the batch size. Other settings followed the default configurations provided by the HuggingFace trainer.

## D LLM-as-a-Judge

### ## System Prompt

Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below.

You should choose the assistant that follows the user’s instructions and answers the user’s question better. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of their responses.

Begin your evaluation by comparing the two responses and provide a short explanation. Avoid any position biases and ensure that the order in which the responses were presented does not influence your decision.

Do not allow the length of the responses to influence your evaluation. Do not favor certain names of the assistants. Be as objective as possible.

After providing your explanation, output your final verdict by strictly following this format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a tie.

### ## Input Statements

You are a helpful and precise assistant for checking the quality of the answer.

[Question]

{question}

[The Start of Assistant 1’s Answer]

{Response From Assistant 1}

[The End of Assistant 1’s Answer]

[The Start of Assistant 2’s Answer]

{Response From Assistant 2}

[The End of Assistant 2’s Answer]

Table 5: Prompt used for training the LLM: For models not supporting system prompts, we combined the system prompt and user prompt into a single input statement.

The prompt we used is presented in Table 5. In




Figure 2: The performance comparison between the two models trained with LIMA and Alpaca-Longest. We train Mistral-7B model with each dataset, We evaluate the data quality when training each dataset with the Mistral-7B model. is depicted on the Y-axis represents the hyperparameter settings used in each experiment. We bolded the settings that consistently demonstrated conclusive results across all three evaluation datasets.

all our experiments, we randomize the order of presented responses to relieve any unintended effects driven by the positional bias. We conducted our experiments with GPT-4o (gpt-4o-2024-08-06), setting the temperature to 0 and top-p to 1.0. The API usage cost for the experiments detailed in Table 2 was \$11.25. Conducting a similar hyperparameter search using GPT4o-mini incurred a cost of \$1.44. While using GPT4o-mini presents a cost-effective option, the Pearson-r correlation score between GPT4o and GPT4o-mini was 0.559 in our experiments. Although this score might be considered reasonably high, we argue using GPT-4o is more effective for establishing a more precise and rigorous setting. We leave experiments with alternative judges for future research.