NEURON-AWARE DATA SELECTION IN INSTRUCTION TUNING FOR LARGE LANGUAGE MODELS

Anonymous authors

000

001

002 003 004

010 011

012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

033

035

037

040

041

042

043

044

046 047

048

051

052

Paper under double-blind review

ABSTRACT

Instruction Tuning (IT) has been proven to be an effective approach to unlock the powerful capabilities of large language models (LLMs). Recent studies indicate that excessive IT data can degrade LLMs performance, while carefully selecting a small subset of high-quality IT data can significantly enhance their capabilities. Therefore, identifying the most efficient subset data from the IT dataset to effectively develop either specific or general abilities in LLMs has become a critical challenge. To address this, we propose a novel and efficient framework called NAIT. NAIT evaluates the impact of IT data on LLMs performance by analyzing the similarity of neuron activation patterns between the IT dataset and the target domain capability. Specifically, NAIT captures neuron activation patterns from in-domain datasets of target domain capabilities to construct reusable and transferable neuron activation features. It then evaluates and selects optimal samples based on the similarity between candidate samples and the expected activation features of the target capabilities. Experimental results show that training on the 10% Alpaca-GPT4 IT data subset selected by NAIT consistently outperforms methods that rely on external advanced models or uncertainty-based features across various tasks. Our findings also reveal the transferability of neuron activation features across different capabilities of LLMs. In particular, IT data with more logical reasoning and programmatic features possesses strong general transferability, enabling models to develop stronger capabilities across multiple tasks, while a stable core subset of data is sufficient to consistently activate fundamental model capabilities and universally improve performance across diverse tasks.

1 Introduction

IT of LLMs has become a foundational technique for activating LLMs instruction and knowledge capabilities (Ouyang et al., 2022). Previous studies have shown that excessive IT data can degrade LLMs performance, while selecting a small amount of high-quality IT data can significantly improve model performance. For instance, LIMA (Zhou et al., 2023) achieved impressive results using only lk IT data. However, a major challenge remains as current approaches lack interpretability in identifying "high-quality" data and fail to enhance the specific (one or more) target domain capabilities of LLMs in an open dataset (Wu et al., 2023; Qin et al., 2024; Wang et al., 2024a). Moreover, existing SOTA IT data selection methods like Instruction Mining (Cao et al., 2023), AlpaGasus (Chen et al., 2024) and SelectIT (Liu et al., 2024b) often rely on surface-level features, external models and data for scoring, or the model's uncertainty, which are computationally expensive. This limits the scalability of these methods and their broader application to large-scale data.

Inspired by previous works on neuron activations and capabilities of LLMs (Voita et al., 2024; Xu et al., 2024c), we address these concerns in NAIT (Neuronal Activation-based efficient IT data selection framework), a novel approach designed to evaluate data quality and enhance specific domain capabilities of LLMs by analyzing neuronal activation patterns related to the target domain capabilities of the LLM. The NAIT framework is built on the core hypothesis that when an LLM processes a given sample, the closer the model's resulting neuronal activation pattern aligns with the activation pattern of the target capability, the more effectively that sample enhances the LLM's performance in that specific domain.

056

060 061

062 063 064

065

066

067

068

069

071

072

073

074

075

076

077

079

081

082

083

084

085

087

090

091

092

094

096

098

099

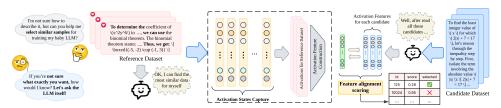
101

102 103

104 105

106

107



(a) Neuron Activation Feature Extraction for Target Capability.

(b) Activation Feature-Guided Data Selection.

Figure 1: **Overall framework of NAIT.** First, we capture the neuron activation of the LLM using in-domain data that we want the model to learn. Next, we construct an activation feature through a dimensionality reduction method. Finally, we evaluate the feature alignment score between this activation feature and the model's activation on each candidate dataset to guide data selection.

NAIT operates in two main stages. First, we extract neuron activation features corresponding to the LLM's target capability. This process begins by providing a small set of in-domain examples that exemplify the specific capability of interest. As the LLM processes these examples, its neuron activation states are recorded to identify reusable activation features linked to the desired capability. These features serve as a representation of the LLM's response and neural activation preferences for the target tasks, forming the foundation for the subsequent data selection phase. Next, we perform capability-driven data selection by identifying and selecting IT data that can further activate and enhance the model's relevant capabilities. This is achieved by comparing the neuron activation features of candidate data within an open and diverse dataset (e.g., Alpaca-GPT4) against the in-domain activation features. By prioritizing data that closely matches the neuron activation patterns of the target capability, NAIT ensures that the selected data is both relevant and effective for improving the model's specialized abilities. In general, NAIT exhibits the following properties: (1) Efficiency (§4.2 & §5.1 & §5.2 & Appendix F): NAIT achieves strong performance with only a small number of in-domain examples and IT data, eliminating the need for external models. It also boasts low computational cost and minimal time requirements. (2) Robustness (§5.1): The framework demonstrates excellent adaptability across different base models, IT datasets, and domain-specific datasets. (3) Interpretability (§5.3): Qualitative analysis shows that NAIT unveils the logical reasoning and programmatic features that possess strong general transferability. Additionally, NAIT identifies a subset of core IT data that remains stable across different task-specific activation features.

Experimental results demonstrate that training LLaMA-2-7b on only 10% of the Alpaca-GPT4 IT dataset selected using NAIT yields an average performance improvement of 3.00% across five tasks compared to training on the entire dataset using IT. Further experiments show that NAIT consistently improves performance across different numbers of in-domain data, base models, diverse IT datasets and selection strategies, highlighting NAIT's strong robustness. The method based on neuron activation features also shows strong interpretability. Further analysis reveals that NAIT tends to select IT data with more logical reasoning and programmatic neuron activation features. Such data can effectively enhance the model's reasoning ability and general ability across various downstream tasks. The key contributions of this paper are as follows:

- Proposing NAIT, the first instruction data selection framework based on neuron activation patterns, introducing a new paradigm for the targeted development of model capabilities.
- Revealing the correlation between neuron activation features and the retention of fundamental model capabilities, while exploring the transferability of these features across different downstream tasks.
- Open-sourcing a cross-task neuron feature library and the Alpaca-NAIT dataset, a high-quality IT dataset curated from the Alpaca-GPT4 dataset using our proposed NAIT.

2 RELATED WORK

IT plays a crucial role in bridging general pretraining and task-specific alignment, as summarized in Appendix C. It enables LLMs to follow instructions, activate latent abilities, and significantly enhance downstream task performance.

Table 1: **Comparison of our method and existing methods.** ✓ indicates the presence of a feature, ✓ indicates its absence, and ○ denotes partial support.

Method	Feature Source	Externally-Independent	Targeted Ability	Cost-Effective	Interpretability
LIMA (Zhou et al., 2023)	Human	×	×	×	×
Instruction Mining (Cao et al., 2023)	PPL & Reward scores &	✓	×	×	0
AlpaGasus (Chen et al., 2024)	ChatGPT Score	×	×	×	×
SelectIT (Liu et al., 2024b)	Multi-granularity Uncertainty	✓	×	×	0
NAIT (Ours)	Neuron Activation	✓	✓	✓	1

2.1 HIGH-QUALITY IT DATA

Recent studies have highlighted the key role of high-quality IT data in enhancing the alignment of LLMs with human preferences as well as generating accurate, relevant, and safe responses (Wang et al., 2024a; Liu et al., 2024a). Typically, IT data consists of instruction-response pairs, which can be created in two main ways: (1) by reformulating traditional NLP task data into the IT format (e.g., FLAN (Wei et al., 2022) and P3 (Sanh et al., 2022)) or (2) by synthesizing new data using self-instruction based on a small set of manually crafted seed IT data (e.g., Alpaca (Taori et al., 2023)). Subsequent work has further explored methods for synthesizing high-quality IT data. For example, Vicuna (Chiang et al., 2023) enhanced data diversity with large-scale dialogue data, while WizardLM (Xu et al., 2024a) used LLMs to automatically generate IT datasets with controllable complexity. Notably, LIMA (Zhou et al., 2023) demonstrated that LLMs fine-tuned on only 1k carefully curated IT data can achieve comparable performance of models trained with large-scale dataset, highlighting the importance of data quality over quantity. These findings suggest that pretrained LLMs already contain extensive world knowledge and require only a small amount of high-quality instruction data to yield strong performance in the IT phase.

However, defining what constitutes "high-quality" IT data remains challenging. Data quality is influenced by complex, multidimensional factors, and there is still a lack of interpretable selection mechanisms and standardized criteria (Wang et al., 2024a).

2.2 EFFICIENT IT DATA SELECTION

Existing IT data selection strategies typically filter subsets based on metrics such as quality, diversity, or importance (Qin et al., 2024). These strategies fall into four main categories based on how these metrics are extracted: (1) Handcrafted Feature-Based Methods rely on manually designed features. For example, DQI (Mishra et al., 2020) selects high-quality data by leveraging lexical features, n-gram frequency, and relational features, while Xie et al. (2023) optimize data distribution using n-gram features. Although these methods more interpretable, they are limited to surfacelevel characteristics and cannot capture deeper data properties; (2) Model Feature-Based Methods extract features such as uncertainty or perplexity from model outputs. For instance, Instruction Mining (Cao et al., 2023) leverages perplexity and reward scores, and SelectIT (Liu et al., 2024b) conducts multi-granularity uncertainty analysis for data selection. However, these methods depend heavily on model outputs, which may introduce bias, and often exhibit "black-box" characteristics that limit interpretability. (3) LLM-as-Scorer Methods utilize advanced LLMs, such as ChatGPT, to score data based on complex criteria. For instance, InsTag (Lu et al., 2024) evaluates diversity and complexity in IT data, while AlpaGasus (Chen et al., 2024) scores data quality similarly. These approaches, though effective, are computationally expensive, lack transparency, and are difficult to control or scale due to reliance on closed-source APIs. (4) Loss and Gradient-based Coreset Sampling Methodsselect data based on loss values or gradient information. For example, Chen et al. (2023) select samples to minimize evaluation loss, and Xia et al. (2024) use gradient signals to identify critical data points. These methods are computationally intensive, especially for large-scale models, and may suffer from approximation errors.

In summary, existing data selection methods often require substantial computational resources, additional model training, or expensive queries, which limit their scalability to large datasets. Furthermore, they may struggle to enhance specific target domain capabilities of LLMs when applied to open datasets.

3 OUR NAIT METHOD

3.1 Preliminary

Neuron activation analysis provides insights into the knowledge storage and operational mechanisms of LLMs (Voita et al., 2024; Durrani et al., 2020; Xu et al., 2024c). Recent studies show that subsets of neurons are activated by specific tasks, playing a critical role in LLMs' ability to process knowledge and solve various tasks (Yu & Ananiadou, 2024; Wang et al., 2024b; Tang et al., 2024).

Motivation Inspired by these findings, we hypothesize that the effectiveness of instruction data may lie in its ability to activate task-relevant neurons. Specifically, when an LLM processes a given sample, the closer the neuron activation pattern aligns with the activation characteristics of the target capability, the more effective the sample is in improving the LLM's performance in that domain.

Comparison Compared to existing methods, the proposed NAIT approach offers significant advantages, as shown in Table 1. Specifically, NAIT directly leverages neuron activation patterns within LLMs, eliminating the need for external models or complex proxy features and greatly improving data selection efficiency. Furthermore, due to the strong correlation between neuron activation and the internal knowledge and tasks of LLMs, NAIT can effectively improve both the model's specialized and general domain capabilities. By capturing the internal neuron activation states during LLM operation, NAIT also provides robust interpretability for data selection. Overall, NAIT outperforms existing methods in efficiency, scalability, and interpretability, offering a more effective solution for IT data selection.

3.2 OUR DATA SELECTION FRAMEWORK

NAIT identifies IT data that effectively enhances specific domain capabilities C in the LLM \mathcal{M} via two main modules: (a) Neuron Activation Feature Extraction and (b) Activation Feature-Guided Data Selection, as illustrated in Figure 1. The proposed framework is detailed below:

3.2.1 Module (a). Neuron Activation Feature Extraction for Target Capabilities $\mathcal C$

To establish the neuron activation features associated with target capabilities C, we implement a three-stage features extraction process:

• In-domain Dataset Preparation Let \mathcal{P} denote the capabilities-specific in-domain dataset containing I representative data from the \mathcal{C} -related tasks:

$$\mathcal{P} = \{P_i\}_{i=1}^I \text{ where } P_i \sim \mathbb{P}_{\mathcal{C}}$$
 (1)

where $\mathbb{P}_{\mathcal{C}}$ represents the intrinsic data distribution of capabilities \mathcal{C} .

• Neuron Activation States Capture For each input sequence $P_i = (t_1, \dots, t_K)$ with K tokens, we record the activation states of all J neurons in \mathcal{M} at each token t_k at the decoder layer, which provide a comprehensive view of the model's internal representations:

$$\mathcal{A}(t_k) = [a_i^{(k)}]_{i=1}^J \in \mathbb{R}^J, \quad k \in 1, \dots, K.$$
 (2)

For each token t_k , we record the neuron-level activation change by comparing the activation state of each neuron N_j with the beginning token of data b:

$$\Delta \mathcal{A}(t_k) = \mathcal{A}(t_k) - \mathcal{A}(b) = [\delta a_j]_{j=1}^J \in \mathbb{R}^J$$
(3)

where $\delta a_j = a_j^{(t_k)} - a_j^{(b)}$ quantifies the activation changes of neuron N_j for token t_k . To summarize the activations for the entire sequence, we compute the mean activation across all K tokens:

$$\Delta \mathcal{A}(P_i) = \frac{1}{K} \sum_{k=1}^{K} \mathcal{A}(t_k) \in \mathbb{R}^J.$$
(4)

Here, $\Delta A(P_i)$ represents the average activation state of all J neurons over the sequence P_i .

• Neuron Activation Feature Construction

To obtain compact feature representations and identify key features, we apply Principal Component Analysis (PCA) for dimensionality reduction:

$$\mathcal{V}_{P_i} = \text{PCA}\left(\Delta \mathcal{A}(P_i), d\right) \in \mathbb{R}^d \tag{5}$$

The characteristic activation feature for capabilities $\mathcal C$ is then derived as the centroid of in-domain features:

$$\mathcal{V}_{\mathcal{P}} = \frac{1}{|\mathcal{P}|} \sum_{i=1}^{I} \mathcal{V}_{P_i} \in \mathbb{R}^d$$
 (6)

3.2.2 MODULE (B). ACTIVATION FEATURE-GUIDED DATA SELECTION

Given candidate dataset $\mathcal{D} = \{D_n\}_{n=1}^N$, we perform IT data selection through neuron activation features matching:

• Feature Alignment Scoring For each candidate $D_n \in \mathcal{D}$, capture its neuron activation states $\mathcal{A}(D_n)$. Then, measure alignment with computed neuron activation features by projecting: the activations features vector of candidate $D_n \in \mathcal{D}$ onto the computed neuron activation features:

$$S(D_n) = A(D_n) \cdot \mathcal{V}_{\mathcal{P}} \tag{7}$$

• Data Ranking & Selection Select top-k candidates maximizing alignment scores:

$$\mathcal{D}_{\text{selected}} = \underset{D_n \in \mathcal{D}}{\text{top-}k} \left(\mathcal{S}(D_n) \right) \tag{8}$$

4 EXPERIMENTS

4.1 SETUPS

Benchmark We evaluate NAIT across five representative capability domains using widely adopted benchmarks: factual knowledge with MMLU (Hendrycks et al., 2021), mathematical reasoning with GSM (Cobbe et al., 2021), general reasoning with BBH (Suzgun et al., 2023), multilingual understanding with TyDiQA (Clark et al., 2020), and coding ability with BigCodeBench (Zhuo et al., 2024) and HumanEval (Chen et al., 2021). We use in-domain splits as capability references and conduct evaluation on the corresponding test sets. Benchmark setups are provided in Appendix D.

Baselines We compare NAIT against a diverse set of representative IT data construction and selection approaches: Alpaca-GPT4 (Peng et al., 2023), a widely used self-instruct dataset synthesized by GPT-4; LIMA (Zhou et al., 2023), which shows that a small amount of carefully curated high-quality IT data can be highly effective; AlpaGasus (Chen et al., 2024), which leverages ChatGPT to score and filter data; Q2Q (Li et al., 2024), which evaluates data quality via loss signals from a precursor model; and SelectIT (Liu et al., 2024b), which selects high-quality data by exploiting uncertainty estimates from base LLMs. Detailed descriptions are provided in Appendix D.

Implementation Details In this study, we adopt LLaMA-2-7b as the foundational model for fine-tuning. The detailed fine-tuning and text generation settings are provided in Appendix D.

4.2 MAIN RESULTS

Overall Performance The NAIT method effectively enhances model's overall performance by leveraging neuron activation features across multiple domain capabilities. As shown in Table 2, System 01 represents the baseline performance of the Alpaca-GPT4, while System 12 demonstrates the proposed NAIT method surpassing Alpaca-GPT4 under full fine-tuning setting. Notably, NAIT achieves up to a 3% performance improvement across five test datasets using only 10% of the original data, significantly surpassing existing IT data selection methods, such as AlpaGasus (System 03), Q2Q (System 04), and SelectIT (System 05). This demonstrates NAIT's advantages in terms of data efficiency and generalization. We also observe that activating certain domain capabilities can affect the model's performance on other downstream tasks. For example, activating only the in-domain capability with TydiQA (System 10) leads to decreased performance on math and code tasks. In contrast, when all in-domain capabilities are activated (System 12), model's mathematical performance improves.

Table 2: **Performance comparison of baslines using 10% of the IT data.** This table shows the results of various tasks across multiple baselines. NAIT (e.g., NAIT (MMLU)) refers to the process where in-domain dataset to guide the IT data selection. Random refers to a randomly sampled 10% subset of the IT data. The **Bold** and <u>Underline</u> represent the best and second performance respectively in each column. Δ (\uparrow) indicates the performance improvement relative to the ID 01.

System ID	Method	MMLU	BBH	GSM	TydiQA	CodeX	AVG	Δ (\uparrow)
Full Fine-tuning								
01	Alpaca-GPT4 (Peng et al., 2023)	46.87	39.94	14.63	39.48	27.87	34.16	
02	LIMA (Zhou et al., 2023)	45.20	39.91	15.76	44.92	27.75	34.71	+1.61%
03	01 + AlpaGasus (Chen et al., 2024)	43.21	39.91	13.34	44.70	23.94	33.02	-3.33%
04	01 + Q2Q (Li et al., 2024)	46.73	40.34	14.50	44.41	25.19	34.23	+0.22%
05	01 + SelectIT (Liu et al., 2024b)	47.90	40.33	15.40	43.91	27.92	35.09	+2.72%
06	01 + Random Baseline	47.14	39.21	14.13	44.16	25.55	34.04	-0.35%
	Our Proposed Met	hod (Indivi	dual Cap	pability F	eatures)			
07	01 + NAIT (MMLU)	47.81	38.52	15.68	47.16	25.23	34.88	+2.11%
08	01 + NAIT (BBH)	47.78	40.46	13.34	45.93	25.15	34.53	+1.09%
09	01 + NAIT (GSM)	46.45	40.28	16.00	46.54	27.84	35.42	+3.70%
10	01 + NAIT (TydiQA)	46.17	40.00	13.80	47.78	25.02	34.55	+1.16%
11	01 + NAIT (CodeX)	47.51	39.72	15.53	44.19	28.49	35.09	+2.72%
	Our Proposed	Method (Al	l Capabi	lity Feati	ures)			
12	01 + NAIT (7 + 8 + 9 + 10 + 11)	46.83	40.02	16.53	46.09	26.44	35.18	+3.00%

NAIT with Distinct Capability Features The experimental results from System 07 to 11 demonstrate that the NAIT method can achieve more consistent and stable performance improvements when targeting individual-specific abilities. Notably, this advantage is particularly pronounced in complex tasks such as GSM, TydiQA, and CodeX, with improvements of 1.87%, 3.62%, and 2.94% over the random baseline, respectively. Further analysis reveals that the contributions of neuron activation features vary across different task domains and exhibit a certain degree of cross-domain transferability. For instance, neuron activation features extracted based on the BBH in-domain dataset can significantly enhance the model's performance on the MMLU task, while those extracted from the CodeX in-domain dataset can likewise effectively improve the model's performance on the GSM task, with the improvement approaching that of their respective original activation states. It is worth mentioning that the performance of the NAIT method on the GSM dataset reaches as high as 3.80%, even surpassing the model under multi-capability activation (System 12). These results fully demonstrate that the NAIT method not only effectively eliminates redundant or highly similar samples, but also maintains data diversity, thereby further validating its effectiveness and generalization capability in optimizing data utilization and enhancing task-specific performance.

5 ANALYSIS

This section aims to analyze further and address the following research questions: (1) How do the number of in-domain datasets, the proportion of IT datasets, model size, and IT data methods affect NAIT? (see § 5.1) (2) What are the advantages of NAIT in cost efficiency? (see § 5.2) (3) How interpretable in NAIT? (see § 5.3) (4) How does NAIT perform in target domain? (§ Appendix F)

5.1 ABLATION STUDY

Proportion of Selected IT Dataset Although NAIT demonstrates excellent performance in data evaluation and ranking, selecting the optimal data proportion from IT datasets with a large amount of redundant information remains challenging. To further investigate this, we adjusted the proportion of the selected IT dataset from 10% to 100% and analyzed the performance of NAIT under different IT data proportions. As shown in Figure 2, the top 30% of IT samples selected by NAIT exhibited the best performance. Notably, as the proportion of training data continued to increase, the model's performance showed an overall downward trend, reaching its lowest point when the entire dataset (100%) was used. These results indicate that excessive redundant data may undermine the model's generalization, further emphasizing the need for data selection to maximize model performance.

325

326

327

328

329

330

331

332

333

334

335

336

337

338

339

340

341

342

343

344

345

347

348

349

350

351

352

353 354

355

356

357

358

359

360

361 362

363

364

366

367

368

369

370

371

372

373

374

375

376

377

Number of In-domain Data The number of indomain data is a critical parameter in our method, as it directly influences the effectiveness of extracting neuron activation features for specific domain capabilities. As shown in Figure 3, we fine-tune the model using IT data selected based on the activation patterns of each in-domain data group and evaluate its performance on corresponding tasks. The results indicate that NAIT outperforms the random selection baseline across all tasks in most cases. Even when the data size is small (e.g., 16, 64, or 256 samples), most cases exceed the random baseline. This suggests that even with limited data, the extracted features effectively capture the required capabilities. As the number of in-domain data increases, performance improvements gradually level off, with most tasks reaching their maximum performance at the largest data size. However, for GSM and TydiQA, peak performance is achieved at 18.35 and 47.57, respectively, when the data size is 4096. This highlights that the quality of in-domain data also plays an important role in enabling NAIT to capture task-specific features.

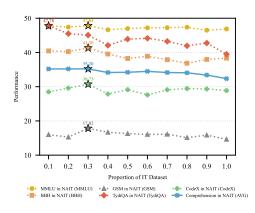


Figure 2: **Performance at different proportions of the IT dataset.** Task performance in NAIT (the in-domain dataset), e.g., MMLU in NAIT, refers to using in-domain data to guide IT data selection and to assess task outcomes. The comprehensive results correspond to System 12.

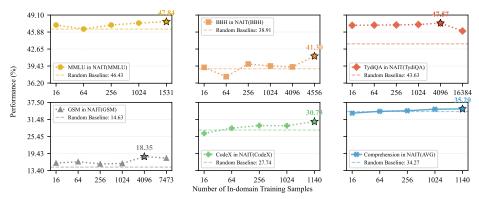


Figure 3: Comparison of NAIT with random data selection across different in-domain data scales at 30% IT dataset. Task performance in NAIT (in-domain dataset), such as MMLU in NAIT (MMLU), refers to using the in-domain dataset to guide IT data selection and evaluate task performance. The dashed line indicates the 30% of the IT dataset random selection baseline.

Various Foundation Models Our method demonstrates significant performance improvements when fine-tuning the *LLaMA-2-7b* model with the curated Alpaca-NAIT dataset. evaluate its robustness and generalizability, we extend our study to other foundational models, including (1) LLaMA-2-13b, representing models of different scales; and (2) Mistral-7b and LLaMA-3-8b, representing models with different architectures. As shown in Table 3, NAIT achieves similar performance improvements on LLaMA-2-13b, demonstrating effectiveness across models of varying sizes. Additionally, NAIT achieves gains of +21.92% on Mistral-7b and +18.65% on LLaMA-3-8b, highlighting its potential for diverse architectures and advanced models.

Table 3: **Effectiveness of NAIT Across Models.** Alpaca and Random are trained on the full dataset and a random of 10% subset, while NAIT selects a task-relevant 10% subset.

Method	MMLU	ввн	GSM	TydiQA	CodeX	o	verall		
		2211	00.11	1,414.1	Coucia	AVG	$\Delta \left(\uparrow \right)$		
			LLam	a-2-13b					
Alpaca	53.90	45.00	20.85	44.13	34.84	39.74			
+Random	52.70	47.59	22.06	44.66	36.79	40.76	+2.57%		
+NAIT	54.10	47.04	24.11	48.89	38.53	42.53	+7.02%		
			Mist	ral-7b					
Alpaca	47.00	40.19	12.89	35.40	34.42	33.98			
+Random	48.37	46.11	17.36	34.80	41.81	37.69	+10.92%		
+NAIT	52.90	49.07	18.95	41.22	45.02	41.43	+21.92%		
Llama-3-8b									
Alpaca	59.60	48.38	29.11	48.38	57.36	48.57			
+Random	48.37	55.93	37.23	54.65	69.63	52.10	+7.27%		
+NAIT	60.80	60.93	43.91	47.72	74.78	57.63	+18.65%		

Different Instruction Tuning **Datasets** To further verify the effectiveness of our method, we extend our IT selection method to other IT datasets, including Evo-Instruct (Xu et al., 2024a) using WizardLM and Orca-GPT4 (Mukherjee et al., 2023). The experimental results are shown in Figure 4, as the proportion of IT data increases, the model performance rises and then declines. Notably, the efficient IT data subsets selected by NAIT achieve optimal performance at 50% (Orca-GPT4) and 80% (Evo-Instruct), significantly outperforming the use of the full dataset. Additionally, Evol-Instruct and Orca-GPT4 require a higher proportion of data, especially Evol-Instruct, which needs up to 80% to achieve optimal performance, reflecting their higher quality, complexity, and information density compared to Alpaca-GPT4.

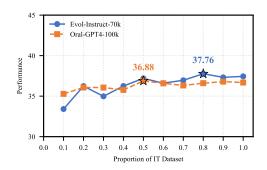


Figure 4: Comparison of NAIT across different **IT dataset.** Task performance in NAIT refers to using the all in-domain dataset to guide IT data selection and evaluate task performance.

Different Selection Strategies We also conducted experiments using context length and perplexity as measures of data complexity, following Liu et al. (2024c), to provide an additional insight into our method's effectiveness, as detailed in Appendix E.

5.2 Cost Efficiency

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394 395 396

397

398 399

400 401

402

403

404

405

406

407

408

409

410

411

412 413

414 415

416

417

418

419

420

421

422

423 424

425

426

427

428 429

430

431

We compare inference time and cost between NAIT and existing approaches, including AlpaGasus (GPT-4o), Q2Q (GPT-4o), and SelectIT, on the Alpaca-GPT4 dataset (52k). As shown in Table 4, NAIT achieves up to 19x and 4× cost reductions compared to AlpaGasus and O2O, respectively, along with significant improvements in inference speed by 17.75 and 2.2 hours. Compared to SelectIT, NAIT reduces cost by 94.3% and achieves a 17.58× speedup. These results highlight the feasibility and effectiveness of NAIT in real-world applications.

Table 4: Efficiency Comparison of Different Methods on NVIDIA A800 80GB with Batch Size set to 8. API costs follow the official OpenAI pricing (\$2.00/million input tokens, \$8/million output tokens for GPT-4.1); GPU costs are estimated based on the Google Cloud pricing (\$1.15 per GPU hour for NVIDIA A800 80GB).

Method	Externally-Independent	Time	Cost
AlpaGasusChen et al. (2024)	Х	19.07h	\$178.02
Q2QLi et al. (2024)	X	3.52h	\$4.05
SelectITLiu et al. (2024b)	✓	23.20h	\$26.68
NAIT	✓	1.32h	\$1.52

5.3 Interpretability Analysis

Transferability of Neural Activation Feature From the results in Table 2, we observe that the neural activation features corresponding to different capabilities exhibit remarkable variation in their transferability. For instance, the neural activation feature extracted based on GSM significantly enhances the overall average performance of LLMs, and it also yields positive effects on BBH and CodeX tasks. To further investigate the relationships between the transferability of neural activation features, we introduce **Transferability** to evaluate each capability during the transfer process. Transferability of capability i reflects to what extent i can enhance other capabilities when activated. It is defined as

Transferability_i =
$$Acc(i, j) - Acc(j, j)$$
 (9)

where Acc(i, j) is the performance on task j activated by feature i, and Acc(j, j) is task j's baseline performance with its own activation feature.

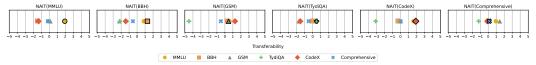


Figure 5: Transferability of neural activation features across capabilities. Each column represents the capability to which the neural activation feature is applied. Icons with an outline border indicate the capability's performance on its own task, which serves as the baseline reference.

As shown in Figure 5, the transferability across capabilities can be directly observed. On the one hand, activation features of single capabilities such as GSM and CodeX exhibit evident positive transfer in cross-task settings, indicating that logical reasoning and programmatic features possess strong general power. In contrast, the activation features of TydiQA show weaker or even negative transfer, and their performance within the task itself is relatively limited, reflecting both a dependence on language-specific patterns and shortcomings in cross-domain adaptability. On the other hand, when multiple capabilities are aggregated into a comprehensive activation direction, the model achieves optimal overall transferability. This finding demonstrates the existence of a universal core feature whose neural activation patterns remain stable across tasks, thereby providing a solid foundation for the model's integrated capabilities.

Relation of Activation Features To further investigate the interpretability of NAIT, we apply T-SNE to project the neuron activation features into representation space (see Figure 6). The figure shows that the feature's directions corresponding to different targeted capabilities form relatively separable clusters in the representation space, indicating their consistency with the task-specific reasoning mechanisms or knowledge domains. TydiQA is farther from the other capability directions, suggesting weaker transferability. Notably, the comprehensive, as an aggregated representation across multiple tasks, reveals the **general core capabilities.** This indicates that there exist general data or features capable of simultaneously activating multiple distinct abilities, thereby supporting cross-task knowledge transfer and capability integration.

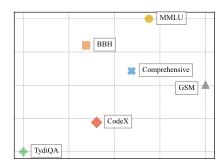


Figure 6: Comparison of the direction of activation feature extracted by NAIT across different capabilities.

Selected IT Data Distribution To further explore this general capability, we conduct a statistical analysis of the IT subsets selected from different activation features. The results show that when the selection proportion is set to 30%, more than 80% of the subsets overlap across different capabilities. Even under a stricter 10% selection proportion, the overlap still exceeds 70%, indicating that NAIT tends to identify a stable set of general core data across tasks. we draw the distribution of the 10% IT dataset selected by NAIT. As shown in Figure 7, a considerable portion of the selected samples are repeatedly shared across multiple targeted capabilities, while a smaller fraction remains task-specific, further validating the universal capability that NAIT achieves in identifying data with broad applicability.

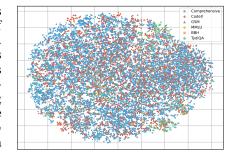


Figure 7: **Distribution of 10% IT dataset selected by NAIT.** Each subset selected by different activation features shows more than 70% overlap with the comprehensive subset.

6 Conclusion

In this paper, we propose NAIT, an efficient framework for selecting high-quality IT data by leveraging neural activation patterns. NAIT identifies the optimal data subset by evaluating the alignment between candidate samples and neural activation features associated with target capabilities. Experimental results demonstrate that models fine-tuned on selected data consistently outperform baseline methods. Ablation studies further verify that only a small number of samples are sufficient to accurately extract activation features that capture model capabilities. In addition, NAIT exhibits strong robustness and reveals significant cross-task transferability of neural activation features. Further analysis shows that data with logical reasoning and programmatic characteristics tend to activate stronger general capabilities in the model, and that a stable core subset of data exists that can universally enhance performance across a variety of tasks.

ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics¹. Our research focuses on improving the efficiency and interpretability of IT for LLMs through neurally informed data selection. We do not introduce new data collection involving human participants, nor do we use private or personally identifiable data. All datasets employed (e.g., Alpaca-GPT4, GSM8K, MMLU, BBH, TyDiQA, CodeX) are publicly available and widely used in the research community. We carefully respect data licenses and ensure compliance with privacy and intellectual property standards. Our method aims to reduce the computational and environmental cost of developing LLMs by selecting smaller subsets of high-quality data, which aligns with the principle of minimizing harm. However, as with other LLM research, there is the potential risk of misuse in generating biased or harmful outputs. To mitigate this, we explicitly analyze fairness and generalization across different task domains and report transferability results that highlight both strengths and limits of our approach. We disclose all relevant details transparently to foster responsible use of our method.

Reproducibility statement

We have taken several steps to ensure reproducibility of our results. Firstly, we describe the construction process of our neuron activation feature framework in § 3 with explicit mathematical definitions. Additionally, detailed descriptions of the model architectures, datasets, hyperparameters, and training setups are provided in § 4.1 and § Appendix D. Further experimental details, including ablation studies and comparisons with baselines, are presented in § 5 and § Appendices E to F. All in-domain datasets and IT datasets used in our experiments (e.g., Alpaca-GPT4, Evo-Instruct, Orca-GPT4) are publicly available, and dataset statistics are summarized in § Appendix D. To facilitate reproducibility, we commit to open-sourcing the implementation of NAIT, the cross-task neuron feature library, and the Alpaca-NAIT dataset upon acceptance. These resources, along with clear documentation, will allow researchers to fully reproduce and extend our experiments.

REFERENCES

Yihan Cao, Yanbin Kang, and Lichao Sun. Instruction mining: High-quality instruction data selection for large language models. *CoRR*, abs/2307.06290, 2023. doi: 10.48550/ARXIV.2307.06290. URL https://doi.org/10.48550/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. Alpagasus: Training a better alpaca with fewer data. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.* OpenReview.net, 2024. URL https://openreview.net/forum?id=FdVXgSJhvz.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. *CoRR*, abs/2107.03374, 2021. URL https://arxiv.org/abs/2107.03374.

Mayee F. Chen, Nicholas Roberts, Kush Bhatia, Jue Wang, Ce Zhang, Frederic Sala, and Christopher Ré. Skill-it! A data-driven skills framework for understanding and training language models. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural

https://iclr.cc/public/CodeOfEthics

Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/70b8505ac79e3e131756f793cd80eb8d-Abstract-Conference.html.

- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See https://vicuna. lmsys. org (accessed 14 April 2023), 2(3):6, 2023.
- Jonathan H. Clark, Jennimaria Palomaki, Vitaly Nikolaev, Eunsol Choi, Dan Garrette, Michael Collins, and Tom Kwiatkowski. Tydi QA: A benchmark for information-seeking question answering in typologically diverse languages. *Trans. Assoc. Comput. Linguistics*, 8:454–470, 2020. doi: 10.1162/TACL\A_00317. URL https://doi.org/10.1162/tacl a 00317.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *CoRR*, abs/2110.14168, 2021. URL https://arxiv.org/abs/2110.14168.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. Analyzing individual neurons in pre-trained language models. In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (eds.), Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020, pp. 4865–4880. Association for Computational Linguistics, 2020. doi: 10.18653/V1/2020.EMNLP-MAIN.395. URL https://doi.org/10.18653/v1/2020.emnlp-main.395.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net, 2021. URL https://openreview.net/forum?id=d7KBjmI3GmQ.
- Ming Li, Yong Zhang, Zhitao Li, Jiuhai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. From quantity to quality: Boosting LLM performance with self-guided data selection for instruction tuning. In Kevin Duh, Helena Gómez-Adorno, and Steven Bethard (eds.), Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024, pp. 7602–7635. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.NAACL-LONG.421. URL https://doi.org/10.18653/v1/2024.naacl-long.421.
- Fuxiao Liu, Xiaoyang Wang, Wenlin Yao, Jianshu Chen, Kaiqiang Song, Sangwoo Cho, Yaser Yacoob, and Dong Yu. MMC: advancing multimodal chart understanding with large-scale instruction tuning. In Kevin Duh, Helena Gómez-Adorno, and Steven Bethard (eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pp. 1287–1310. Association for Computational Linguistics, 2024a. doi: 10.18653/V1/2024.NAACL-LONG.70. URL https://doi.org/10.18653/v1/2024.naacl-long.70.
- Liangxin Liu, Xuebo Liu, Derek F. Wong, Dongfang Li, Ziyi Wang, Baotian Hu, and Min Zhang. Selectit: Selective instruction tuning for llms via uncertainty-aware self-reflection. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 15, 2024, 2024b. URL http://papers.nips.cc/paper_files/paper/2024/hash/b130a5691815f550977e331f8bec08ae-Abstract-Conference.html.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. What makes good data for alignment? A comprehensive study of automatic data selection in instruction tuning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria,*

May 7-11, 2024. OpenReview.net, 2024c. URL https://openreview.net/forum?id=BTKAeLqLMw.

- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. #instag: Instruction tagging for analyzing supervised fine-tuning of large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.* OpenReview.net, 2024. URL https://openreview.net/forum?id=pszewhybU9.
- Swaroop Mishra, Anjana Arunkumar, Bhavdeep Singh Sachdeva, Chris Bryan, and Chitta Baral. DQI: measuring data quality in NLP. *CoRR*, abs/2005.00816, 2020. URL https://arxiv.org/abs/2005.00816.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. Orca: Progressive learning from complex explanation traces of GPT-4. *CoRR*, abs/2306.02707, 2023. doi: 10.48550/ARXIV.2306.02707. URL https://doi.org/10.48550/arXiv.2306.02707.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 December 9, 2022, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/blefde53be364a73914f58805a001731-Abstract-Conference.html.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with GPT-4. *CoRR*, abs/2304.03277, 2023. doi: 10.48550/ARXIV.2304.03277. URL https://doi.org/10.48550/arXiv.2304.03277.
- Yulei Qin, Yuncheng Yang, Pengcheng Guo, Gang Li, Hang Shao, Yuchen Shi, Zihan Xu, Yun Gu, Ke Li, and Xing Sun. Unleashing the power of data tsunami: A comprehensive survey on data assessment and selection for instruction tuning of language models. *CoRR*, abs/2408.02085, 2024. doi: 10.48550/ARXIV.2408.02085. URL https://doi.org/10.48550/arXiv.2408.02085.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net, 2022. URL https://openreview.net/forum?id=9Vrb9D0WI4.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging bigbench tasks and whether chain-of-thought can solve them. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics:* ACL 2023, Toronto, Canada, July 9-14, 2023, pp. 13003–13051. Association for Computational Linguistics, 2023. doi: 10.18653/V1/2023.FINDINGS-ACL.824. URL https://doi.org/10.18653/v1/2023.findings-acl.824.
- Tianyi Tang, Wenyang Luo, Haoyang Huang, Dongdong Zhang, Xiaolei Wang, Xin Zhao, Furu Wei, and Ji-Rong Wen. Language-specific neurons: The key to multilingual capabilities in large language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16, 2024, pp. 5701–5715. Association for

- Computational Linguistics, 2024. doi: 10.18653/V1/2024.ACL-LONG.309. URL https://doi.org/10.18653/v1/2024.acl-long.309.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models. https://crfm. stanford. edu/2023/03/13/alpaca. html*, 3(6):7, 2023.
- Elena Voita, Javier Ferrando, and Christoforos Nalmpantis. Neurons in large language models: Dead, n-gram, positional. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024, pp. 1288–1301. Association for Computational Linguistics, 2024. doi: 10.18653/V1/2024.FINDINGS-ACL.75. URL https://doi.org/10.18653/v1/2024.findings-acl.75.
- Jiahao Wang, Bolin Zhang, Qianlong Du, Jiajun Zhang, and Dianhui Chu. A survey on data selection for LLM instruction tuning. *CoRR*, abs/2402.05123, 2024a. doi: 10.48550/ARXIV.2402.05123. URL https://doi.org/10.48550/arXiv.2402.05123.
- Yifei Wang, Yuheng Chen, Wanting Wen, Yu Sheng, Linjing Li, and Daniel Zeng. Unveiling factual recall behaviors of large language models through knowledge neurons. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pp. 7388–7402. Association for Computational Linguistics, 2024b. URL https://aclanthology.org/2024.emnlp-main.420.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022.* OpenReview.net, 2022. URL https://openreview.net/forum?id=gEZrGCozdqR.
- Zhengxuan Wu, Atticus Geiger, Thomas Icard, Christopher Potts, and Noah D. Goodman. Interpretability at scale: Identifying causal mechanisms in alpaca. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/f6a8b109d4d4fd64c75e94aaf85d9697-Abstract-Conference.html.
- Mengzhou Xia, Sadhika Malladi, Suchin Gururangan, Sanjeev Arora, and Danqi Chen. LESS: selecting influential data for targeted instruction tuning. In Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024. OpenReview.net, 2024. URL https://openreview.net/forum?id=PG5fV50maR.
- Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy Liang. Data selection for language models via importance resampling. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/6b9aa8f418bde2840d5f4ab7a02f663b-Abstract-Conference.html.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. Wizardlm: Empowering large pre-trained language models to follow complex instructions. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024.* OpenReview.net, 2024a. URL https://openreview.net/forum?id=CfXh93NDgH.
- Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. A paradigm shift in machine translation: Boosting translation performance of large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024b. URL https://openreview.net/forum?id=farT6XXntP.

Haoyun Xu, Runzhe Zhan, Derek F. Wong, and Lidia S. Chao. Let's focus on neuron: Neuron-level supervised fine-tuning for large language model. *CoRR*, abs/2403.11621, 2024c. doi: 10.48550/ARXIV.2403.11621. URL https://doi.org/10.48550/arXiv.2403.11621.

Zeping Yu and Sophia Ananiadou. Neuron-level knowledge attribution in large language models. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, pp. 3267–3280. Association for Computational Linguistics, 2024. URL https://aclanthology.org/2024.emnlp-main.191.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. LIMA: less is more for alignment. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/ac662d74829e4407ce1d126477f4a03a-Abstract-Conference.html.

Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, Thong Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kaddour, Ming Xu, Zhihan Zhang, Prateek Yadav, Naman Jain, Alex Gu, Zhoujun Cheng, Jiawei Liu, Qian Liu, Zijian Wang, David Lo, Binyuan Hui, Niklas Muennighoff, Daniel Fried, Xiaoning Du, Harm de Vries, and Leandro von Werra. Bigcodebench: Benchmarking code generation with diverse function calls and complex instructions. *CoRR*, abs/2406.15877, 2024. doi: 10.48550/ARXIV.2406.15877. URL https://doi.org/10.48550/arXiv.2406.15877.

A LIMITATION

Model Scale Due to computational limitations, our analysis is restricted to models with fewer than 20B parameters. Future studies could evaluate this method on larger models to provide insights into its scalability and broader implications for optimizing LLMs.

B THE USE OF LARGE LANGUAGE MODELS (LLMS)

Polishing the Writing with a Large Language Model In preparing this paper, we used a large language model to refine the writing. Typical applications included: (1) enhancing grammar and style; (2) ensuring consistency of terminology; and (3) improving the quality of translations.

C COMPARATIVE ANALYSIS BEFORE AND AFTER IT

Table 5: Comparison of model functions before and after instruction tuning.

Function	Before Instruction Tuning	After Instruction Tuning
Enhancing instruction-following ability	Limited to text completion	Capable of following instructions
Activating latent abilities	Relies on knowledge from pertaining	Activates abilities not evident in pretrain- ing, such as logical reasoning or task de- composition
Enhancing downstream task performance	General knowledge capabilities	Improved downstream task capabilities

The following table 5 compares the main functional differences of models before and after IT. As shown, IT significantly improves instruction-following abilities, activates latent skills such as reasoning and task decomposition, and leads to notable gains in downstream task performance.

Benchmark Settings We evaluated our approach on various downstream tasks using their data as in-domain data (detailed dataset descriptions are in Appendix D):

- Factual Knowledge: We evaluate the Massive Multitask Language Understanding (MMLU) dataset (Hendrycks et al., 2021), splitting it into in-domain and test sets, and report 5-shot results.
- *Math Capability*: We use the Grade School Math (GSM) dataset (Cobbe et al., 2021), divided into in-domain and test sets, and report 5-shot results under the CoT setting.
- *Reasoning*: We assess the Big-Bench-Hard (BBH) dataset (Suzgun et al., 2023), with in-domain and test splits, and report 5-shot results under the CoT setting.
- *Multilingual Understanding*: We evaluate the TyDiQA benchmark (Clark et al., 2020), a multilingual QA dataset, using in-domain and test splits, and report 5-shot results with the gold-passage setup.
- Coding Capability: We evaluate coding ability using BigCodeBench (Zhuo et al., 2024) and HumanEval (Chen et al., 2021), reporting pass@10 results with a sampling temperature of 0.8.

D DETAILS OF TRAINING AND EVALUATION SETUP

Statistics of the in-domain, IT, and Test dataset Table 6 summarizes the statistics of the Indomain, IT, and Test datasets.

Table 6: Statistics of the	Number of in-domain	. IT. and Test dat	asets Used in Experiment.

DownStream Task	MMLU	BBH	GSM	Tydiqa	CodeX
In-domain Dataset	1,531	4,556	7,473	49,400	1140
Alpaca-GPT4 IT Dataset			52,002		
Evol-Instruct IT Dataset	Evol-Instruct IT Dataset 70,000				
Oral-GPT4 IT Dataset 100,000					
Test Dataset	14042	1080	1319	900	164
Test few-Shot	5	1	8	1	0
Test Dataset		1080 1	1319		10.

Baselines Settings Our comparison baselines cover a variety of advanced methods for evaluating and improving IT datasets:

- *Alpaca-GPT4* (Peng et al., 2023): A widely utilized IT dataset that leverages a self-instruct methodology, enabling the autonomous generation of instructions through the advanced capabilities of GPT-4.
- *LIMA* (Zhou et al., 2023): It primarily comprises 1,000 meticulously crafted, high-quality IT data points designed to enhance the alignment capabilities of LLMs.
- *AlpaGasus* (Chen et al., 2024): This approach employs the advanced capabilities of ChatGPT to evaluate and selectively curate data from the original Alpaca dataset.
- Q2Q (Li et al., 2024): It functions by training a precursor model and assessing the quality of the
 instructional data based on two distinct loss values derived from this model.
- **SelectIT** (Liu et al., 2024b): A novel method that leverages the intrinsic uncertainty of LLMs to select high-quality IT data without requiring external resources.

Fine-tuning Parameters The model is fine-tuned over 3 epochs with a batch size of 128 to ensure efficient learning while avoiding overfitting. The optimization process employs the Adam optimizer with hyperparameters set to $\beta_1=0.9$ and $\beta_2=0.999$, which are standard values for stable and effective training. The learning rate follows a cosine decay schedule, starting at 2e-5 and decreasing gradually to 0, a strategy known to improve convergence by reducing oscillations during later training stages. To maximize the model's performance, we utilize an input sequence length of 2048 tokens, as this configuration has been demonstrated to be effective for LLMs in handling long-context tasks. Table 7 provides an overview of the key hyperparameters used during training on 4 A800 GPUs with 1 node.

Table 7: Key hyperparameter settings for supervised fine-tuning.

Parameter Key	Value
Learning rate	2.0e-5
batch size	128
cutoff length	2048
Ir scheduler type	cosine
bf16	True
warmup ratio	0.03
weight decay	0.1
epoch	3

Text Generation Settings For text generation, we employ greedy search for computational efficiency, and set the temperature to 0.8 and the top-p to 0.95 following the HumanEval setting (Chen et al., 2021), to enhance the creativity and diversity of the generated content while maintaining its accuracy and contextual relevance.

E THE DETAILS OF DIFFERENT SELECTION STRATEGIES

Table 8: Performance comparison of different selection strategies using 10% of the IT data. The Bold represents the best performance respectively in each column. Δ (\uparrow) indicates the performance improvement relative to the ID 01 baseline.

Method	MMLU	ввн	GSM	M TvdiOA CodeX _		O	Overall		
				-J Q		AVG	Δ (\uparrow)		
Alpaca-GPT4	46.87	39.94	14.63	39.48	27.87	34.16	-		
+ Hard (PPL)	45.52	22.22	8.72	27.87	20.56	24.98	-26.87%		
+ Easy (PPL)	46.78	38.15	15.39	33.12	28.93	32.47	-5.53%		
+ Hard (Length)	46.10	36.94	16.00	34.32	30.89	32.85	-3.83%		
+ Easy (Length)	44.15	10.28	6.37	28.07	17.00	21.17	-38.03%		
+ NAIT	46.83	40.02	16.53	46.09	26.44	35.18	+3.00%		

F NAIT FOR CROSS-LINGUAL

Table 9: **Overall results on machine translation with LLMs using the ALMA-7b.** NAIT indicates selecting the top 30% of data by NAIT, Rand. refers to randomly selecting 30% of data, and Full uses the full dataset.

Method	En⇒	De	De⇒	En	Zh⇒	En	En⇒Zh		Overall	
	COMET	BLEU	COMET	BLEU	COMET	BLEU	COMET	BLEU	COMET	BLEU
+ Full	84.54	28.24	83.13	29.06	79.80	22.89	84.15	34.58	82.91	28.69
+ Rand.	82.44	28.29	81.12	27.20	78.53	21.11	82.12	33.91	81.05	27.63
+ Nait	85.14	29.02	82.50	28.70	80.22	23.11	83.89	35.92	82.94	29.19

As shown in Table 9, we further evaluate the cross-lingual data selection capability of NAIT on the multilingual IT dataset in the context of machine translation. We adopt the powerful ALMA-7b (Xu et al., 2024b) model as the backbone and conduct experiments on four representative translation directions: English-German, German-English, Chinese-English, and English-Chinese. The experimental results demonstrate that the NAIT method outperforms both random selection (Rand.) and

using the full dataset (Full) on most evaluation metrics. Notably, NAIT achieves the highest overall COMET (82.94) and BLEU (29.19) scores, further highlighting its superior generalization ability and robustness in both target domains and cross-lingual scenarios.