

From Quantity to Quality: Boosting LLM Performance with Self-Guided Data Selection for Instruction Tuning

Anonymous ACL submission

Abstract

In the realm of Large Language Models, the balance between instruction data quality and quantity has become a focal point. Recognizing this, we introduce a self-guided methodology for LLMs to autonomously discern and select cherry samples from vast open-source datasets, effectively minimizing manual curation and potential cost for instruction tuning an LLM. Our key innovation, the Instruction-Following Difficulty (IFD) metric, emerges as a pivotal metric to identify discrepancies between a model’s expected responses and its autonomous generation prowess. Through the adept application of IFD, cherry samples are pinpointed, leading to a marked uptick in model training efficiency. Empirical validations on datasets like Alpaca and WizardLM underpin our findings; with a mere 10% of conventional data input, our strategy showcases improved results. This synthesis of self-guided cherry-picking and the IFD metric signifies a transformative leap in the optimization of LLMs, promising both efficiency and resource-conscious advancements. Codes, data, and models will be available.

1 Introduction

Large Language Models (LLMs) have revolutionized the landscape of artificial intelligence (Touvron et al., 2023a,b; Penedo et al., 2023; Scao et al., 2022). Notable models such as GPT-3 (Brown et al., 2020) and GPT-4 (OpenAI, 2023) leverage extensive datasets and advanced training methodologies to exhibit high-level text understanding and generation capabilities. The applications of these models extend across diverse domains, including interactive systems, automated content generation, and support for scientific inquiries.

Instruction fine-tuning (Wei et al., 2022; Longpre et al., 2023) is a method employed to refine the performance of LLMs by providing specific guidelines or instructions during the model’s training phase. It operates by supplying the LLM with

explicit training instructions to produce the corresponding outputs that are more congruent with the desired responses. A well-formulated instruction or prompt provides essential contextual information, refining the model’s capability to generate relevant and task-specific outputs (Taori et al., 2023; Ouyang et al., 2022).

Based on the findings from Wang et al. (2022b) and Self-Instruct (Wang et al., 2023b) early experiments, reducing the number of instances per task does not degrade the model’s generalization performance to unseen tasks. While conventionally, instruction tuning is predominantly relied on amassing vast datasets. A seminal revelation from the LIMA (Zhou et al., 2023) highlights the art of instruction tuning: rather than a sheer volume of data, it’s the quality of the data that dictates the model’s performance. LIMA’s findings emphasize that even a limited amount of manually curated, high-quality data can elevate the model’s instruction-following prowess. While it underscores the efficacy of data overabundance, the question of how to automatically identify high-quality data from a vast ocean of available datasets remains under investigation.

To bridge this gap, our study puts forth a method for autonomously sifting through expansive open-source datasets to discover the most impactful training samples. We coin these samples as “cherry data,” designating those data fragments that hold the potential to exponentially enhance LLM instruction tuning. At the heart of our research is the hypothesis that during their preliminary training stages with carefully chosen instruction data, LLMs can develop an intrinsic capability to discern instructions. This foundational understanding equips them with the discernment to assess the quality of broader datasets, thus making it possible to estimate the instruction-following difficulty in a self-guided manner.

Our methodology encompasses a self-guided approach to extract cherry data from the target dataset,

subsequently training a more refined cherry model. This process pivots from the prevailing data quantity emphasis to prioritizing data quality. Initially, the model is familiarized with a fraction of the target dataset during the “Learning from Brief Experience” phase. This preliminary knowledge paves the way for the subsequent “Evaluating Based on Experience” phase, where we meticulously evaluate the model’s response generation.

Under the hypothesis that LLMs have acquired all critical knowledge during pre-training, our goal is to devise a metric for evaluating instruction-following difficulty. Adhering to the objective of minimizing cross-entropy loss in model training, we use this metric to gauge the training data’s impact, with a specific focus on isolating the instructional component from the answer’s influence. We achieve this by comparing the loss when the model generates a response with and without instructional context. This comparison forms the basis of our Instruction-Following Difficulty (IFD) score, which assesses the extent of instructional assistance in response generation. A higher IFD score, indicating minimal instructional help, suggests greater difficulty. We propose that selecting samples with moderate IFD scores during instruction tuning strikes an optimal balance, effectively improving the model’s ability to process and follow complex instructions, essential for its learning and adaptation to advanced tasks.

Harnessing these insights, the “Retraining from Self-Guided Experience” phase utilizes cherry data with standout IFD scores to hone the model, culminating in our superior cherry models. The net result is a model that aligns more adeptly with instructions, ensuring enhanced performance. This self-guided approach starkly contrasts with existing techniques (Chen et al., 2023), which typically engage external models for data curation.

Extensive experimental results validate the efficacy of our method. By applying our methodology to the Alpaca and WizardLM instruction tuning datasets, our model outperforms the official Alpaca model with only approximately 5% data selected and outperforms the reimplemented WizardLM model with approximately 10% data selected. The key contributions of this paper:

- We propose a self-guided approach enabling models to autonomously “select cherry data” from vast open-source datasets. This innovation minimizes manual curation and optimizes the use of existing data resources, reducing

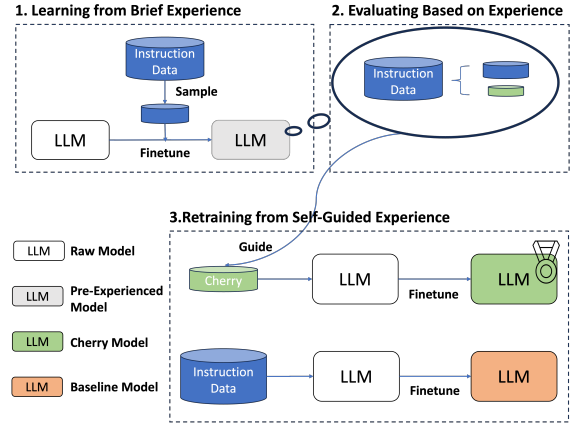


Figure 1: Overview of our proposed method.

costs and streamlining training.

- We introduce the Instruction-Following Difficulty (IFD) metric as a tool to identify gaps in a model’s responses versus its autonomous generation capability. Using the IFD metric, we can pinpoint these cherry samples, optimizing model training efficiency.
- Backed by validation on datasets like Alpaca and WizardLM, our strategy demonstrates enhanced outcomes with only 10% of the typical data input, emphasizing our approach’s efficiency and transformative impact.
- We provide a different model-specific view in measuring the difficulty of new instructions, which may benefit future instruction data generation work.

2 Methodology

As illustrated in Figure 1, our methodology is divided into three core phases: Learning from Brief Experience, Evaluating Based on Experience, and Retraining from Self-Guided Experience. The initial phase emphasizes equipping the model with a basic instruction-following capability using select portions of the dataset. The subsequent phase introduces a novel metric to evaluate the instruction-following difficulty score of each sample based on the previously trained pre-experienced model. Finally, after obtaining difficulty scores in the target dataset, the cherry samples are defined and sampled to train our final model, which we call the cherry models. In our experiments, the underlying model used is the Meta LLaMA (Touvron et al., 2023a), complemented by the target dataset.

2.1 Learning from Brief Experience

This phase aims to equip the initial model with a basic instruction-following capability by forc-

ing the model to first experience a subset of the target dataset. Specifically, for the initial full target dataset, D_0 contains n triplets $x = (Instruction, [Input], Answer)$, we define the string $Question = map(Instruction, [Input])$ as the complete instruction. The map function is aligned with the original target dataset. Each word in $Question(Q)$ and $Answer(A)$ is denoted as x_i^Q and x_i^A respectively. Let LLM_θ denote the LLM we use and θ represent the weight of LLMs, specifically, θ_0 represents the pre-trained base LLM model. Then the instruction embeddings for each sample x_j are obtained by:

$$[h_{j,1}^Q, \dots, h_{j,m}^Q] = LLM_{\theta_0}(w_{j,1}^Q, \dots, w_{j,m}^Q) \quad (1)$$

$$h_j^Q = \frac{\sum_{i=1}^m h_{j,i}^Q}{m} \quad (2)$$

where $w_{j,i}^Q$ represents the i_{th} word of $Question$ strings of sample j and $h_{j,i}^Q$ represents its corresponding last hidden states.

To ensure the diversity of instructions exposed to the initial model, the basic clustering technique K-Means on these instruction embeddings is utilized. Motivated by LIMA’s finding, we are motivated to make this experience process as brief as possible by sampling only a few instances in each cluster which we call pre-experienced samples. Specifically, we generate 100 clusters on instruction embeddings and sample 10 instances in each cluster. Then the initial model is trained for only 1 epoch with these samples to obtain our brief pre-experienced model.

2.2 Evaluating Based on Experience

In the instruction-tuning process, the loss of a sample pair (Q, A) is calculated by continuously predicting the next tokens given the instruction Q and their proceeding words:

$$L_\theta(A|Q) = -\frac{1}{N} \sum_{i=1}^N \log P(w_i^A|Q, w_1^A, w_2^A, \dots, w_{i-1}^A; \theta) \quad (3)$$

where N is the number of words of the ground-truth answer A . We denote this averaged cross-entropy loss as the **Conditioned Answer Score** $s_\theta(A|Q) = L_\theta(A|Q)$.¹ This metric evaluates the model’s capability to generate appropriate responses based on provided instructions. It measures the extent to which the model’s output aligns

¹We use different symbols to differentiate the loss when used as an objective function and the loss when used as a score.

with both the instruction and the corresponding correct answer.

Under this circumstance, a higher $s_\theta(A|Q)$ does not mean a harder instruction to follow, it may simply be caused by the inherent factor of string A itself. In the pre-LLM era, when models are required to learn both the knowledge and instruction-following ability during finetuning, it is reasonable to simply use $s_\theta(A|Q)$ as an indicator for the difficulty of a sample. However, things change a little for current LLMs, which have learned most of the knowledge in the pre-training phase and only need to learn to align and follow the instructions. To estimate the difficulty of following instructions of a given sample, we introduce the **Direct Answer Score** $s_\theta(A)$:

$$s_\theta(A) = -\frac{1}{N} \sum_{i=1}^N \log P(w_i^A|w_1^A, \dots, w_{i-1}^A; \theta). \quad (4)$$

which measures LLM’s ability to generate this answer alone. This metric gauges the inherent difficulty or challenge posed by the answer in isolation, without the contextual guidance from its corresponding instruction. A higher direct answer score may suggest that the answer is inherently more challenging or intricate for the model to generate.

Further, analyzing the balance between a sample’s inherent challenge and the model’s capabilities in following it sheds light on the intricacies of estimating the difficulty of the instruction of a given sample. Specifically, we try to estimate the **Instruction-Following Difficulty (IFD)** scores $IFD_\theta(Q, A)$ on following instruction of a given (Q, A) pairs by calculating the ratio between $s_\theta(A)$ and $s_\theta(A|Q)$:

$$IFD_\theta(Q, A) = \frac{s_\theta(A|Q)}{s_\theta(A)} \quad (5)$$

Under this circumstance, the influence of LLM’s intrinsic ability to fit the answer string is partially alleviated. The score measures the degree how a given instruction benefits the alignment of the corresponding response. High IFD scores infer the inability of the model to align responses to the given corresponding instructions, which in turn indicates the difficulty of an instruction. It is worth noting that this $IFD_\theta(Q, A)$ is a model-specific value, and we use our pre-experienced model to obtain all these values in the target dataset.

2.3 Retraining from Self-Guided Experience

In alignment with the Introduction’s concept of “cherry data” — high-impact training samples that significantly enhance LLM performance — and its impact on LLM performance, this phase involves systematically labeling each instance in the target dataset and sorting them based on their Instruction-Following Difficulty (IFD) scores. Aligned with previous work (Xu et al., 2023) which emphasizes the difficulty of samples, we define those with relatively higher IFD scores as the cherry samples. Moreover, the dense scores from our method provide better flexibility in deciding how much data to be used. Then these cherry samples are used upon initial pre-trained LLM to obtain the cherry models. The whole process is in a self-guided manner and no external models are required.

3 Experimental Setup

3.1 Datasets

Training Datasets The Alpaca dataset (Taori et al., 2023) from Stanford University, encompasses 52002 instruction-following samples. Developed using the self-instruct (Wang et al., 2023b) approach with text-davinci-003. Though initially competitive, its dependence on text-davinci-003 posed data quality concerns. WizardLM dataset (Xu et al., 2023) leverages the Evol-Instruct algorithm to improve the quality of instruction data. Furthermore, the incorporation of ChatGPT during the reformulation guarantees high fidelity of data. Of its 250,000 instructions, we primarily utilized the WizardLM-7b subset, consisting of 70,000 samples.

Test Datasets To ensure comprehensive and unbiased assessment, we employed 5 diverse test sets: Vicuna (Chiang et al., 2023), Koala (Vu et al., 2023), WizardLM (Xu et al., 2023), Self-instruct (Wang et al., 2023b), and LIMA (Zhou et al., 2023). These test sets contain approximately 1000 human curated instructions, closed-domain or closed-domain for different tasks from different sources. Among them, Vicuna and WizardLM further provide the specific sub-category for each instruction, making it possible for in-depth analysis. Although all of these test sets are introduced to guarantee the variety of testing, we select all of these sets to offer a broader palette of instruction types than the typical one or two test sets.

3.2 Implementation Details

Rooted in the Llama-7b pre-trained model, our training framework aligns with protocols from Alpaca and WizardLM datasets. The Adam optimizer (Kingma and Ba, 2017), with a 2×10^{-5} learning rate and a batch size of 128, steers the training across three epochs. Our pre-experienced models, however, undergo just a single epoch of training. Training on the Alpaca dataset necessitated a max input length of 512. For WizardLM, we opted for a 1024 input length due to hardware constraints while its original model used 2048, which offers an inherent edge to the original model. Another challenge with WizardLM was “AI censure” instances. Taking a leaf from the Vicuna strategy, we filtered these samples, resulting in a streamlined WizardLM subset with 63655 entries. Our data selection methodology was then applied to this subset. Samples with IFD scores higher than 1 will be filtered out before selection. For experiments on llama2-7b and llama2-13b models, we utilize the instruction prompt from Vicuna (Chiang et al., 2023). Thanks to the Flash Attention mechanism (Dao et al., 2022), all models on llama2 use the max length of 2048.

3.3 Evaluation Metric

3.3.1 Pair-wise Comparison

Assessing the instruction-following capabilities of LLMs is challenging. While extensive research is dedicated to creating automated evaluation metrics for LLMs (Chang et al., 2023), human judgment remains unmatched. However, it’s both labor-intensive and potentially influenced by biases. Leveraging the recent advancements in independent LLM evaluations (Zheng et al., 2023; Chiang et al., 2023; Li et al., 2023), we utilize GPT4 and ChatGPT for comparative evaluations. Following (Chen et al., 2023), for each instruction in the test dataset, models that need to be compared are prompted to generate responses respectively. Then an API model, either GPT4 or ChatGPT, assigns scores for their responses. The model is regarded to be better in this dataset only if its answer is preferred by the judging model.

In the evaluation, each model’s response is rated by the judge on a scale from 1 to 10, reflecting attributes like relevance and accuracy. To further address the positional bias (Ko et al., 2020; Wang et al., 2023a), we send the responses of two models to the judge twice with different orders and com-

pare their scores. Thus we define one model to be seen as winning only if it does not lose in both the ordering, specifically:

- **Wins:** outperforms in both or wins in one and ties in the other.
- **Tie:** ties in both or wins in one and loses in the other.
- **Loses:** lags in both or ties in one and loses in the other.

This evaluation serves as the foundation of our experimental outcomes.

3.3.2 Benchmarks

The performances on two recently popular benchmarks for LLMs are also provided: Huggingface Open LLM Leaderboard² and AlpacaEval Leaderboard³. Huggingface Open LLM Leaderboard evaluates LLMs using (Gao et al., 2021), a unified framework to test generative language models on a large number of different evaluation tasks, on 4 key benchmarks including ARC (Clark et al., 2018), HellaSwag (Zellers et al., 2019), MMLU (Hendrycks et al., 2021) and TruthfulQA (Lin et al., 2022). AlpacaEval Leaderboard provides an LLM-based automatic evaluation based on AlpacaFarm (Dubois et al., 2023) evaluation set, in which the model responses are compared with responses of Davinci003 by GPT4.

3.3.3 Human Evaluation

To better illustrate the efficacy of our method, further human evaluation is conducted. Specifically, we randomly sampled 20 instructions from each test set to generate a new random set containing 100 instructions in total. Then 3 human participants are asked to compare the responses generated by the models to be compared. For each comparison, 3 options are given (Win, Tie, and Loss) and the final results are determined by the majority voting of the participants.

4 Experimental Results

4.1 Main Results

In this section, we first present our primary **pair-wise evaluation results** in Figure 2. (a) our model trained with only approximately 5% of the original Alpaca data beats the Alpaca model trained with full data. (b) our model trained with only approximately 10% of the original WizardLM data

²https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard

³https://tatsu-lab.github.io/alpaca_eval




Figure 2: Comparing our models trained on selected data with full data. (a) Comparison between our model with 5% Alpaca data and the official Alpaca model. (b) Comparison between our model with 10% WizardLM data and the reimplemented WizardLM model. Both (a) and (b) use GPT4 as the judge. Each horizontal bar represents a comparison in a specific test set.

beats the reimplemented WizardLM model under the same training configuration which is described in the Implementation Details.

Moreover, we craft subsets containing the top 5%, 10%, 15%, and 20% of the target datasets and train models on these distinct subsets, enabling us to investigate the performance changes. As shown in Figure 3, we draw the overall winning rate changes across the data growth, which is calculated as $(\text{Num}(\text{Win}) - \text{Num}(\text{Lose})) / \text{Num}(\text{All}) + 1$, providing a direct indicator on the comparison with the full-data trained models. A consistent observation across both datasets is that with merely 10% of selectively chosen data, our models manage to exceed the results of models trained on the full dataset. These findings not only highlight the efficiency of our data selection strategy but also underscore the potential of training powerful models with significantly reduced data requirements. By validating our approach on the renowned Alpaca dataset and the more intricate WizardLM dataset, we emphasize the wide applicability and robustness of our proposed method.

The comparison between our cherry models




Figure 3: The winning score changes over data growth comparing our models with corresponding full-data models. The winning score is calculated as $(\text{Num}(\text{Win}) - \text{Num}(\text{Lose})) / \text{Num}(\text{All}) + 1$. The Number of Wins, Losses, and All are calculated across all five test sets we used. When the value is higher than 1.0, it means this model performs better than the comparison. Comparisons in this figure are performed by ChatGPT.

with baseline models on **Huggingface Open LLM Leaderboard** and **AlpacaEval Leaderboard** are presented in Table 6 where we can see our cherry model using 5% Alpaca data outperforms the official Alpaca on both benchmarks, our cherry model using 10% WizardLM data has a closed performance compared with our re-implemented WizardLM. These results further showcase the effectiveness of our automatically selected data.

Moreover, the **human evaluation** results also showcase the usefulness of our method. When comparing the Cherry Alpaca (5%) and the Alpaca (100%), there are 49/100 wins for our cherry alpaca, 25/100 ties, and 26/100 losses. When comparing the Cherry WizardLM (10%) and the re-implemented WizardLM (100%), there are 37/100 wins for our Cherry WizardLM, 32/100 ties, and 31/100 losses.

4.2 Ablation on Cherry Data Selection

4.2.1 Data Randomly Selected

We train various LLaMA-7B models using randomly chosen data and juxtaposed their performance with that of our models, which employed a difficulty ratio. As shown in Figure 4 (labeled as *Random*), models trained on 5%, 10%, or 15% random data consistently underperformed against the official Alpaca model. Notably, with an equivalent amount of data, our model surpasses the performance of models using randomly selected data, underlining our method’s superiority.




Figure 4: The overall winning score changes over data growth comparing models using different data selection strategies with the official Alpaca model. All the comparison in this figure is performed by ChatGPT.

4.2.2 Data with Diversity

In this experiment, we train a series of models only considering the diversity of the data samples. Specifically, we utilize the same method for obtaining diverse samples from the full dataset, in which a k-means algorithm is first implemented, and then data is sampled from each cluster. It is a direct baseline for the situation where only the diversity of data is considered. As illustrated in Figure 4 (labeled as *Diversity*), these models render subpar performance and are similar to the random trained models. This result highlights the necessity of using difficult samples over pure diverse samples.

4.2.3 Data with Low IFD Score

In this experiment, we aim to underscore the efficacy of our proposed IFD score. We train a model using data chosen based on low IFD scores on the pre-experienced model, a direct antithesis to our primary experimental setting. As illustrated in Figure 4 (labeled as *Low IFD score*), models trained using low IFD scores render subpar performance. This observation highlights the prowess of our metric in sifting through high-quality data: a higher score consistently yields superior results compared to the baseline, while a lower score deteriorates the model’s intrinsic performance.

4.2.4 Data with High CA Scores

For this comparison, we juxtapose our model against one trained on data selected by higher Conditioned Answer scores which is **equivalent to the loss or perplexity**, and is a commonly accepted baseline. As Figure 4 (labeled as *High CA score*) elucidates, models in this group trail the official

	Huggingface Open LLM Leaderboard					AlpacaEval
	Average	ARC	HellaSwag	MMLU	TruthfulQA	AlpacaEval
Official Alpaca	50.21	42.65	76.91	41.73	39.55	26.46
Ours (5% Alpaca)	52.06	53.92	79.49	36.51	38.33	34.74
Reimplemented WizardLM*	52.79	53.07	77.44	37.75	42.90	61.99
Ours (10% WizardLM)	51.59	52.90	78.95	33.08	41.41	61.44

Table 1: The comparison of performance on Huggingface Open LLM Leaderboard and AlpacaEval Leaderboard.

Alpaca model significantly. The salient difference between these models and ours rests on the elimination of Direct Answer scores. In models relying solely on CA scores, the underlying comprehension of the pre-trained LLM towards original answer texts isn’t factored in, rendering high CA scores ineffective in gauging the intricate nuances of the instruction following.

4.3 Results on Other Models

In this section, experiments on newer LLaMA2-7B and LLaMA2-13B models are conducted as shown in Table 2. In these experiments, the IFD score of each sample is calculated directly based on the corresponding LLaMA2 pre-trained models by using prompts from Vicuna (Chiang et al., 2023). On both LLaMA2-7B and LLaMA2-13B models, our cherry models trained with much less data outperform the models trained with original full data. These experimental results illustrate the consistent advantages of our method and further verify the generalizability of our method.

5 Cherry Data Characteristics

5.1 Validity of Cherry Data

Furthermore, we orchestrate a human-centric study focused on data selection on the Alpaca data. Here, 200 samples, comprising both 100 samples from the top 5% IFD scores and 100 samples from the least 5% IFD scores, are randomly mixed. Three human evaluators are tasked to evaluate the difficulty of each instruction based on the same six aspects mentioned above. Then they are asked to categorize each sample pair as *Hard* or *Easy*, without the necessity to provide specific scores on the six aspects. The final results are decided by majority voting, which shows that 100 out of 100 samples from the least 5% IFD scores were labeled as *Easy*. 89 of the 100 top 5% samples are branded as *Hard*, and 11 of the 100 top 5% samples are also branded as *Easy*. This rudimentary human study lends further weight to the efficacy of our IFD scores in singling out challenging instances.

5.2 Distributional Characteristics

In this segment, our focus is on understanding the distributional properties of the cherry data within the original dataset. Specifically, we first compute the embedding of each instruction in the Alpaca dataset and employ t-SNE for dimensionality reduction, mapping high-dimensional embeddings to 2D space. The visualized vectors, color-coded based on the top or least 5% difficulty ratios, are showcased in Figure 5. Contrary to conventional beliefs, our cherry data isn’t uniformly scattered. Instead, a palpable demarcation exists between samples of high and low difficulty, challenging prior assumptions that selected data should span the entire instruction spectrum and maximize diversity.

To delve deeper into the distributional intricacies of instruction embeddings, we utilize naive K-means (K=100) for clustering. We home in on 10 representative clusters, half of which displayed a significant overlap with the top 5% samples and the other half with the least 5% samples. Clusters dominated by low IFD score samples are replete with rudimentary tasks like editing punctuation, words, or sentences. In contrast, high IFD score clusters are typified by deeper, more intricate tasks such as storytelling or elucidation of phenomena. We posit that these in-depth tasks are paramount for aligning large language models, compelling them to rearrange and access their intrinsic knowledge repositories. Our methodology lends partial credence to this hypothesis, leaving room for further exploration.

6 Related Work

6.1 Data-driven Instruction Tuning

Previous instruction tuning collections are typically handcrafted or task-related (Khashabi et al., 2020; Ye et al., 2021; Wei et al., 2022; Wang et al., 2022a; Du et al., 2022; Honovich et al., 2023), (Wang et al., 2023b) utilized GPT3 (Brown et al., 2020) to generate 52k distinct instructions which do not directly relate to each task, which paves the way to generating instruction data set by distilling from teacher

	Huggingface Open LLM Leaderboard					AlpacaEval
	Average	ARC	HellaSwag	MMLU	TruthfulQA	AlpacaEval
Alpaca llama2 7b	55.25	54.35	78.65	47.02	40.98	27.75
Ours (5% Alpaca)	55.78	57.94	80.37	44.19	40.62	36.78
Ours (10% Alpaca)	56.31	58.02	80.42	46.64	40.18	-
Ours (15% Alpaca)	56.37	57.42	80.68	46.40	40.95	-
Alpaca llama2 13b	58.78	57.59	81.98	54.05	41.49	35.00
Ours (5% Alpaca)	61.21	62.37	84.00	55.65	42.82	46.82
Ours (10% Alpaca)	61.02	62.97	83.88	55.29	41.93	-
Ours (15% Alpaca)	61.23	62.37	83.48	55.56	43.42	-

Table 2: The comparison of performance on Huggingface Open LLM Leaderboard and AlpacaEval Leaderboard for




Figure 5: Visualization using t-SNE on instruction embeddings from the Alpaca dataset. Red points represent samples with the top 5% IFD scores and Blue points represent samples with the least 5% IFD scores.

models. After the release of Meta LLaMA(Touvron et al., 2023a), the world witnessed a surge of open-sourced instruction tuning datasets and LLMs.

6.2 Instruction Data Selection

Though consensus has been made that “quality is all you need” (Touvron et al., 2023b) for instruction tuning, how to find data of good quality other than human curation is still an under-explored topic. Two recent papers were proposed to bridge this gap and shared similar motivations to us: Instruction Mining (Cao et al., 2023) and ALPAGASUS(Chen et al., 2023). Instruction Mining evaluates a number of different indicators and applies a statistical regression model for selecting data. However, it does not present the performance compared with the model trained with full data and its method is overly complicated since it has to split all data into several bins and fully-trained the models. ALPAGASUS directly utilizes an outside fully-trained strong LLM (ChatGPT) to score each sample and finally selected 9k Alpaca data to beat the official

Alpaca trained on full data. Though effective, it neglects the intrinsic ability of the base model and relies too much on extra models.

6.3 Data Selection in Active Learning

Another topic related to our work is active learning (Cohn et al., 1996), which is a kind of approach where the algorithm selectively queries the oracles to label instances that are most informative. It actively decides which data points they want to learn from, optimizing the learning process when labeling data is costly or time-intensive. Though having slightly different goals with our work, the acquisition functions from active learning can be enlightening to our method. Most previous active learning methods work on text classification tasks (Ein-Dor et al., 2020; Yuan et al., 2020; Schröder et al., 2022; Yu et al., 2023), and recently, active learning on text generation tasks was also studied(Perlitiz et al., 2023), especially on neural machine translation(Zhao et al., 2020; Gupta et al., 2021; Vashistha et al., 2022) and text summarization (Tsvigun et al., 2022). In active learning, the Uncertainty-based approaches share a similar motivation to us, selecting the most informative samples for the model which is the least certain.

7 Conclusion

This study has illuminated the potential of harnessing the innate capabilities of LLMs for selecting high-quality instruction tuning data that fit the model. Through our innovative self-guided approach, LLMs demonstrate the ability to discern and cherry-pick the most pertinent data samples, a concept we’ve aptly termed cherry data. Central to our methodology is the Instruction-Following Difficulty metric, a novel tool adept at gauging the nuanced differences between a model’s autonomous outputs and expected responses. Our findings not only emphasize the importance of data quality over quantity but also underscore the potential for cost-effective and streamlined LLM training.

Limitation

The main limitation of this method is the inconvenience of training the pre-experienced model. The concept of the Instruction-Following Difficulty score proposed by us is simple and effective, while the inconvenient pre-experienced phase makes it hard to directly put our method into usage in real-world scenarios. Though experiments on LLaMA2 models show that calculating IFD scores directly on the base LLaMA2 models also promises a good selection, we believe using the pre-experienced phase is valuable since it equips base models with the basic instruction-following ability, making the calculation of Conditioned Answer Score more reasonable. As a result, we believe the use of the pre-experienced phase could be a tradeoff: From the **Research Viewpoint**, using pre-experienced models is more reasonable and performs better. From the **Real-world Implementation Viewpoint**, directly using the base model is more efficient and at the same time effective as well.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yihan Cao, Yanbin Kang, and Lichao Sun. 2023. [Instruction mining: High-quality instruction data selection for large language models](#).
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. [A survey on evaluation of large language models](#).
- Lichang Chen, Shiyang Li, Jun Yan, Hai Wang, Kalpa Gunaratna, Vikas Yadav, Zheng Tang, Vijay Srivasan, Tianyi Zhou, Heng Huang, and Hongxia Jin. 2023. [Alpagasus: Training a better alpaca with fewer data](#).
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan

- Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#).
- David A. Cohn, Zoubin Ghahramani, and Michael I. Jordan. 1996. [Active learning with statistical models](#). In *NIPS*.
- Tri Dao, Daniel Y. 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*.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. [GLM: General language model pretraining with autoregressive blank infilling](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335, Dublin, Ireland. Association for Computational Linguistics.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Alpaca-farm: A simulation framework for methods that learn from human feedback](#).
- Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. [Active Learning for BERT: An Empirical Study](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7949–7962, Online. Association for Computational Linguistics.
- Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. [A framework for few-shot language model evaluation](#).
- Kamal Gupta, Dhanvanth Boppana, Rejwanul Haque, Asif Ekbal, and Pushpak Bhattacharyya. 2021. [Investigating active learning in interactive neural machine translation](#). In *Proceedings of Machine Translation Summit XVIII: Research Track*, pages 10–22, Virtual. Association for Machine Translation in the Americas.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.

741	Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2023. Unnatural instructions: Tuning language models with (almost) no human labor . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 14409–14428, Toronto, Canada. Association for Computational Linguistics.	797	Yotam Perlitz, Ariel Gera, Michal Shmueli-Scheuer, Dafna Sheinwald, Noam Slonim, and Liat Ein-Dor. 2023. Active learning for natural language generation .	798
742		799		800
743				
744				
745				
746				
747				
748	Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. UNIFIEDQA: Crossing format boundaries with a single QA system . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 1896–1907, Online. Association for Computational Linguistics.	801	Teven Le Scao, Angela Fan, Christopher Akiki, Elizabeth-Jane Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, Francois Yvon, Matthias Galle, Jonathan Tow, Alexander M. Rush, Stella Rose Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurencon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa Etxabe, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris C. Emezue, Christopher Klammer, Colin Leong, Daniel Alexander van Strien, David Ifeoluwa Adedani, Dragomir R. Radev, Eduardo Gonzalez-Ponferrada, Efrat Levkovich, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady ElSahar, Hamza Benyamina, Hieu Trung Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jorg Froberg, Josephine L. Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro von Werra, Leon Weber, Long Phan, Loubna Ben Allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, Mar’ia Grandury, Mario vSavsko, Max Huang, Maximin Coavoux, and Mayank Singh. 2022. Bloom: A 176b-parameter open-access multilingual language model . <i>ArXiv</i> , abs/2211.05100.	802
749		803		804
750		805		806
751		807		808
752		809		810
753		811		812
754		813		814
755	Diederik P. Kingma and Jimmy Ba. 2017. Adam: A method for stochastic optimization .	815		816
756		817		818
757	Miyoung Ko, Jinhyuk Lee, Hyunjae Kim, Gangwoo Kim, and Jaewoo Kang. 2020. Look at the first sentence: Position bias in question answering . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 1109–1121, Online. Association for Computational Linguistics.	819		820
758		821		822
759		823		824
760		825		826
761		827		828
762		829		830
763		831		832
764	Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval .	833		834
765		835		836
766		837		838
767		839		840
768		841		842
769	Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. TruthfulQA: Measuring how models mimic human falsehoods . In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.	843		844
770		845		846
771		847		848
772		849		850
773		851		852
774		853		854
775	S. Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. The flan collection: Designing data and methods for effective instruction tuning . <i>ArXiv</i> , abs/2301.13688.	855		856
776		857		858
777		859		860
778		861		862
779		863		864
780	OpenAI. 2023. Gpt-4 technical report .	865		866
781		867		868
782	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback . In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 27730–27744. Curran Associates, Inc.	869		870
783		871		872
784		873		874
785		875		876
786		877		878
787		879		880
788		881		882
789		883		884
790		885		886
791	Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only .	887		888
792		889		890
793		891		892
794		893		894
795		895		896
796		897		898

857	Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. Llama 2: Open foundation and fine-tuned chat models .	
874	Akim Tsvigun, Ivan Lysenko, Danila Sedashov, Ivan Lazichny, Eldar Damirov, Vladimir Karlov, Artemy Belousov, Leonid Sanochkin, Maxim Panov, Alexander Panchenko, Mikhail Burtsev, and Artem Shelmanov. 2022. Active learning for abstractive text summarization . In <i>Findings of the Association for Computational Linguistics: EMNLP 2022</i> , pages 5128–5152, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	
883	Neeraj Vashistha, Kriti Singh, and Ramakant Shakya. 2022. Active learning for neural machine translation .	
885	Thuy-Trang Vu, Xuanli He, Gholamreza Haffari, and Ehsan Shareghi. 2023. Koala: An index for quantifying overlaps with pre-training corpora .	
888	Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023a. Large language models are not fair evaluators .	
891	Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khoshabi, and Hannaneh Hajishirzi. 2023b. Self-instruct: Aligning language models with self-generated instructions . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.	
899	Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022a. Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	
917	Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, et al. 2022b. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 5085–5109.	917
918		918
919		919
920		920
921		921
922		922
923		923
924		924
925	Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. Finetuned language models are zero-shot learners . In <i>International Conference on Learning Representations</i> .	925
926		926
927		927
928		928
929		929
930	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 .	930
931		931
932		932
933		933
934	Qinyuan Ye, Bill Yuchen Lin, and Xiang Ren. 2021. CrossFit: A few-shot learning challenge for cross-task generalization in NLP . In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 7163–7189, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	934
935		935
936		936
937		937
938		938
939		939
940		940
941	Yue Yu, Rongzhi Zhang, Ran Xu, Jieyu Zhang, Jiaming Shen, and Chao Zhang. 2023. Cold-start data selection for better few-shot language model fine-tuning: A prompt-based uncertainty propagation approach . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 2499–2521, Toronto, Canada. Association for Computational Linguistics.	941
942		942
943		943
944		944
945		945
946		946
947		947
948		948
949	Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. 2020. Cold-start active learning through self-supervised language modeling . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 7935–7948, Online. Association for Computational Linguistics.	949
950		950
951		951
952		952
953		953
954		954
955	Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 4791–4800, Florence, Italy. Association for Computational Linguistics.	955
956		956
957		957
958		958
959		959
960		960
961	Yuekai Zhao, Haoran Zhang, Shuchang Zhou, and Zhihua Zhang. 2020. Active learning approaches to enhancing neural machine translation . In <i>Findings of the Association for Computational Linguistics: EMNLP 2020</i> , pages 1796–1806, Online. Association for Computational Linguistics.	961
962		962
963		963
964		964
965		965
966		966
967	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena .	967
968		968
969		969
970		970
971		971
972	Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu,	972
973		973

974 Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis,
975 Luke Zettlemoyer, and Omer Levy. 2023. [Lima: Less](#)
976 [is more for alignment](#).

A Ablation on Pre-Experienced Data Selection

A.1 Number of Pre-Experience Data

Following the findings from LIMA that 1000 high-quality samples are enough to train a reasonably good model, we set the amount of data used for our pre-experienced model as 1000. However, it is still under-investigated how many data samples are required to equip the model with basic instruction-following ability. Thus this section analyzes the significance of employing experience-augmented models and how the number of pre-experienced data affects the final performance of our cherry models. For these comparisons, we conduct the experiments where 0, 100, 300, and 500 pre-experienced samples are utilized to train the pre-experienced models for 1 epoch. Using 0 pre-experienced samples represents direct use of the initial raw model as the pre-experienced model. We calculate the IFD scores from these different pre-experienced models and select the top 5%, 10%, and 15% samples for training while keeping other experimental conditions constant.

As depicted in Figure 6 *0 Pre-Experienced Samples*, when no pre-experienced samples are utilized, the corresponding cherry models have the least performance. These results underline the indispensability of an experience-augmented model equipped with foundational instruction-following capabilities. Moreover, even in the absence of a pre-experienced model, our IFD score remains effective in identifying optimal training data as it outperforms the Alpaca model when using 10% of the data. When 100 samples are utilized as shown in *100 Pre-Experienced Samples*, the corresponding cherry models are slightly better than no samples used but with a similar trend, which indicates that 100 samples are not enough for the model to acquire the basic instruction-following ability. When adding the number of pre-experienced samples to 300, a distinct performance gain is discovered, and further addition of samples does not make the performance of corresponding cherry models better. We hypothesize this is when the model is equipped with the basic instruction-following capability and thus can better illustrate the instruction-following score of each instance.

A.2 Distribution of Pre-Experience Data

To better illustrate what kinds of data are required in the pre-experience process, extensive experi-

ments are conducted where we selected 1000 pre-experienced samples by calculating the IFD scores based on the initial raw model and utilized these samples to train the pre-experienced model and further get the cherry samples and the cherry model. Different from our main method where the pre-experienced samples are selected based on the diversity of instruction distribution, this experiment is used to figure out what is the better strategy for the pre-experienced model, considering the diversity of instructions or difficulty of instructions. Another baseline method is using randomly selected data for the training of pre-experienced models. The performance of using 5%, 10%, and 15% cherry data is shown in Table 3 compared with the Alpaca model. Comparing random selection or considering embedding distributions or instruction difficulties, they all surpass the Alpaca model and are comparable to each other, indicating the effectiveness of both strategies and further proving that our IFD metric is robust across different pre-experienced models. This experiment further illustrates that what matters is this pre-experience process, rather than the sampling strategies for this process.

	5%	10%	15%	100%
Random (1000)	1.007	1.047	1.077	1
Distribution (1000)	1.050	1.097	1.064	1
Difficulty (1000)	1.057	1.072	1.096	1

Table 3: The overall winning score changes over the data growth comparing models with different strategies of selecting pre-experienced samples with the official Alpaca model. All the comparison in this table is performed by ChatGPT.




Figure 6: The overall winning score changes over data growth comparing models with different numbers of pre-experienced samples with the official Alpaca model. All the comparison in this figure is performed by ChatGPT.

B Performance across Sub-Categories

To evaluate the performance variations of our model, we scrutinize the capabilities across diverse instruction tasks. To accomplish this, we compare the response of our cherry models, trained with 5% Alpaca data and 10% WizardLM data, to their corresponding comparing models, the official Alpaca and the reimplemented WizardLM across sub-categories in the WizardLM and Vicuna test sets, as displayed in Table 4 and Table 5.

Our cherry model trained on Alpaca data exhibits superior or at least comparable performance to the official Alpaca model on most of the subcategories in the Vicuna and WizardLM test sets. Notably, exceptions are observed in the Math and Coding categories, corroborating the observations made by (Chen et al., 2023). We surmise that the base 7B models inherently perform sub-optimally on these two tasks, necessitating a greater volume of data samples to effectively learn the alignment.

Our cherry model trained on WizardLM data also has a better or comparable performance compared with the reimplemented WizardLM model on most of the subcategories. Specifically, Our model underperforms in Math, Code, Complex Format, and Counterfactual. The main reason our model loses in these categories is the abundance of training data for these categories in the original dataset and the supreme abilities of the original WizardLM in these tasks, which is mentioned in (Xu et al., 2023). As a consequence, when we reduce the number of data used, our model can not be trained on these data-needed categories as much as the original model, thus leading to a relatively incomparable performance.

	Math	Coding	Writing	Generic	Knowledge	Roleplay	Common-sense	Fermi	Counterfactual
Alpaca	0.33 / 0.33 / 0.33	0.14 / 0.43 / 0.43	0.60 / 0.20 / 0.20	0.60 / 0.10 / 0.30	0.60 / 0.20 / 0.20	0.60 / 0.10 / 0.30	0.60 / 0.20 / 0.30	0.50 / 0.30 / 0.20	0.70 / 0.10 / 0.20
WizardLM	0.00 / 0.33 / 0.67	0.00 / 0.29 / 0.71	0.60 / 0.20 / 0.20	0.40 / 0.40 / 0.20	0.70 / 0.30 / 0.00	0.40 / 0.30 / 0.30	0.70 / 0.30 / 0.00	0.50 / 0.30 / 0.20	0.20 / 0.20 / 0.60

Table 4: The comparison between our cherry models and their corresponding comparing models on sub-categories in Vicuna test sets, using GPT4 as the judge.

	Math	CodeGeneration	Writing	Computer	Reasoning	ComplexFormat	CodeDebug	CommonSense	Counterfactual
Alpaca	0.21 / 0.37 / 0.42	0.28 / 0.33 / 0.39	0.56 / 0.17 / 0.28	0.40 / 0.33 / 0.27	0.31 / 0.54 / 0.15	0.50 / 0.25 / 0.25	0.50 / 0.50 / 0.00	0.55 / 0.11 / 0.33	1.00 / 0.00 / 0.00
WizardLM	0.42 / 0.37 / 0.21	0.33 / 0.28 / 0.39	0.50 / 0.44 / 0.06	0.33 / 0.40 / 0.27	0.38 / 0.23 / 0.38	0.25 / 0.25 / 0.50	0.40 / 0.40 / 0.20	0.56 / 0.44 / 0.00	0.00 / 0.38 / 0.62
	Multilingual	Roleplay	Biology	Technology	Ethics	TruthfulQA	Sport	Law	Medicine
Alpaca	0.29 / 0.29 / 0.42	0.67 / 0.17 / 0.17	0.50 / 0.00 / 0.50	0.83 / 0.17 / 0.00	0.67 / 0.00 / 0.33	0.60 / 0.00 / 0.40	1.00 / 0.00 / 0.00	0.40 / 0.00 / 0.60	0.80 / 0.00 / 0.20
WizardLM	0.14 / 0.71 / 0.14	0.33 / 0.33 / 0.33	0.17 / 0.50 / 0.33	0.50 / 0.50 / 0.00	0.17 / 0.83 / 0.00	0.80 / 0.20 / 0.00	0.20 / 0.60 / 0.20	0.20 / 0.60 / 0.20	0.80 / 0.00 / 0.20

Table 5: The comparison between our cherry models and their corresponding comparing models on sub-categories in WizardLM test sets, using GPT4 as the judge.

C Results with Official WizardLM

In this section, we provide the results of using 40% of the WizardLM data to have a comparable performance with the official WizardLM model in a relatively unfair setting. The official WizardLM is uncensored and trained with the max token size of 2048, while our model is trained with the max token size of 1024, representing an inherent disadvantage of our model. However, even with this situation, our model can still reach a comparable performance with the official WizardLM model, inferring the effectiveness of our method.




Figure 7: Comparing our models trained on cherry data with official WizardLM trained on full data using GPT4 as the judge.

	Huggingface Open LLM Leaderboard					AlpacaEval
	Average	ARC	HellaSwag	MMLU	TruthfulQA	AlpacaEval
Official WizardLM	54.18	51.60	77.70	42.70	44.70	67.64
Ours (40% WizardLM)	52.83	53.07	77.79	35.29	45.17	65.09

Table 6: The comparison of performance on Huggingface Open LLM Leaderboard and AlpacaEval Leaderboard.

D Validity of Cherry Data

Our goal in this section is to determine if the data selected based on the IFD scores aligns with known characteristics of high-quality training data. To this end, we randomly sample 100 instances from data with the top 5% scores and the least 5% scores. Utilizing ChatGPT, we evaluate each instruction on six aspects: Scope, Complexity, Clarity, Depth, Simplicity, and Knowledge Required. The results are depicted in Figure 8. Data with a higher IFD score generally scored higher in Scope, Complexity, Depth, and Knowledge Required, but lower in Clarity and Simplicity. Simplicity, in particular, have the most pronounced discrepancy. This lends credence to our assertion that our IFD scores aptly gauge instruction complexity. Consequently, our method gravitates towards selecting more intricate samples.




Figure 8: The comparison between data instances with top 5% and least 5% IFD scores from Alpaca data. We prompt ChatGPT to score the instruction of each data instance with respect of Scope, Complexity, Clarity, Depth, Simplicity, and Knowledge Required.

As mentioned in the previous section, we try to evaluate each instruction into six aspects, Scope, Complexity, Clarity, Depth, Simplicity, and Knowledge Required. We define these aspects as follows:

- **Scope:** The instruction encompasses the breadth and range of actions or information necessary for successful completion.
- **Complexity:** The instruction integrates multiple steps or concepts that require careful attention and understanding.
- **Clarity:** The instruction is articulated straightforwardly, ensuring it's easily understood

without ambiguity.

- **Depth:** The instruction provides thorough details and nuances, ensuring a comprehensive understanding of the task at hand.
- **Simplicity:** While thorough, the instruction avoids unnecessary jargon or convolutions, making it accessible and easy to follow.
- **Knowledge Required:** The instruction acknowledges and, if necessary, provides the foundational knowledge or context the user needs for successful execution.

From the previous Figure 8, we can see samples selected with top IFD scores have larger scores in the aspects that reflect the difficulty of instruction, including Scope, Complexity, Depth, and Knowledge Required. These samples only underscore samples with the lowest IFD scores on the aspect of Clarity and Simplicity. This experiment detailedly illustrates the difference between samples with high or low IFD scores and verifies the effectiveness of our method in measuring the difficulty of an instruction.

E Cherry Example Analysis

To illustrate the implications of our findings and demonstrate the characteristics of the data selected by our method, we provide several examples in Figure 9.

The first positive example presents the situation that both the direct answer score (DA) and the conditioned answer score (CA) are relatively high. In this situation, the high DA means that it is hard for the initial pre-trained LLM to generate this poem, and the high CA means given the instruction does not make the generation of this poem much easier. So it is valuable for LLM to learn this sample. The second positive example presents the situation that both the CA score and DA score are relatively low. The low DA score means that LLM has learned this knowledge it is easy for LLM to generate this sentence. However, providing the corresponding instruction does not change the situation much, indicating the poor ability to follow this instruction.

The first negative example presents a situation where the response is too short. Due to the intrinsic nature of next token prediction that longer texts tend to have lower perplexity, the DA score is relatively high for the response that is too short and thus causes the IFD Score large, which we believe is a good feature of our method. The second negative example presents a situation where the DA score and CA score are relatively small. In this example, the response is quoted from a book that LLM must have read, thus as a known knowledge, it is easy for LLM to reproduce this sentence. However, with an instruction included, the CA score becomes even much lower, indicating LLM has gained quite a good ability in following this instruction. The third example presents the most common situation, where the instruction is simply not difficult enough.

Positive Examples	Negative Examples
<p>[Instruction] Create a poem that tells the story of a struggle against an unseen force.</p> <p>[Input]</p> <p>[Response]</p> <p>A distant enemy looms on the horizon, Selves of the light fight despite fear and fatigue. Innumerable obstacles barring their mission, Though seeming to weaken, strengths within achieve.</p> <p>Roots of bravery sustained their cause, Though many obstacles pushed and pulled. Heart's noise drowned out every pause, Escaping each struggle stronger and full.</p> <p>The reward of their labors thought clear and light, The force of the shadows still kept them at bay. Yet further pushed and tired through the night, Triumph steadily came in a growing display.</p> <p>Peace of their struggle found in the morning's glow, Rising higher than any unseen force's ebb and flow.</p> <p>[Direct Answer score (DA)]3.970 [Conditioned Answer score (CA)]3.337 [IFD score]0.928</p>	<p>[Instruction] What emotion is expressed in this tweet?</p> <p>[Input] I'm so frustrated right now!!</p> <p>[Response] Frustration</p> <p>[Direct Answer score (DA)] 6.593 [Conditioned Answer score (CA)] 0.601 [IFD score] 0.091</p>
<p>[Instruction] Give a brief description of the coronavirus.</p> <p>[Input]</p> <p>[Response]</p> <p>Coronavirus is a type of virus that can cause illnesses ranging from the common cold to severe diseases such as Middle East Respiratory Syndrome (MERS) and Severe Acute Respiratory Syndrome (SARS). The virus is named for the crown-like spikes on its surface. It spreads primarily through close contact with an infected person, through small droplets from the nose or mouth after coughing or sneezing. Symptoms can include fever, tiredness, and dry cough. In severe cases, it can cause difficulty breathing, pneumonia, kidney failure, and even death.</p> <p>[Direct Answer score (DA)]0.761 [Conditioned Answer score (CA)]0.696 [IFD score]0.914</p>	<p>[Instruction] Identify a famous quote from the given book</p> <p>[Input] Book: Pride and Prejudice by Jane Austen</p> <p>[Response] "It is a truth universally acknowledged, that a single man in possession of a good fortune, must be in want of a wife."</p> <p>[Direct Answer score (DA)] 0.497 [Conditioned Answer score] 0.026 [IFD score] 0.053</p>
	<p>[Instruction] Generate a report summarizing a conversation about car maintenance.</p> <p>[Input] Joe and Lisa were discussing the importance of regular car maintenance such as oil changes, tire rotations, and brake inspections. Joe believes that regular car maintenance is essential for keeping the car operating properly and avoiding major repairs. Lisa believes that some car maintenance procedures are unnecessary and are a waste of money.</p> <p>[Response]</p> <p>Joe and Lisa discussed the importance of regular car maintenance such as oil changes, tire rotations, and brake inspections. Joe believes that regular car maintenance is essential for keeping the car operating properly and avoiding major repairs while Lisa believes that some procedures are unnecessary and are a waste of money. It is agreed that regular oil changes and tire rotations are important to maintain the health of the car but that some other maintenance procedures can be skipped depending on the age and condition of the car.</p> <p>[Direct Answer score (DA)] 1.667 [Conditioned Answer score] 0.599 [IFD score] 0.359</p>

Figure 9: Example instances selected by our proposed method. 2 positive examples and 3 negative examples are provided.

F Additional Discussion

F.1 Fully-trained Model as Pre-Experienced Model?

In our method, efforts are conducted to keep the pre-experience process as simple as possible, however, there still exists a question of whether the fully-trained model can be the pre-experienced model for selecting the cherry samples. To better illustrate this question, the fully-trained Alpaca model is utilized as the pre-experienced model for selecting the cherry data, 5%, 10%, and 15% of the cherry data are selected and the corresponding cherry models are trained. The performances are shown in Table 7, in which the models with the fully-trained Alpaca hardly surpass the Alpaca with fewer data and our models. This experiment proves that the fully-trained model is not appropriate in selecting samples for the initial raw model, which is caused by the overly distribution gap between the fully-trained models and raw models.

	5%	10%	15%	100%
Ours	1.050	1.097	1.064	1
Fully-trained Alpaca	0.968	0.999	1.005	1

Table 7: The overall winning score changes over the data growth comparing models with fully-trained Alpaca as the pre-experienced model with the official Alpaca model. All the comparison in this table is performed by ChatGPT.

F.2 How Many Cherry Samples are Required?

While extensive experiments with our method on Alpaca and WizardLM prove the effectiveness of our method in selecting high-quality samples from the original target dataset automatically, it is still under-exploring how much data is optimal. Unlike (Chen et al., 2023) in which the scores of target samples are scarce, the dense scores from our method provide better flexibility in deciding how much data you can use. However, this flexibility is also a curse that makes it hard to conclude the optimal number of data to select, which is influenced by various factors including the absolute values of the IFD scores, the distribution of hard examples, and the number of data in original datasets. However, from our empirical study, we think selecting samples with the top 10% IFD scores would be a safe and reasonable choice.

G Prompt for Evaluation

In this section, we provide the detailed prompt we used for evaluating the performance of two responses for the same instruction as shown in Figure 10.

Prompt for Performance Evaluation

System Prompt

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

User Prompt

[Question]

Question

[The Start of Assistant 2's Answer]

Answer 2

[The End of Assistant 2's Answer]

[The Start of Assistant 2's Answer]

Answer 2

[The End of Assistant 2's Answer]

We would like to request your feedback on the performance of two AI assistants in response to the user question displayed above.

Please rate the helpfulness, relevance, accuracy, level of details of their responses. Each assistant receives an overall score on a scale of 1 to 10, where a higher score indicates better overall performance. Please first output a single line containing only two values indicating the scores for Assistant 1 and 2, respectively. The two scores are separated by a space. In the subsequent line, please provide a comprehensive explanation of your evaluation, avoiding any potential bias and ensuring that the order in which the responses were presented does not affect your judgment.

Figure 10: The prompt we used to request ChatGPT or GPT4 to evaluate the responses.

H Detailed Main Comparison

H.1 Comparison with the Official Alpaca

As shown in Figure 11, we present the detailed comparison between our cherry models with the official Alpaca (7B) model across different test set with different percentage of cherry data, from 5% to 15%, using ChatGPT as the judge. Starting from 5% of the full data, our cherry models outperform the official Alpaca model in all these data scales.

H.2 Comparison with the Reimplemented WizardLM

As shown in Figure 12, we present the detailed comparison between our cherry models with the reimplemented WizardLM (7B) model across different test set with different percentage of cherry data, from 5% to 15%, using ChatGPT as the judge. Our cherry models begin outperforming the reimplemented WizardLM from the scale of 10% of the data.

H.3 Comparison with the Official WizardLM

As shown in Figure 13, we show the detailed comparison between our cherry models with the reimplemented WizardLM (7B) model across different test set with different percentage of cherry data, from 5% to 40%, using ChatGPT as the judge. When compared with the official WizardLM data, our cherry model achieves a comparable performance when using 40% of the WizardLM data, which is positive considering the inherent disadvantage of our training configuration.




Figure 11: Comparing our cherry models with the official Alpaca model from 5% to 15% of the data using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.




Figure 12: Comparing our cherry models with the reimplemented WizardLM model from 5% to 15% of the data using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.

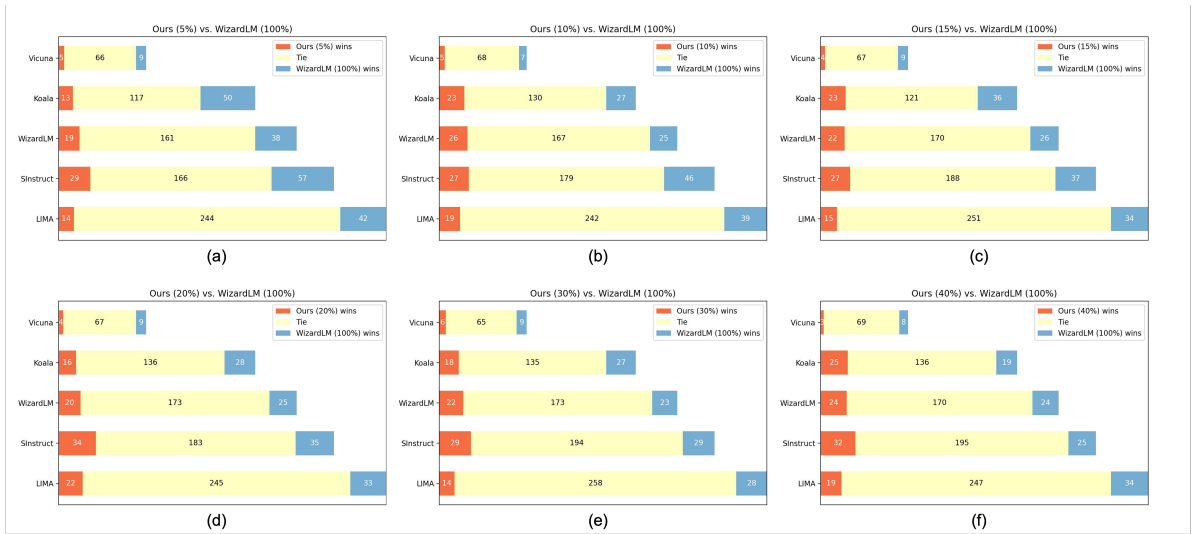


Figure 13: Comparing our cherry models with the official WizardLM model from 5% to 40% of the data using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.

I Detailed Ablation Comparison

I.1 Data Randomly Selected

As shown in Figure 14(a)(b)(c), we show the detailed comparison between the models trained with randomly selected data with our cherry models across different test set with different percentage of data, from 5% to 15%, using ChatGPT as the judge. From 5% to 15% of the data, our cherry models consistently outperform the random models.

I.2 Data with Low IFD Score

As shown in Figure 15, we show the detailed comparison between the models trained with data selected with low IFD scores with our cherry models across different test set with different percentage of data, from 5% to 15%, using ChatGPT as the judge. From 5% to 15% of the data, our cherry models consistently have better performances.

I.3 Data with High CA Scores

As shown in Figure 16, we show the detailed comparison between the models trained with data selected with high conditioned answer scores with our cherry models across different test set with different percentage of data, from 5% to 15%, using ChatGPT as the judge. From 5% to 15% of the data, our cherry models consistently have better performances.

I.4 Number of Pre-Experienced Data

Figure 17 shows the comparisons when different numbers of pre-experienced samples are utilized to train the pre-experienced model.

I.5 Distribution of Pre-Experience Data

Figure 18 shows the comparisons when IFD scores are used as the strategy to select pre-experienced data to train the pre-experienced model.

I.6 Fully-trained Model as Pre-Experienced Models

Figure 19 shows the detailed comparisons when the fully-trained official Alpaca is utilized as the pre-experienced model for selecting cherry data.




Figure 14: Comparing our cherry models with models utilizing randomly selected data from 5% to 15%, using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.




Figure 15: Comparing our cherry models with models trained with data selected with low IFD score from 5% to 15%, using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.




Figure 16: Comparing our cherry models with models trained with data selected with high conditioned answer scores from 5% to 15%, using ChatGPT as the judge. Each horizontal bar represents a comparison in a specific test set.



Figure 17: Comparing our cherry models with different numbers of pre-experienced samples with the Alpca model from 5% to 15%, using ChatGPT as the judge. (a)(b)(c) and (d) represent the situation where 0, 100, 300 and 500 pre-experienced samples are utilized respectively.




Figure 18: Comparing models whose pre-experienced data are selected by considering the difficulty of instructions with the Alpca model from 5% to 15%, using ChatGPT as the judge.




Figure 19: Comparing models with fully-trained Alpaca as the pre-experienced model with the Alpca model from 5% to 15%, using ChatGPT as the judge.

J More Examples

In this section, some positive examples with top 5% IFD scores in the Alpaca dataset are presented in Figure 20 and 21. Negative examples with the least 5% IFD scores are presented in Figure 22.

Positive Example

[Instruction] Design a wellness plan for the given audience

[Input] Expectant Mothers

[Response]

This wellness plan for expectant mothers should include strategies for physical, emotional and mental health. Physical health should include exercise, a healthy diet and adequate rest. Expectant mothers should strive for at least 30 minutes of moderate exercise such as walking or swimming every day. A nutritious diet should include a variety of fruits, vegetables and proteins. Adequate sleep is essential for the body and mind to recharge, so make sure to get seven to eight hours per night.

Stress management is important for emotional and mental health. This can include mindfulness activities, journaling or talking to a friend. Participating in activities like yoga, meditation or guided visualization can also be beneficial. Make sure to take time every day to treat yourself, whether it be by getting a massage or simply taking a long bubble bath.

Make sure to take prenatal vitamins as prescribed by your doctor and attend all of your obstetrician checkups. It's also important to stay hydrated by drinking plenty of water. Schedule regular check-ins with your doctor to ensure the health of both you and your baby.

[Instruction] Identify the types of bias in machine learning

[Input]

[Response]

The types of bias in machine learning include sampling bias, data bias, selection bias, measurement bias, model bias and inference bias. Sampling bias occurs when the dataset is not representative of the population; data bias happens when data is skewed or incomplete; selection bias arises when data is chosen in a non-random way; measurement bias is caused by incorrect labels or noisy data; model bias arises when the model is overly reliant on specific features or data; and inference bias is caused by the model being too simplistic.

[Instruction] Generate a story about a journey in space.

[Input]

[Response]

It was a beautiful day on Earth, but Andrew was not content to stay on land. He wanted to explore and experience the vastness of space. So, on that fateful day, he set off on an intergalactic voyage to the farthest reaches of the stars.

...

Exploring this new universe felt like a dream, and Andrew found himself discovering more wonders than he had ever imagined. One of his greatest adventures was traveling to a star that had never been seen by anyone before him.

Years passed, and Andrew eventually returned to Earth, full of stories and memories of his travels. He had learned a great deal about the universe and made many friends along the way. His journey had taught him that space is full of boundless possibilities and wonders, and he will never forget his time in the stars.

[Instruction] Provide an analysis of the impact of carbohydrates on human health.

[Input]

[Response]

Carbohydrates are an important part of a balanced human diet and have been linked to a variety of health benefits when consumed in appropriate amounts. Carbohydrates are the main source of energy for the body and are important for the brain, organs, and muscles to function properly. ... Choosing complex carbohydrates, such as fruits and vegetables, whole grains, and legumes, instead of simple carbohydrates that are found in processed and refined foods, can help to provide essential vitamins, minerals and other nutrients. Ultimately, carbohydrates are a necessary part of a healthy diet and should be consumed in moderation and in a balanced way.

Figure 20: Positive examples sampled from the alpaca dataset with top 5% IFD scores.

Positive Example

[Instruction] Outline the steps a startup would need to take to increase brand awareness

[Input]

[Response]

Steps to Increase Brand Awareness for a Startup:

1. Identify target customers: Understand the target demographic and research their needs and interests.
2. Develop a content strategy: Use content to engage, inform or entertain customers in order to establish a connection with them.
3. Share the story: Create a website and create social media platforms to share the story, mission or purpose of the startup with customers.
4. Make it interactive: Encourage audience participation via competitions, surveys, or polls.
5. Track and analyze data: Measure the response of customers and use the insights to inform any changes or adjustments."

[Instruction] Create a poem that is related to the season of Fall.

[Input]

[Response]

The leaves are falling, leaves of autumn gold,
A gentle breeze, a chill that blows,
The sun that slowly sets, before night comes to fold,
A season of beauty, with dreams to behold.
The days grow shorter, the sky turns gray,
The land it sighs as summer slips away,
Orange, yellow, and crimson begin to display,
Bringing in a vibrant beauty, during this time of day.
The winds blow cold, as Frost begins to bite,
The air whispers, chill and full of fright,
It is the season of change, and so we all must write,
The story of our lives, in the midst of this night."

[Instruction] What are the major components of the Java programming language?

[Input]

[Response]

The major components of the Java programming language are the Java Virtual Machine (JVM), the Java Runtime Environment (JRE), the Java Platform, the Java Development Kit (JDK), and the Java Libraries. These components provide core components such as threading, garbage collection, memory management and reflection, as well as APIs such as the I/O package, networking, graphics, and more. The JDK is the most fundamental component and is required to compile Java programs. It contains a compiler, interpreter, and other tools needed to create Java applications. The JRE, on the other hand, contains the necessary components to execute the Java programs. The JVM takes the compiled Java programs and executes them. Lastly, the Java Libraries are a vast collection of software packages written in Java that provide the necessary components for programmers to develop applications.

[Instruction] Given a historical event, narrate an account of it in one paragraph.

[Input] The sinking of the Titanic

[Response]

The sinking of the Titanic was one of the most devastating maritime disasters in history. On April 15th, 1912, the illustrious luxury liner set sail on its maiden voyage from Southampton, England, bound for New York City. Despite being touted as the unsinkable ship, the Titanic struck an iceberg and sunk in the early morning hours of the 15th. It took over two hours for the ship to completely disappear beneath the surface of the Atlantic Ocean. Over 1,500 souls were lost, and the disaster shocked the world. To this day, the tragedy of the Titanic serves as a somber reminder of the power of the sea and the fragility of human life.

Figure 21: Positive examples sampled from the alpaca dataset with top 5% IFD scores.

Negative Example

[Instruction] What are the three primary colors?

[Input]

[Response]

The three primary colors are red, blue, and yellow.

[Instruction] Determine if the following sentence is in the simple or the perfect tense: I have celebrated my birthday.

[Input]

[Response]

The sentence is in the perfect tense.

[Instruction] Evaluate this sentence for spelling and grammar mistakes

[Input] He finished his meal and left the resturant

[Response]

He finished his meal and left the restaurant.

[Instruction] Rewrite the following sentence using active voice.

[Input] The news report was read by the captain.

[Response]

The captain read the news report.

[Instruction] Based on the description, determine if the following text is positive or negative sentiment.

[Input] It was a disastrous event.

[Response]

Negative sentiment

[Instruction] What is the force on a 1 kg mass due to the gravitational force?

[Input]

[Response]

The force on a 1 kg mass due to the gravitational force is 9.8 N.

[Instruction] Transform the following sentence into the passive voice

[Input] I bought a book

[Response]

A book was bought by me.

[Instruction] Edit the given text so it reads correctly in terms of grammar, sentence structure and spelling.

[Input] Me and my friend went to the cinema last niught and we saw a realy great movei.

[Response]

My friend and I went to the cinema last night and we saw a really great movie.

[Instruction] Rewrite the following sentence in a more concise way: "The dog that I saw in the park, it was black in color."

[Input] The dog that I saw in the park, it was black in color.

[Response]

I saw a black dog in the park.

[Instruction] Categorize the following text into one of the two sentiment classes: positive or negative.

[Input] I didn't like the product very much because it was too slow and not user friendly.

[Response]

Negative.

[Instruction] What type of capital did the company have?

[Input] The company received an equity funding of \$5 million.

[Response]

The company had equity capital.

[Instruction] Name the three main mountain ranges in the United States.

[Input]

[Response]

The three main mountain ranges in the United States are the Rocky Mountains, the Appalachian Mountains, and the Sierra Nevada Mountains.

[Instruction] Edit the following sentence so it uses the word 'licentious': Bill and Ted's behavior was inappropriate.

[Input]

[Response]

Bill and Ted's licentious behavior was inappropriate.

Figure 22: Negative examples sampled from the alpaca dataset with the least 5% IFD scores.