

LangGPS: Language Separability Guided Data Pre-Selection for Joint Multilingual Instruction Tuning

Yangfan Ye¹, Xiaocheng Feng^{1,2*}, Xiachong Feng^{3*}, Lei Huang¹, Weitao Ma¹, Qichen Hong⁴, Yunfei Lu⁴, Duyu Tang⁴, Dandan Tu^{4*}, Bing Qin^{1,2}

¹Harbin Institute of Technology

²Peng Cheng Laboratory

³The University of Hong Kong

⁴Huawei Technologies Co., Ltd

{yfy,xcfc,lfhuang,wtma,qinb}@ir.hit.edu.cn, fengxc@hku.hk, {hongqichen,luyunfei6,tangduyu,tudandan}@huawei.com

Abstract

Joint multilingual instruction tuning is a widely adopted approach to improve the multilingual instruction-following ability and downstream performance of large language models (LLMs), but the resulting multilingual capability remains highly sensitive to the composition and selection of the training data. Existing selection methods, often based on features like text quality, diversity, or task relevance, typically overlook the intrinsic linguistic structure of multilingual data. In this paper, we propose *LangGPS*, a lightweight two-stage pre-selection framework guided by *language separability*—a signal that quantifies how well samples in different languages can be distinguished in the model’s representation space. *LangGPS* first filters training data based on separability scores and then refines the subset using existing selection methods. Extensive experiments across six benchmarks and 22 languages demonstrate that applying *LangGPS* on top of existing selection methods improves their effectiveness and generalizability in multilingual training, especially for understanding tasks and low-resource languages. Further analysis reveals that highly separable samples facilitate the formation of clearer language boundaries and support faster adaptation, while low-separability samples tend to function as bridges for cross-lingual alignment. Besides, we also find that language separability can serve as an effective signal for multilingual curriculum learning, where interleaving samples with diverse separability levels yields stable and generalizable gains. Together, we hope our work offers a new perspective on data utility in multilingual contexts and support the development of more linguistically informed LLMs.

Code — <https://github.com/YYF-Tommy/LangGPS>

Extended version — <https://arxiv.org/pdf/2511.10229>

1 Introduction

Large language models (LLMs) have demonstrated remarkable generalization and adaptability across a wide range of downstream tasks in multilingual scenarios (Ye, Tao, and Kong 2023; Qin et al. 2024; Zhang et al. 2023; Ye

et al. 2024a,b). A central enabler of such capabilities is joint multilingual instruction tuning, in which LLMs are trained on instruction-following data covering multiple languages (Ouyang et al. 2022; Shaham et al. 2024; Huo et al. 2025; Ye et al. 2025). This paradigm allows models to understand and follow instructions in multilingual scenarios, fostering more inclusive access to language technologies.

Recent work has emphasized the importance of multilingual instruction datasets that encompass diverse languages, domains, and task formats (Muennighoff et al. 2022; Zhu et al. 2023; Li et al. 2023; Singh et al. 2024). These efforts have laid a solid foundation for developing LLMs that can effectively align with user intent across linguistic boundaries. Despite these advances, the success of multilingual instruction tuning remains highly sensitive to the composition and selection of the training data. Consequently, **Data Selection** has emerged as a critical strategy for balancing effectiveness and efficiency in large-scale model training. Existing approaches can be typically grouped into two main categories: feature-based methods, which prioritize training samples based on characteristics such as data quality and diversity (Li et al. 2024; Bukharin et al. 2024); and target-dependent methods, which select data according to its estimated relevance to specific target tasks (Xie et al. 2023; Xia et al. 2024). However, little attention has been paid to the intrinsic linguistic structure inherent in multilingual data.

In multilingual scenarios, where the model must learn to follow instructions across diverse and often low-resource languages, we argue that a key prerequisite for effective multilingual instruction following is the model’s ability to form and maintain clear linguistic boundaries. Without sufficiently distinguishing between languages, the model risks conflating linguistic patterns, especially in typical low-resource languages. To formalize this intuition, we introduce the notion of language separability, which quantifies the extent to which the data samples in different languages are distinguishable in the model’s representation space. Highly separable samples tend to exhibit well-structured, language-specific representations, whereas poorly separable ones often entangle with other languages in the representation space, lacking clear linguistic boundaries. As shown in Fig-

*Corresponding Authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

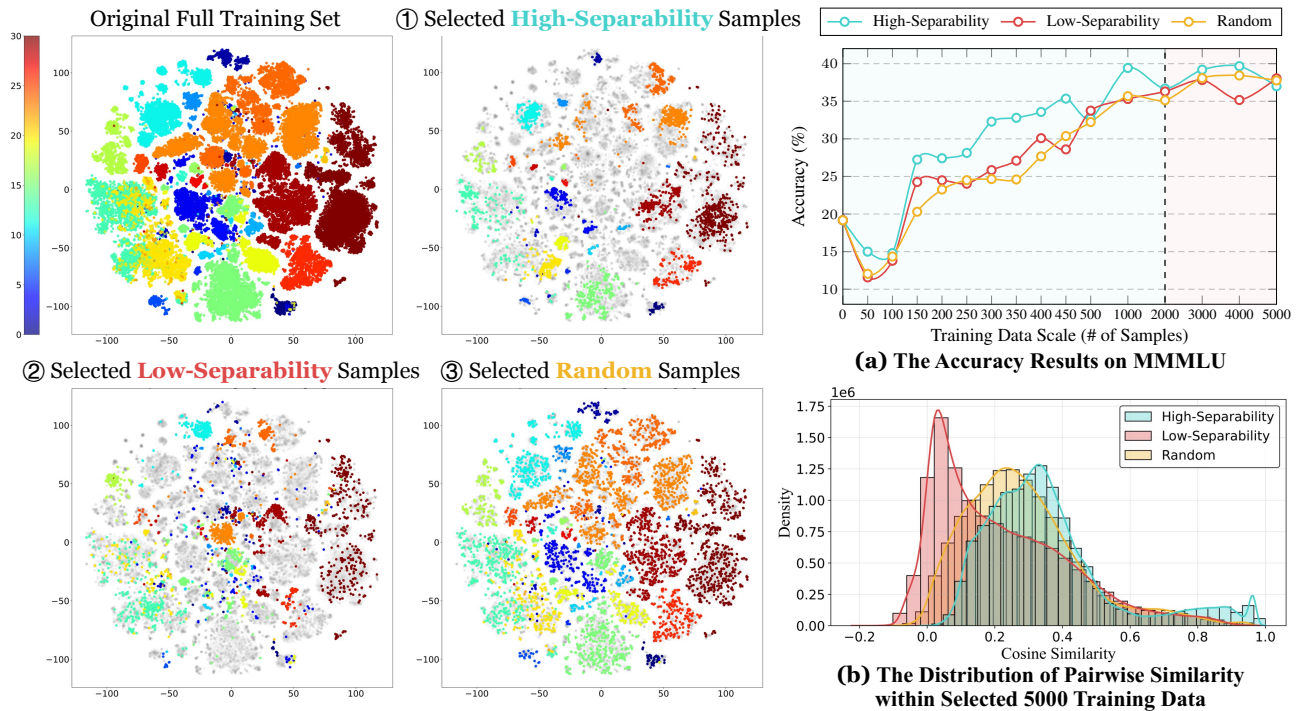


Figure 1: The left shows the t-SNE visualization of samples selected by different strategies under same language distribution; each color denotes one of 31 languages. The representations are last-token hidden states from *LLaMA-3.1-8B*. On the right, panel (a) presents the performance of *LLaMA-3.1-8B* after training with the three types of samples on *MMMLU* across training sizes; while panel (b) illustrates the pairwise similarity distributions of selected data. (See detailed settings in Appendix)

Figure 1 (a), our preliminary experiments (see details in Appendix) on the multilingual MMLU dataset reveal that, under limited training data (< 2000), training with highly separable samples leads to significantly better performance compared to low-separability or randomly selected samples. However, as Figure 1 (b) illustrates, highly separable samples may also suffer from over-similarity (excessive resemblance among samples), introducing a lack of diversity, which is another critical factor in data selection.

Building on these insights, we propose *LangGPS*, a **Language separability Guided data Pre-Selection** strategy for multilingual instruction tuning. *LangGPS* adopts a two-stage framework: (1) a lightweight pre-selection stage that filters training data based on estimated language separability, and (2) a fine selection stage that refines the selected subset using existing data selection methods, such as diversity-based, task relevance-based, or gradient-based approaches. This design integrates scalability, linguistic structure awareness, and full compatibility with existing selection strategies that focus on different aspects, offering a more principled and flexible framework for multilingual instruction tuning.

We conduct extensive experiments across six benchmarks covering both understanding and generation tasks, spanning 22 evaluated languages and two representative LLMs. Our results find that existing data selection methods often overlook the intrinsic linguistic structure inherent in multilingual training data, leading to inconsistent gains in joint multilingual training. By leveraging language separability as a

guiding signal, *LangGPS* can be seamlessly integrated with existing selection methods and improves their effectiveness and generalizability, especially for understanding tasks and low-resource languages. Further analysis reveals that highly separable samples are critical for language-specific modeling and fast adaptation, helping models form and maintain clearer language boundaries, while low-separability samples may serve as valuable bridges for cross-lingual alignment due to their entangled representations across languages. Moreover, we demonstrate that language separability not only informs data selection, but also enables effective curriculum learning: interleaving training samples with diverse separability levels allows models to absorb both structured and entangled linguistic signals, improving training effectiveness. These findings highlight the importance of considering linguistic structure in multilingual training pipelines, and we hope they inspire future work on more linguistically aware approaches in multilingual community.

2 Related Work

Joint Multilingual Instruction Fine-Tuning. While early work shows that fine-tuning on English-only data can yield cross-lingual generalization (Muennighoff et al. 2022), multilingual instruction tuning—i.e., fine-tuning on data from multiple languages—has emerged as a more widely adopted strategy for improving multilingual performance (Ouyang et al. 2022). Recent studies often incorporate data augmentation or distillation techniques

to enhance multilingual capabilities. For example, Li et al. (2023) addressed the issue of "translationese" by generating multilingual responses using Google Translate and ChatGPT. Some other approaches incorporate multilingual examples into English-centric tuning (Shaham et al. 2024), apply translation-based finetuning to improve semantic alignment (Ranaldi, Pucci, and Freitas 2024), combine translation data, cross-lingual tasks, and scaling laws (Zhu et al. 2023), or leverage self-distillation from high-resource languages (Zhang et al. 2024). Beyond data-level interventions, Ye et al. (2025) explored latent-level cross-lingual interactions via a cross-lingual connection mechanism. Empirical results consistently show that incorporating multilingual examples during instruction tuning leads to significantly stronger multilingual generalization compared to purely monolingual or English-centric approaches (Shaham et al. 2024; Chen et al. 2024).

Data Selection. A central goal of data selection in model fine-tuning is to minimize the amount of training data while maintaining—or even improving—model performance. Early work primarily focused on ensuring data quality and diversity. For example, Touvron et al. (2023) and Zhou et al. (2023) employed elaborate filtering pipelines to emphasize the role of high-quality supervision. Data diversity has also gained increasing attention, as it is crucial for training robust and generalizable models (Bukharin et al. 2024). Lu et al. (2023) proposed measuring diversity via intention tags of instructions, while Yang et al. (2025) achieved diversity-driven data selection through sparse autoencoder. Other selection criteria include features such as sequence length (Zhao et al. 2024), task complexity (Xu et al. 2024), etc. Beyond data-intrinsic properties, a growing line of work adopts model-aware selection, identifying samples that align closely with the target domain based on gradient similarity or other interaction signals between training and target sets (Xia et al. 2024; Zhao et al. 2025).

While these approaches have advanced instruction tuning through data-level and model-level interventions, little attention has been paid to the intrinsic linguistic structure present in multilingual training data. Our work addresses this gap by investigating language separability as a guiding signal for multilingual data selection.

3 Measurement of Language Separability

To quantify how well the training samples from different languages are separated in representation space, we use the silhouette score (Rousseeuw 1987). Originally designed to assess clustering quality, this metric favors compact intra-group and distant inter-group structures, making it well-suited for evaluating language separability.

Let $D = \{d_1, d_2, \dots, d_L\}$ denote a multilingual supervised instruction dataset, where each $d_l = \{p_i^l\} = \{(x_i^l, y_i^l)\}$ represents a language-specific subset corresponding to language $l \in \{1, \dots, L\}$. Here, each sample pair p_i^l consists of an input instruction x_i^l and its corresponding ground-truth response y_i^l . Given a model, we obtain the representation for each data point by formatting the instruction–response pair (x_i^l, y_i^l) using the model’s training template, feeding it into

the model, and extracting the hidden state of the last token. All representations are then grouped by language labels.

For each data point p_i^l , we compute its average distance $a(p_i^l)$ to all other points in the same language cluster d_l :

$$a(p_i^l) = \frac{1}{|d_l| - 1} \sum_{p_j^l \in d_l, j \neq i} \text{dist}(p_i^l, p_j^l) \quad (1)$$

where $|d_l|$ denotes the number of data points in d_l , and $\text{dist}(p_i^l, p_j^l)$ represents the Euclidean distance between p_i^l and p_j^l in their representation space. The value of $a(p_i^l)$ quantifies how well the sample p_i^l fits within its language cluster d_l (the smaller the value, the better the alignment).

Next, we measure the dissimilarity between p_i^l and other clusters by computing its average distance to the samples in its nearest neighboring cluster:

$$b(p_i^l) = \min_{m \neq l} \frac{1}{|d_m|} \sum_{p_j^m \in d_m} \text{dist}(p_i^l, p_j^m) \quad (2)$$

A larger $b(p_i^l)$ indicates that p_i^l is more distinct from samples in other language clusters.

The silhouette score for each data point p_i^l jointly considers intra-cluster compactness $a(p_i^l)$ and inter-cluster separation $b(p_i^l)$, and is defined as:

$$s(p_i^l) = \frac{b(p_i^l) - a(p_i^l)}{\max\{a(p_i^l), b(p_i^l)\}} \quad (3)$$

The silhouette score ranges from -1 to 1 , with higher values indicating that the sample p_i^l is well-clustered and clearly separated from other languages in the representation space. Our language separability guided data pre-selection strategy prioritizes the top $\rho\%$ of samples with the highest silhouette scores within each language cluster, encouraging the model to form and maintain clearer linguistic boundaries by training on selected highly separable data.

4 Experiments

4.1 Setups

Models. We selected two representative LLMs for our main experiments: (1) *LLaMA-3.1-8B* (Dubey et al. 2024) and (2) *Qwen2.5-7B* (Yang et al. 2024).

Training Corpus. We totally select 97,696 multilingual instruction pairs from *aya dataset* (Singh et al. 2024) as our full training corpus and the training corpus covers 31 languages, ensuring extensive multilingual coverage (see detailed statistics in Appendix). Our training processes are conducted on $8 * A100-80GB$ GPUs with the following settings: *batch size=16*, *epochs=3*, *learning rate=1.0e-5*, *warmup ratio=0.1*, and *bf16=true*. The implementation is based on *LLaMA-Factory* (Zheng et al. 2024).

Evaluation Datasets. We conduct experiments on 6 benchmarks, which can be categorized into:

- **Multilingual Understanding:** (1) *XNLI* (Conneau et al. 2018), a multilingual natural language inference (NLI) dataset, (2) *XStoryCloze* (Lin et al. 2022), a multilingual

commonsense reasoning dataset for evaluating story understanding and (3) *MMMLU*, the multilingual version of *MMLU* (Hendrycks et al. 2020), designed to evaluate models’ general knowledge.

- **Multilingual Generation:** (1) *MKQA* (Longpre, Lu, and Daiber 2021), an open-domain multilingual question answering evaluation dataset, (2) *XQuAD* (Artetxe, Ruder, and Yogatama 2020), a question answering dataset and (3) *XLSum* (Hasan et al. 2021), a multilingual abstractive summarization benchmark comprising professionally annotated article-summary pairs.

For each dataset, we evaluate on 10 languages, covering a total of 22 languages. *Accuracy* metric is used for *XNLI*, *XStoryCloze*, *MMMLU*, *MKQA* and *XQuAD* datasets. And for *XLSum* dataset, *ROUGE-L* scores are reported. We use greedy decoding with a max of 40 new tokens for each model. More details are provided in Appendix.

Baselines. Detailed implementations are in Appendix.

- **Random (*Rand*):** samples are randomly selected from the full training set. Results are averaged over three different random seeds to ensure robustness.
- **KMeans Clustering (*KMC*)** performs kmeans in the embedding space of model *all-MiniLM-L6-v2* and selects the closest sample to each cluster centroid.
- **Lexical Diversity (*MTLD*)** (McCarthy and Jarvis 2010): selects samples with higher lexical diversity.
- **Naturalness (*Nat*)** (Zhong et al. 2022) selects samples that better resemble natural human-written text.
- **Coherence (*Coh*)** (Zhong et al. 2022) selects samples where the response serves as a coherent continuation of the paired question.
- **Understandability (*Und*)** (Zhong et al. 2022) selects samples that are more understandable.
- **Importance Resampling (*DSIR*)** (Xie et al. 2023) selects samples most relevant to the target set by estimating their importance weights.
- **Gradient Similarity (*LESS*)** (Xia et al. 2024): selects samples most relevant to the target set by low-rank gradient similarity search.

The target set for the target-dependent baselines *DSIR* and *LESS* is constructed by sampling 3 examples from each language subset across all 6 benchmarks.

Detailed Implementations of *LangGPS*. As a pre-selection strategy, *LangGPS* first selects the top $\rho\%$ of samples with the highest language separability scores from each language cluster in the original multilingual training set (with $\rho = 20$ in our main experiments). These selected samples are then passed to existing data selection methods for further fine-grained filtering. In our main experiments, we apply *LangGPS* on top of both *Random* selection and the three best-performing baselines for each model. We consider data selection settings of 1%, 3%, and 5%—i.e., the final training set contains 1%, 3%, or 5% of the full data—and report the average performance across these three settings.




Figure 2: Average relative gains or declines (in percentage) when applying *LangGPS* on top of *Rand* and the three strongest-performing baselines, shown separately for high-resource and low-resource languages.

4.2 Main results

Table 1 reports average results across languages for each dataset, and Figure 2 shows the average relative gains or declines on high- and low-resource languages. The detailed results for each language can be found in Appendix.

(1) Existing Methods Are Not Universally Effective. Existing feature-based and target-dependent data selection methods fail to deliver consistent improvements in joint multilingual instruction tuning. For example, the trending method *LESS* achieves the best performance on *LLaMA-3.1-8B* ($\Delta = +4.90\%$), yet ranks the worst on *Qwen2.5-7B* ($\Delta = -7.35\%$). Moreover, *Nat*, *Coh*, and *Und*—the three text quality-based selection methods—perform significantly worse on understanding tasks compared to generation tasks. And on *Qwen2.5-7B*, we observe that all baselines underperform even the *Random* selection strategy. These observations highlight the challenge of designing universally effective selection strategies in multilingual settings.

(2) *LangGPS* Yields a Performance Boost, Especially on Understanding Tasks and Low-Resource Languages. Overall, applying *LangGPS* on top of existing selection methods achieves consistent performance improvements, as evidenced by the increase in Δ values. And the gains are more pronounced on understanding tasks. Besides, we also observe that for *LLaMA-3.1-8B*, which exhibits weaker inherent multilingual capabilities, both the baselines and *LangGPS* yield more substantial improvements compared to *Qwen2.5-7B*, whose multilingual competence is already stronger. In Figure 2, we present the average performance gains or declines when applying *LangGPS*, separately for high-resource and low-resource languages¹. Overall, *LangGPS* brings positive improvements on both language groups and the gains are more pronounced on low-resource lan-

Model: LLaMA-3.1-8B																									
Method	Multilingual Understanding												Multilingual Generation												
	XNLI				XStoryCloze				MMMLU				MKQA				XQuAD				XLSum		Δ		
Full	36.6			64.2			37.7			15.1			60.8			23.3		%	1%	3%	5%	Avg.			
(%)	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%						3%	5%	Avg.
Rand	25.1	33.0	35.4	31.2	43.7	61.4	61.1	55.4	34.2	37.4	38.1	36.5	14.6	14.4	14.4	14.5	63.6	59.9	59.6	61.1	21.9	22.8	22.8	22.5	-6.44%
<i>Feature-based Baselines</i>																									
KMC*	24.3	35.5	34.1	31.3	44.1	63.7	65.6	57.8	38.1	39.6	37.2	38.3	15.1	14.8	15.0	15.0	59.7	58.9	59.3	59.3	21.0	23.1	22.4	22.2	-5.18%
MTLD	28.8	34.9	34.4	32.7	34.3	52.9	47.1	44.8	25.1	32.5	29.9	29.1	15.7	15.9	14.5	15.4	60.3	62.5	57.2	60.0	19.5	21.9	24.2	21.8	-11.55%
Nat	20.4	40.1	36.2	32.2	13.3	40.0	54.8	36.0	24.5	35.3	37.6	32.4	18.5	17.3	17.1	17.6	59.3	61.3	59.5	60.0	19.6	20.9	22.1	20.8	-10.80%
Coh*	28.3	37.1	36.3	33.9	42.1	54.5	51.5	49.4	34.6	35.9	36.4	35.6	18.2	18.1	18.4	18.2	64.1	64.6	64.8	64.5	22.4	24.0	23.8	23.4	-1.47%
Und	23.5	38.1	35.0	32.2	25.3	47.5	59.6	44.1	28.8	35.0	38.4	34.1	18.2	16.9	17.4	17.5	60.6	60.6	59.5	60.2	20.6	21.2	22.2	21.3	-7.70%
<i>Target-dependent Baselines</i>																									
DSIR	29.5	27.7	36.9	31.4	36.7	39.7	49.5	42.0	27.3	30.1	35.1	30.8	17.7	16.9	16.4	17.0	64.7	63.4	62.4	63.5	22.1	23.1	23.5	22.9	-8.66%
LESS*	33.9	33.7	37.8	35.1	65.5	71.2	68.3	68.3	40.4	40.6	40.8	40.6	18.2	18.0	17.2	17.8	65.2	65.6	64.8	65.2	21.2	21.7	22.8	21.9	+4.90%
<i>Applying LangGPS (to random selection and three best well-performing baselines marked with *)</i>																									
Rand	26.4	34.7	36.2	32.5	48.7	72.5	72.5	64.6	33.7	38.2	39.7	37.2	15.4	15.0	14.7	15.0	62.6	60.2	62.0	61.6	21.9	22.8	23.8	22.9	-2.21%
				(+1.3)				(+9.2)				(+0.8)				(+0.5)				(+0.5)				(+0.4)	
KMC	26.1	38.0	37.6	33.9	42.4	66.3	69.2	59.3	38.3	40.2	38.8	39.1	15.3	15.1	13.8	14.7	61.8	60.3	60.6	60.9	21.5	22.2	23.8	22.5	-2.86%
				(+2.6)				(+1.5)				(+0.7)				(-0.3)				(+1.6)				(+0.3)	
Coh	33.7	32.8	38.1	34.9	47.1	47.2	60.8	51.7	37.6	35.4	37.5	36.9	19.1	18.2	17.6	18.3	63.4	62.6	63.7	63.2	22.2	24.3	24.3	23.6	-0.01%
				(+1.0)				(+2.3)				(+1.3)				(+0.1)				(-1.3)				(+0.2)	
LESS	35.6	37.8	37.7	37.0	66.1	76.8	74.0	72.3	39.9	40.0	42.2	40.7	17.1	16.5	17.0	16.9	65.6	64.3	64.3	64.7	22.5	22.8	23.4	22.9	+6.37%
				(+1.9)				(+4.0)				(+0.1)				(-0.9)				(-0.5)				(+1.0)	

Model: Qwen2.5-7B																									
Method	Multilingual Understanding												Multilingual Generation												
	XNLI				XStoryCloze				MMMLU				MKQA				XQuAD				XLSum		Δ		
Full	50.8			64.1			43.2			14.3			62.0			21.8		%	1%	3%	5%	Avg.			
(%)	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%	3%	5%	Avg.	1%						3%	5%	Avg.
Rand	45.6	53.0	54.7	51.1	54.2	76.0	72.7	67.7	49.3	45.6	44.1	46.4	15.5	15.1	14.5	15.0	65.4	64.0	63.9	64.4	22.1	22.2	22.7	22.3	+4.18%
<i>Feature-based Baselines</i>																									
KMC*	51.1	49.8	51.3	50.8	74.3	64.4	73.0	70.6	43.9	43.6	46.8	44.8	13.1	14.2	13.4	13.6	61.2	62.0	62.4	61.9	22.5	22.0	21.8	21.9	+1.52%
MTLD*	50.1	53.9	51.5	51.8	64.2	73.6	72.0	69.9	46.5	46.1	43.7	45.4	15.2	14.2	14.6	14.7	65.5	63.5	62.8	63.9	21.7	22.6	22.4	22.2	+4.03%
Nat	37.9	42.4	49.1	43.1	56.7	51.3	58.8	55.6	46.5	44.1	42.5	44.4	15.2	14.5	14.4	14.7	59.2	62.2	61.5	60.9	21.4	20.9	20.5	20.9	-4.72%
Coh	38.4	50.3	44.0	44.2	63.7	61.4	53.6	59.6	40.7	44.5	38.0	41.0	15.3	14.7	14.5	14.8	60.2	61.8	61.1	61.0	21.7	21.8	22.4	22.0	-3.66%
Und	39.5	48.8	49.3	45.9	49.1	57.7	58.5	55.1	43.8	45.1	42.2	43.7	14.9	14.4	14.7	14.7	58.3	61.8	62.2	60.8	21.1	20.3	20.9	20.8	-4.44%
<i>Target-dependent Baselines</i>																									
DSIR*	53.5	51.4	47.2	50.7	67.3	65.9	65.4	66.2	45.9	45.0	45.4	45.4	16.4	13.6	14.2	14.7	62.1	61.6	61.7	61.8	21.6	21.3	22.5	21.8	+2.61%
LESS	35.9	41.9	54.2	44.0	56.9	21.1	72.7	50.2	42.9	41.1	45.2	43.0	13.9	14.4	12.8	13.7	65.0	64.2	68.0	65.7	18.0	19.4	20.9	19.5	-7.35%
<i>Applying LangGPS (to random selection and three best well-performing baselines marked with *)</i>																									
Rand	48.8	52.0	54.9	51.9	59.1	72.2	73.2	68.2	48.8	46.2	44.2	46.4	15.4	14.8	14.5	14.9	64.7	65.0	64.9	64.9	22.0	21.7	22.4	22.0	+4.33%
				(+0.8)				(+0.5)				(+0.03)				(-0.1)				(+0.5)				(-0.3)	
KMC	50.3	53.7	53.6	52.5	78.8	70.1	72.2	73.7	43.9	46.5	45.2	45.2	13.8	14.4	13.9	14.0	62.7	62.3	63.4	62.8	21.6	23.0	23.0	22.5	+4.32%
				(+1.7)				(+3.1)				(+0.4)				(+0.4)				(+0.9)				(+0.6)	
MTLD	51.8	52.7	53.7	52.7	69.8	70.5	68.4	69.6	47.2	47.6	43.3	46.1	15.2	15.0	14.9	15.0	65.5	64.3	63.9	64.6	21.1	21.8	22.4	21.8	+4.64%
				(+0.9)				(-0.3)				(+0.7)				(+0.3)				(+0.7)				(-0.4)	
DSIR	49.8	54.1	50.6	51.5	62.1	67.2	69.8	66.4	48.9	50.2	46.9	48.7	15.9	14.8	15.0	15.2	64.2	63.7	64.1	64.0	21.0	22.2	22.5	21.9	+4.62%
				(+0.8)				(+0.2)				(+3.3)				(+0.5)				(+2.2)				(+0.1)	

Table 1: Main results averaged over all languages for each dataset. **Avg.** denotes the average performance across the 1%, 3%, and 5% training settings. For each model, *LangGPS* is applied on top of both random selection and the three best well-performing baselines (with highest Δ values) under that model. **Blue cell** indicates better performance than the vanilla baseline under the same training setting, while **Gray cell** indicate a performance drop. Underline numbers indicate the best performance among vanilla baselines, **Bold** numbers indicate the best performance achieved after applying *LangGPS*. Δ represents the average relative gain or decline (in percentage) of each method across all datasets compared to training on the full dataset.

$$\Delta(\text{Method}) = \frac{1}{|\text{AllDatasets}|} \sum_{d \in \text{AllDatasets}} \frac{\text{Result}(\text{Method}, d) - \text{Result}(\text{Full}, d)}{\text{Result}(\text{Full}, d)}$$

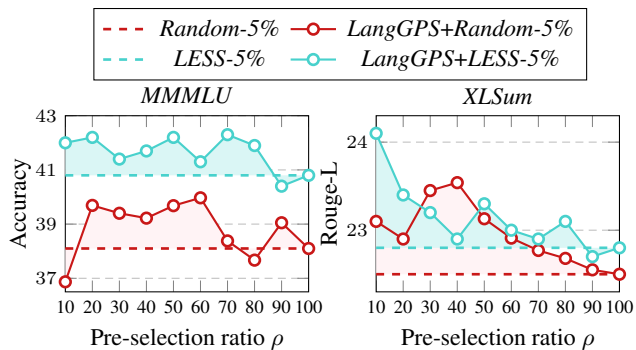


Figure 3: Effect of the pre-selection ratio ρ on performance (MMMLU and XLSum, 5% setting, LLaMA-3.1-8B). We vary ρ from 10% to 100% and report results using both *Rand* and *LESS* as downstream selectors.

guages, indicating that *LangGPS* can help mitigate the performance gap between high- and low-resource languages.

5 Further Analysis

5.1 Analysis of the Pre-selection Ratio ρ

We vary the pre-selection ratio ρ from 10% to 100% and report the performance on MMMLU, XLSum in Figure 3 (5% setting, LLaMA-3.1-8B) using both *Rand* and *LESS* as downstream selectors. Note that when $\rho = 100%$, *LangGPS+baseline* becomes equivalent to the vanilla baseline.

We observe that *LangGPS* is not highly sensitive to the choice of ρ , and generally performs well when ρ is in the range of 20% to 70%. A very small ρ (e.g., 10%) tends to over-focus on highly separable samples, leading to reduced diversity (as shown in Figure 1 (b)), and potential performance drops. On the other hand, as ρ approaches 100%, the pre-selection effect of *LangGPS* gradually converges toward those of the vanilla baselines. Moreover, the choice of ρ also presents a trade-off between data diversity and selection efficiency. A smaller ρ means that subsequent selection methods—especially those with higher computational cost (e.g., *LESS*)—operate on a much smaller candidate pool, reducing overhead at the cost of potentially reduced diversity.

Note that computational cost is analyzed in Appendix.

5.2 Analysis of Linguistic Boundary Representation

To provide a more intuitive understanding of *LangGPS*, we employ t-SNE (Van der Maaten and Hinton 2008) to visualize the representations of 200 sentences sampled from XNLI in parallel across English, Chinese and Arabic.

As shown in Figure 5 (1) (2) (3), multilingual SFT leads to clearer linguistic boundaries and more compact clusters

¹We follow the taxonomy in <https://microsoft.github.io/linguisticdiversity/assets/lang2tax.txt>, where languages rated 4 or 5 are considered relatively high-resource, and others low-resource. And since the 10 languages selected from *MKQA* are not typical low-resource, *MKQA* dataset is excluded from the figure.

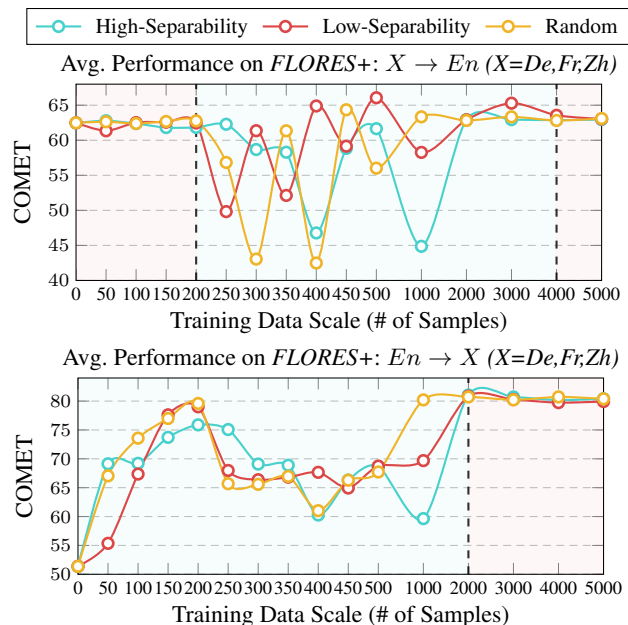


Figure 4: Average translation performance of LLaMA-3.1-8B on FLORES+ ($X \rightarrow En$ and $En \rightarrow X$, where $X = De, Fr, Zh$) under varying training sizes, comparing models trained on top separability-scoring, bottom separability-scoring, and randomly selected samples.

in the model’s representation space, reflected by increased average silhouette scores. Moreover, comparing Figure 5 (2) vs. (4) and (3) vs. (5), we observe that applying *LangGPS* further amplifies this effect: models trained with highly separable samples exhibit even higher silhouette scores, suggesting that *LangGPS* helps the model better form and maintain clear language boundaries.

5.3 The Role of Low-Separability Samples

Our experiments show that prioritizing highly separable samples during training improves multilingual performance. **But are low-separability samples entirely unhelpful?**

Low-separability samples often entangle with other language clusters in representation space, lacking clear language-specific features—but they may instead serve as implicit bridges for cross-lingual alignment and information sharing. To investigate this, we revisit the setup in Figure 1 (a) and Appendix, training models with (1) top separability-scoring, (2) bottom separability-scoring, and (3) randomly selected samples (with varying sizes), and evaluate them on the FLORES+ translation task to assess the role of low-separability data in cross-lingual scenarios.

The results in Figure 4, together with Figure 1(a), reveal several observations in common: (1) The three settings differ notably when training with limited samples, but their performances converge as training data increases. (2) Low-separability samples lead to a severe cold-start problem—when trained on just 50 examples, models perform significantly worse than those trained on high-separability or random data, likely due to the lack of clear language-specific




Figure 5: t-SNE visualizations of output representations from *LLaMA-3.1-8B* under different training settings: before SFT, after SFT with *Rand-5%*, *LESS-5%*, *LangGPS+Rand-5%*, and *LangGPS+LESS-5%*. Average silhouette scores are also reported.

signals in a small set of entangled samples.

Furthermore, we observe that the trends differ significantly between the “ $X \rightarrow En$ ” and “ $En \rightarrow X$ ” translation directions. (1) In “ $X \rightarrow En$ ”, none of the training settings yields improvement in the stable phase (red shaded area)—suggesting that vanilla multilingual SFT offers limited benefit for $X \rightarrow En$ translation. Nevertheless, low-separability samples result in the least degradation and occasionally yield marginal gains (e.g., at 400, 500, or 3000 samples). (2) In contrast, the overall trend of “ $En \rightarrow X$ ” direction shows clear gains from multilingual SFT. And within the 200-2000 sample dip (the cyan-shaded area), model trained with low-separability samples experiences a milder decline compared to those trained with high-separability samples.

5.4 Language Separability Guided Multilingual Curriculum Learning

Curriculum learning shares an underlying intuition with data selection: both aim to improve model training by leveraging differences in sample utility. While data selection focuses on identifying and retaining a subset of training samples deemed beneficial, curriculum learning instead retains the full dataset but modulates the order in which the samples are presented. In essence, both strategies prioritize how the model interacts with training data—either by selecting what to learn from, or by deciding when to learn from it.

In this section, we explore whether language separability can serve as a guiding signal for multilingual curriculum learning. We design three curriculum strategies based on language separability: (1) *Ascending*: training progresses from low to high separability samples; (2) *Descending*: from high to low separability; (3) *Balanced*: training data are interleaved across different separability ranges to maintain an approximately uniform separability distribution throughout training (implementations and algorithm are in Appendix).

The results in Table 2 show that naively presenting data in strictly ascending or descending separability order leads to unsatisfactory performance. Instead, the *Balanced* strategy, where samples of varying separability are mixed evenly throughout training, yields the most stable and generalizable improvements across different models and datasets. This

Datasets	XNLI	XStory Cloze	MMMLU	MKQA	XQuAD	XLSum
Strategy <i>Model: LLaMA-3.1-8B</i>						
<i>Random</i>	33.9	56.8	36.8	11.1	59.8	22.1
<i>Ascending</i>	37.9	67.5	32.5	11.0	58.7	21.3
<i>Descending</i>	33.1	63.4	33.2	10.8	<u>61.0</u>	22.2
<i>Balanced</i>	<u>36.8</u>	68.8	34.4	12.1	61.0	22.0
Strategy <i>Model: Qwen2.5-7B</i>						
<i>Random</i>	45.3	74.3	48.6	11.8	61.4	21.5
<i>Ascending</i>	36.4	76.1	44.6	11.5	<u>61.8</u>	21.5
<i>Descending</i>	44.5	61.6	45.1	12.2	61.7	<u>21.9</u>
<i>Balanced</i>	50.8	<u>75.9</u>	49.7	<u>11.9</u>	62.7	22.3

Table 2: Results of multilingual curriculum learning. We compare the three curriculum strategies—*Ascending*, *Descending*, and *Balanced*—with *Randomly Shuffled*.

finding suggests that maintaining a diversity of separability levels during training helps the model benefit from both highly structured (language-specific) and entangled (cross-lingual) samples, striking a better trade-off.

6 Conclusion

We present *LangGPS*, a lightweight and broadly compatible data pre-selection framework that leverages language separability as the guiding signal for multilingual data selection. By quantifying how well different languages are distinguished in the model’s representation space, *LangGPS* prioritizes highly separable samples that help models form clearer linguistic boundaries. Extensive experiments show that applying *LangGPS* on top of existing selection methods improves their effectiveness and generalizability in multilingual training, especially for understanding tasks and low-resource languages. Further analysis highlights the complementary roles of high- and low-separability samples, and demonstrates that language separability also benefits multilingual curriculum learning. Our findings position language separability as a principled signal for organizing multilingual data, offering a new perspective for improving the effectiveness and generalization of multilingual LLMs.

Acknowledgements

Xiaocheng Feng, Xiachong Feng and Dandan Tu are the co-corresponding authors of this work. We thank the anonymous reviewers for their insightful comments. This work was supported by the National Natural Science Foundation of China (NSFC) (grant 62522603, 62276078, U22B2059), the Key R&D Program of Heilongjiang via grant 2022ZX01A32, the Fundamental Research Funds for the Central Universities (XNJKKGYDJ2024013) and the State Key Laboratory of Micro-Spacecraft Rapid Design and Intelligent Cluster (MS01240122).

References

- Artetxe, M.; Ruder, S.; and Yogatama, D. 2020. On the Cross-lingual Transferability of Monolingual Representations. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4623–4637. Online: Association for Computational Linguistics.
- Bukharin, A.; Li, S.; Wang, Z.; Yang, J.; Yin, B.; Li, X.; Zhang, C.; Zhao, T.; and Jiang, H. 2024. Data Diversity Matters for Robust Instruction Tuning. *arXiv:2311.14736*.
- Chen, P.; Ji, S.; Bogoychev, N.; Kutuzov, A.; Haddow, B.; and Heafield, K. 2024. Monolingual or Multilingual Instruction Tuning: Which Makes a Better Alpaca. In Graham, Y.; and Purver, M., eds., *Findings of the Association for Computational Linguistics: EACL 2024*, 1347–1356. St. Julian’s, Malta: Association for Computational Linguistics.
- Conneau, A.; Rinott, R.; Lample, G.; Williams, A.; Bowman, S. R.; Schwenk, H.; and Stoyanov, V. 2018. XNLI: Evaluating Cross-lingual Sentence Representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hasan, T.; Bhattacharjee, A.; Islam, M. S.; Mubasshir, K.; Li, Y.-F.; Kang, Y.-B.; Rahman, M. S.; and Shahriyar, R. 2021. XL-Sum: Large-Scale Multilingual Abstractive Summarization for 44 Languages. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 4693–4703. Online: Association for Computational Linguistics.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Huo, W.; Feng, X.; Huang, Y.; Fu, C.; Li, B.; Ye, Y.; Zhang, Z.; Tu, D.; Tang, D.; Lu, Y.; et al. 2025. Enhancing Non-English Capabilities of English-Centric Large Language Models Through Deep Supervision Fine-Tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 24185–24193.
- Li, H.; Koto, F.; Wu, M.; Aji, A. F.; and Baldwin, T. 2023. Bactrian-x: Multilingual replicable instruction-following models with low-rank adaptation. *arXiv preprint arXiv:2305.15011*.
- Li, M.; Zhang, Y.; Li, Z.; Chen, J.; Chen, L.; Cheng, N.; Wang, J.; Zhou, T.; and Xiao, J. 2024. From Quantity to Quality: Boosting LLM Performance with Self-Guided Data Selection for Instruction Tuning. In Duh, K.; Gomez, H.; and Bethard, S., eds., *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 7602–7635. Mexico City, Mexico: Association for Computational Linguistics.
- Lin, X. V.; Mihaylov, T.; Artetxe, M.; Wang, T.; Chen, S.; Simig, D.; Ott, M.; Goyal, N.; Bhosale, S.; Du, J.; Pasunuru, R.; Shleifer, S.; Koura, P. S.; Chaudhary, V.; O’Horo, B.; Wang, J.; Zettlemoyer, L.; Kozareva, Z.; Diab, M.; Stoyanov, V.; and Li, X. 2022. Few-shot Learning with Multilingual Generative Language Models. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 9019–9052. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Longpre, S.; Lu, Y.; and Daiber, J. 2021. MKQA: A Linguistically Diverse Benchmark for Multilingual Open Domain Question Answering. *Transactions of the Association for Computational Linguistics*, 9: 1389–1406.
- Lu, K.; Yuan, H.; Yuan, Z.; Lin, R.; Lin, J.; Tan, C.; Zhou, C.; and Zhou, J. 2023. # instag: Instruction tagging for analyzing supervised fine-tuning of large language models. *arXiv preprint arXiv:2308.07074*.
- McCarthy, P. M.; and Jarvis, S. 2010. MTL-D, vocd-D, and HD-D: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2): 381–392.
- Muennighoff, N.; Wang, T.; Sutawika, L.; Roberts, A.; Biderman, S.; Scao, T. L.; Bari, M. S.; Shen, S.; Yong, Z.-X.; Schoelkopf, H.; et al. 2022. Crosslingual generalization through multitask finetuning. *arXiv preprint arXiv:2211.01786*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Qin, L.; Chen, Q.; Zhou, Y.; Chen, Z.; Li, Y.; Liao, L.; Li, M.; Che, W.; and Yu, P. S. 2024. Multilingual large language model: A survey of resources, taxonomy and frontiers. *arXiv preprint arXiv:2404.04925*.
- Ranaldi, L.; Pucci, G.; and Freitas, A. 2024. Empowering cross-lingual abilities of instruction-tuned large language models by translation-following demonstrations. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 7961–7973. Bangkok, Thailand: Association for Computational Linguistics.
- Rousseeuw, P. J. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20: 53–65.

- Shaham, U.; Herzig, J.; Aharoni, R.; Szpektor, I.; Tsarfaty, R.; and Eyal, M. 2024. Multilingual Instruction Tuning With Just a Pinch of Multilinguality. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 2304–2317. Bangkok, Thailand: Association for Computational Linguistics.
- Singh, S.; Vargus, F.; Dsouza, D.; Karlsson, B. F.; Mahendiran, A.; Ko, W.-Y.; Shandilya, H.; Patel, J.; Mataciunas, D.; OMahony, L.; Zhang, M.; Hettiarachchi, R.; Wilson, J.; Machado, M.; Moura, L. S.; Krzemiński, D.; Fadaei, H.; Ergün, I.; Okoh, I.; Alaagib, A.; Mudannayake, O.; Alyafeai, Z.; Chien, V. M.; Ruder, S.; Guthikonda, S.; Alghamdi, E. A.; Gehrmann, S.; Muennighoff, N.; Bartolo, M.; Kreutzer, J.; Üstün, A.; Fadaee, M.; and Hooker, S. 2024. Aya Dataset: An Open-Access Collection for Multilingual Instruction Tuning. *arXiv preprint arXiv:2402.06619*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Xia, M.; Malladi, S.; Gururangan, S.; Arora, S.; and Chen, D. 2024. LESS: Selecting Influential Data for Targeted Instruction Tuning. In *International Conference on Machine Learning (ICML)*.
- Xie, S. M.; Santurkar, S.; Ma, T.; and Liang, P. S. 2023. Data selection for language models via importance resampling. *Advances in Neural Information Processing Systems*, 36: 34201–34227.
- Xu, C.; Sun, Q.; Zheng, K.; Geng, X.; Zhao, P.; Feng, J.; Tao, C.; Lin, Q.; and Jiang, D. 2024. WizardLM: Empowering large pre-trained language models to follow complex instructions. In *The Twelfth International Conference on Learning Representations*.
- Yang, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Li, C.; Liu, D.; Huang, F.; Wei, H.; Lin, H.; Yang, J.; Tu, J.; Zhang, J.; Yang, J.; Yang, J.; Zhou, J.; Lin, J.; Dang, K.; Lu, K.; Bao, K.; Yang, K.; Yu, L.; Li, M.; Xue, M.; Zhang, P.; Zhu, Q.; Men, R.; Lin, R.; Li, T.; Xia, T.; Ren, X.; Ren, X.; Fan, Y.; Su, Y.; Zhang, Y.; Wan, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; and Qiu, Z. 2024. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*.
- Yang, X.; Nie, S.; Liu, L.; Gururangan, S.; Karn, U.; Hou, R.; Khabsa, M.; and Mao, Y. 2025. Diversity-driven data selection for language model tuning through sparse autoencoder. *arXiv preprint arXiv:2502.14050*.
- Ye, J.; Tao, X.; and Kong, L. 2023. Language versatilities vs. specialists: An empirical revisiting on multilingual transfer ability. *arXiv preprint arXiv:2306.06688*.
- Ye, Y.; Feng, X.; Feng, X.; Ma, W.; Qin, L.; Xu, D.; Yang, Q.; Liu, H.; and Qin, B. 2024a. GlobeSumm: A Challenging Benchmark Towards Unifying Multi-lingual, Cross-lingual and Multi-document News Summarization. In Al-Onaizan, Y.; Bansal, M.; and Chen, Y.-N., eds., *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 10803–10821. Miami, Florida, USA: Association for Computational Linguistics.
- Ye, Y.; Feng, X.; Feng, X.; Qin, L.; Huang, Y.; Huang, L.; Ma, W.; Hong, Q.; Zhang, Z.; Lu, Y.; et al. 2024b. Exploring Cross-lingual Latent Transplantation: Mutual Opportunities and Open Challenges. *arXiv preprint arXiv:2412.12686*.
- Ye, Y.; Feng, X.; Yuan, Z.; Feng, X.; Qin, L.; Huang, L.; Ma, W.; Huang, Y.; Zhang, Z.; Lu, Y.; Yan, X.; Tang, D.; Tu, D.; and Qin, B. 2025. CC-Tuning: A Cross-Lingual Connection Mechanism for Improving Joint Multilingual Supervised Fine-Tuning. *arXiv:2506.00875*.
- Zhang, X.; Li, S.; Hauer, B.; Shi, N.; and Kondrak, G. 2023. Don't Trust ChatGPT when your Question is not in English: A Study of Multilingual Abilities and Types of LLMs. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 7915–7927. Singapore: Association for Computational Linguistics.
- Zhang, Y.; Wang, Y.; Liu, Z.; Wang, S.; Wang, X.; Li, P.; Sun, M.; and Liu, Y. 2024. Enhancing Multilingual Capabilities of Large Language Models through Self-Distillation from Resource-Rich Languages. *arXiv preprint arXiv:2402.12204*.
- Zhao, H.; Andriushchenko, M.; Croce, F.; and Flammarion, N. 2024. Long is more for alignment: A simple but tough-to-beat baseline for instruction fine-tuning. *arXiv preprint arXiv:2402.04833*.
- Zhao, Y.; Du, L.; Ding, X.; Ouyang, Y.; Wang, H.; Xiong, K.; Gao, J.; Sun, Z.; Xu, D.; Qing, Y.; et al. 2025. Beyond similarity: A gradient-based graph method for instruction tuning data selection. *arXiv preprint arXiv:2502.11062*.
- Zheng, Y.; Zhang, R.; Zhang, J.; Ye, Y.; Luo, Z.; Feng, Z.; and Ma, Y. 2024. LlamaFactory: Unified Efficient Fine-Tuning of 100+ Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*. Bangkok, Thailand: Association for Computational Linguistics.
- Zhong, M.; Liu, Y.; Yin, D.; Mao, Y.; Jiao, Y.; Liu, P.; Zhu, C.; Ji, H.; and Han, J. 2022. Towards a Unified Multi-Dimensional Evaluator for Text Generation. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2023–2038. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Zhou, C.; Liu, P.; Xu, P.; Iyer, S.; Sun, J.; Mao, Y.; Ma, X.; Efrat, A.; Yu, P.; Yu, L.; et al. 2023. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36: 55006–55021.
- Zhu, W.; Lv, Y.; Dong, Q.; Yuan, F.; Xu, J.; Huang, S.; Kong, L.; Chen, J.; and Li, L. 2023. Extrapolating large language models to non-english by aligning languages. *arXiv preprint arXiv:2308.04948*.