

# CrossIn: An Efficient Instruction Tuning Approach for Cross-Lingual Knowledge Alignment

Geyu Lin<sup>♡</sup>, Bin Wang<sup>♡</sup>, Zhengyuan Liu<sup>♡</sup>, Nancy F. Chen<sup>♡, †</sup>

<sup>♡</sup>Institute for Infocomm Research (I<sup>2</sup>R), A\*STAR, Singapore

<sup>†</sup>Centre for Frontier AI Research (CFAR), A\*STAR, Singapore

{lin\_geyu, wang\_bin, liu\_zhengyuan, nfychen}@i2r.a-star.edu.sg

## Abstract

Multilingual proficiency presents a significant challenge for large language models (LLMs). English-centric models are usually suboptimal in other languages, particularly those that are linguistically distant from English. This performance discrepancy mainly stems from the imbalanced distribution of training data across languages during pre-training and instruction tuning stages. To address this problem, we propose a novel approach called CrossIn, which utilizes a mixed composition of cross-lingual instruction tuning data. Our method leverages the compressed representation shared by various languages to efficiently enhance the model’s task-solving capabilities and multilingual proficiency within a single process. In addition, we introduce a multi-task and multi-faceted benchmark to evaluate the effectiveness of CrossIn. Experimental results demonstrate that our method substantially improves performance across tasks and languages, and we provide extensive insights into the impact of cross-lingual data volume and the integration of translation data on enhancing multilingual consistency and accuracy.<sup>1</sup>

## 1 Introduction

The advancement of large language models (LLMs) like ChatGPT (Achiam et al., 2023) and Gemma (Team et al., 2023) has been a game-changer in the field of natural language processing (NLP), revolutionizing tasks such as language generation and commonsense reasoning (Naveed et al., 2024). Nevertheless, most state-of-the-art LLMs are English-centric, and their performance on non-English languages is usually suboptimal, especially on languages that are dissimilar to English (Blevins and Zettlemoyer, 2022; Mehrabi et al., 2022; Gao et al., 2024). This challenge mainly stems from the imbalanced distribution of multilingual data at both

the pre-training and instruction tuning stages. The exposure bias toward major languages results in an imbalanced capability, where models excel in languages with plentiful data while under-performing in those with limited resources (Dac Lai et al., 2023; Feng et al., 2023). Bridging the language gap is a fundamental step to unlock the full potential of these general-purpose models and ensure that the benefits are accessible to people across the linguistic spectrum (Zhu et al., 2023a).

Efforts to improve the multilingual capabilities of English-centric LLMs have included continue pre-training using extensive language-specific datasets. Yet, mastering languages through additional pre-training could require vast amounts of data and significant computational resources (Workshop et al., 2022). On the other hand, despite the limited proportion of non-English data at the pre-training stage, their absolute volume builds a solid knowledge base of various languages. In each iteration, LLMs are exposed to samples in several languages simultaneously, and the compressed representation encourages models to share linguistic features and generalize across different languages (Workshop et al., 2022). However, this ability is not fully retained through the use of datasets that only include English in follow-up tuning steps.

In this work, we propose an efficient approach based on a mixed composition of cross-lingual instruction tuning data to exploit LLMs’ underlying multilingual capability, which particularly improves the cross-lingual knowledge alignment (Qi et al., 2023; Wang et al., 2023). Instruction tuning is to boost the task solving capability of pre-trained language backbones (Taori et al., 2023; Luo et al., 2023; Touvron et al., 2023b). The task and prompt diversity are crucial in both data preparation and the training process, and a small high-quality set is sufficient to achieve state-of-the-art zero-shot performance (Ouyang et al., 2022). However, the language diversity of instruction tuning

<sup>1</sup>Our datasets and models will be released after the anonymity period.

is often overlooked in English-centric LLMs. We thus aim to enrich instruction tuning from the language perspective. Since all languages share the compressed representation space, cross-lingual instruction tuning can efficiently boost the model’s task-solving and multilingual capabilities within a single process. Unlike previous work that involved a multi-task setting by adding machine translation and mixing monolingual samples of each language (Zhu et al., 2023b), we integrate two languages at the sample level and combine various languages at the corpus level. Moreover, we compare various mixing strategies to identify the impact of different data formulations.

To extensively evaluate the cross-lingual knowledge alignment (Qi et al., 2023; Wang et al., 2023), we establish a benchmark of three tasks (i.e., reading comprehension, commonsense question-answering, and logic reasoning). Consistency is measured by analyzing an LLM’s responses to the same question in different languages, and our benchmark encompasses multiple ability aspects and difficulty levels. Moreover, since exact match and F1 score cannot precisely evaluate system outputs in the generative setting, we unify all three tasks in a multiple-choice format for quantitative and reproducible evaluation. The experimental results demonstrate that our mixed cross-lingual tuning can significantly improve performance in all aspects (up to 40% relative gain), followed by a detailed analysis of the influence of data quantity on language consistency and knowledge accuracy.

The main contributions of our research are:

- **A Multi-faceted Benchmark.** We present a multi-lingual, multi-capability benchmark for assessing the cross-lingual knowledge consistency of language models. In particular, we build a parallel multiple-choice version of the XQuAD dataset (Artetxe et al., 2019) - Cross-XQuAD for machine comprehension, and combining it with commonsense QA and logic reasoning.
- **Mixed Cross-Lingual Instruction Tuning.** We introduce CrossIn, a cross-lingual instruction tuning approach aimed at aligning knowledge across languages to stimulate the model’s full multilingual capability after pre-training. It offers a more efficient way of improving the model’s performance in various linguistic contexts.
- **CrossIn Data Insights.** We conduct extensive experiments with representative LLMs on three tasks, and show the effectiveness of our proposed approach. We provide detailed analysis to study the optimal amount of cross-lingual data and the necessity of sample translation in enhancing models’ cross-lingual consistency.

## 2 Related Work

### 2.1 Multilingual Large Language Model

Multilingual Large Language Models (MLLMs) have experienced significant advancements in recent years. Recently, Qin et al. (2024), as a comprehensive review, summarizes various methodologies for training MLLMs. BLOOM (Workshop et al., 2022), Jais (Sengupta et al., 2023), and Sailor (Dou et al., 2024) are representative models that target improved multilingualism in the pretraining stage. For fine-tuning, ChatGLM employs a reward model trained under a multilingual setting (Zeng et al., 2022), while the x-LLM utilizes a translated version of the Alpaca dataset, combined with supervised translation data and instruction finetuning, to enhance the model’s multilingual capabilities (Zhu et al., 2023b).

Instruction tuning on English datasets can introduce zero-shot capabilities in other languages as well (Wei et al., 2022; Chung et al., 2022). Further studies have explored the use of diverse training sets in multiple languages can improve cross-lingual generalization, suggesting that incorporating data from various languages can significantly enhance the model’s ability to generalize across linguistic boundaries (Muennighoff et al., 2023; Kew et al., 2023; Shaham et al., 2024). In our work, we build upon these findings and focus on improving multilingual consistency through targeted instruction finetuning. By refining the instruction processing mechanism, we aim to enforce the alignment across different languages to improve multilingual capabilities.

### 2.2 Multilingual Evaluation Benchmark

Evaluating the multilingual capabilities of LLMs is crucial for their global applicability, as it ensures that these models can understand and generate text effectively across different languages. Benchmarks such as MMLU (Hendrycks et al., 2021), TruthfulQA (Lin et al., 2021) have been developed to access the general capability of the




Figure 1: An illustration of the dataset construction process of the Cross-XQuAD dataset. The original XQuAD dataset, although multilingual, is not adapted specifically to evaluate LLMs and their cross-lingual consistency.

LLMs in English. XQuAD (Artetxe et al., 2019) and MLQA (Lewis et al., 2019) are popular extractive question-answering datasets that have been developed to evaluate the models’ multilingual performance. However, they focus on language-specific performance without considering the knowledge-sharing capabilities. Recently, Cross-MMLU and Cross-LogiQA (Wang et al., 2023) are proposed to assess the multilingual capability of LLMs with an emphasis on cross-lingual consistency. However, the number of samples is limited which could generally lead to less stable evaluation results.

### 3 Cross-Lingual Consistency Benchmark

Since traditional multilingual evaluations often fail to cater specifically to LLMs or overlook the assessment of cross-lingual consistency in multilingual contexts, in this section, we present a targeted multilingual evaluation benchmark for cross-lingual knowledge alignment.

#### 3.1 Datasets and Metrics

Even though there are multilingual evaluation datasets with parallel samples including MLQA (Lewis et al., 2019) and XQuAD (Artetxe et al., 2019), they are tailored for supervised extractive question-answering tasks and are unsuitable for less structured outputs of LLMs (Schuster et al., 2023). Therefore, recently, two evaluation datasets have been developed for multilingual evaluation with cross-lingual consistency measures (Wang et al., 2023). Specifically, Cross-MMLU and Cross-LogiQA are designed to use multiple-choice questions, presenting parallel samples to assess the knowledge alignment capability of LLMs. These datasets focus on commonsense question answering and logical reasoning. However, as they are

crafted by humans, the number of parallel samples they offer is relatively limited due to the high cost of human labor involved. This limitation could lead to less robust evaluation results.

Considering this, in our work, we enhance the cross-lingual consistency evaluation benchmark by introducing another task type: reading comprehension. Furthermore, we utilize existing high-quality parallel datasets to automatically generate new ones that are tailored for LLM evaluation. Table 1 summarizes the complete benchmark.

For evaluation metrics, we leverage the same concept as presented in Wang et al. (2023). In addition to assessing the overall accuracy of each language, we also integrate cross-lingual consistency metrics, measured by “Consistency” and “AC3”. The consistency score is designed to determine whether the model provides consistent responses to parallel questions across different languages. A higher consistency score suggests that LLMs can apply common knowledge across languages and deliver uniform responses, regardless of correctness. Specifically, for the Cross-XQuAD dataset that spans four languages, the multilingual consistency metric is defined as

$$M_{\{l_1, l_2, \dots, l_s\}} = \frac{\sum_{i=1}^N 1\{a_i^{l_1} = a_i^{l_2} = \dots = a_i^{l_s}\}}{N} \quad (1)$$

where  $a_i^{l_s}$  is the answer for sample index  $i$  from language  $s$ . Then, the consistency is computed as:

$$Consistency_s = \frac{\sum_{\{l_1, l_2, \dots, l_s\} \in C(s, g_i)} M_{\{l_1, l_2, \dots, l_s\}}}{C_4^s} \quad (2)$$

Similar to Wang et al. (2023), we use  $s = 3$  as the default tolerant for consistency metrics, where

Dataset	MCQs	Number of Samples	Supported Language	Consistency Metric
MLQA (Lewis et al., 2019)	✗	5,500 (36×)	7 - Eng, Zho, Spa, Vie, ...	NA
XQuAD (Artetxe et al., 2019)	✗	1,190 (7.9×)	10 - Eng, Zho, Spa, Vie, ...	NA
Cross-MMLU (Wang et al., 2023)	✓	150 (1×)	7 - Eng, Zho, Spa, Vie, ...	✓
Cross-LogiQA (Wang et al., 2023)	✓	176 (1.2×)	7 - Eng, Zho, Spa, Vie, ...	✓
Cross-XQuAD (ours)	✓	1,190 (7.9×)	4 - Eng, Zho, Spa, Vie	✓

Table 1: A list of multilingual datasets. Multi-choice questions (MCQs) are more suitable for quantitative evaluation of large language models and evaluation for multilingual consistency. Traditional metrics such as the F1 score or Exact Match for extractive question answering can introduce unintended biases in evaluating large language models.

the consistency between any three languages is computed. AC3 enhances the traditional accuracy metric by incorporating consistency, offering a more comprehensive evaluation. This approach is adopted because relying solely on consistency or accuracy does not yield a robust assessment.

$$AC3_s = 2 \cdot \frac{Accuracy \cdot Consistency_s}{Accuracy + Consistency_s} \quad (3)$$

By converting the datasets into MCQ (Multiple Choice Question) format, we can better quantify the model’s ability to select the correct answer from a set of options, thereby offering a clearer measure of its understanding and reasoning capabilities.

### 3.2 Cross-XQuAD Construction

Figure 1 indicates the process of constructing the Cross-XQuAD dataset from the original XQuAD dataset. It involves three steps, 1) English MCQ construction with distractive choices, 2) Parallel MCQ construction, and 3) Post-processing and quality check.

First, the original ground-truth answer from the XQuAD dataset can directly be used as the correction choice. As the XQuAD is for an extractive question-answer task, we extract the incorrect options from the provided context corpus as much as possible. Otherwise, the solution would be highly trivial with simple matching techniques. To achieve this, we prompt *ChatGPT-3.5* to get the other three choices as shown in Figure 1b.

Second, using the prepared English sample as a base, we prompt the generation of equivalent samples in the other languages. We discovered that direct translation without specific context can result in deviated interpretations due to polysemy, potentially leading to a biased evaluation. To counter this, we prompt the model with the English sample alongside its contextual counterpart in the target language to generate new samples. This approach has resulted in samples that are highly aligned across multiple languages.

Third, although LLMs can perform as a reasonable automated method for creating parallel samples (Li et al., 2023), we found that human intervention is essential to ensure higher accuracy. Consequently, each constructed sample undergoes a round of human review to confirm its integrity.

Following the above procedure, we construct the Cross-XQuAD dataset with 1,190 parallel samples in four languages which results in 4,760 samples in total. It is by far the largest multilingual evaluation dataset with cross-lingual consistency assessment capabilities.

## 4 CrossIn Method

To address the imbalance across languages in English-centric LLMs pre-training and fine tuning, we explore strategies to enhance multilingual proficiency through cross-lingual instruction tuning. Typically, instruction tuning relies on monolingual training samples (e.g., English) (Zhu et al., 2023a), which limits the potential of massive multilingual exposure at the pre-training stage. Therefore, we aim to enrich instruction tuning from the language perspective. Since all languages share the compressed representation space, cross-lingual instruction tuning can efficiently boost the model’s task-solving and multilingual capabilities within a single process. Specifically, we propose CrossIn approach to mix language compositions at the sample level to enforce the information flow across languages. We hypothesize that our method can further boost multilingual performance by encouraging language-level generalization.

The training data can be divided into three main aspects: **Base**, **CrossIn**, **Trans**.

- **Base**: This part includes the foundational instruction tuning datasets, where the model is relying on to learn all basic capabilities. Given that English datasets are the most resource-rich and of the highest quality, we focus



Type	Instruction	Output
Original Sample	Explain what a "food desert" is and why they are a problem.	A "food desert" is an area with little or no access to affordable, healthy food options....
CrossIn <sub>en2x</sub>	Explain what a "food desert" is and why they are a problem.	“食物沙漠”是一个几乎没有或根本没有负担得起的健康食品选择的地区...
CrossIn <sub>x2en</sub>	解释什么是“食物沙漠”，以及为什么它们是一个问题	A "food desert" is an area with little or no access to affordable, healthy food options....
CrossIn <sub>x2x</sub> (zho-spa)	解释什么是“食物沙漠”，以及为什么它们是一个问题	Un "desierto alimentario" es un área con poco o ningún acceso a opciones de alimentos saludables y asequibles...
CrossIn <sub>x2x</sub> (eng-spa)	Explain what a "food desert" is and why they are a problem.	Un "desierto alimentario" es un área con poco o ningún acceso a opciones de alimentos saludables y asequibles...
CrossIn <sub>x2x</sub> (spa-zho)	Explique qué es un "desierto alimentario" y por qué son un problema.	“食物沙漠”是一个几乎没有或根本没有负担得起的健康食品选择的地区...
Translation	Translate the following sentence into English. 解释什么是“食物沙漠”，以及为什么它们是一个问题	Explain what a "food desert" is and why they are a problem.

Table 2: One example from the Alpaca dataset. It is further transformed into cross-lingual instruction tuning datasets and translation tasks.

on using common English instruction tuning datasets.

- **CrossIn**: It comprises cross-lingual instruction tuning datasets, where instruction and output are featured in two different languages. The dataset aims to align representations across languages, thus enforcing the knowledge-sharing ability of LLMs.
- **Trans**: It consists of translation pairs for instructions. We hypothesize that if the model concurrently learns these translation tasks, it could facilitate the transfer of knowledge between languages.

For **Base**, we leverage existing datasets, we create the **CrossIn** and **Trans** datasets, where we use the Alpaca (Taori et al., 2023) dataset as the source. Examples are shown in Table 2.

For **CrossIn** dataset, we create three variants as the following recipes:

- CrossIn<sub>en2x</sub>: Instructions are provided in English, and we choose the output language randomly. Given the rich prior knowledge available in English, this approach aims to transfer English knowledge to other languages.
- CrossIn<sub>x2en</sub>: Instruction language is chosen randomly, and output is fixed in English. This approach aims to unify multilingual instructions into responses centered around English.
- CrossIn<sub>x2x</sub>: The languages for both the instruction and the output are selected randomly. This approach seeks to facilitate bi-directional alignment across all languages.

---

#### Algorithm 1 CrossIn<sub>x2x</sub> with translation

---

```

 $\mathcal{S} \leftarrow$  Total number of samples
 $\mathcal{L} \leftarrow$  {"English", "Spanish", "Chinese", "Vietnamese"}
 $\mathcal{D} \leftarrow$  Seed Parallel Instructions Dataset
 $\mathcal{C} \leftarrow \emptyset$ 
 $\mathcal{T} \leftarrow \emptyset$ 
 $t_p \leftarrow$  Translation Prompt
for  $i \leftarrow 1$  to  $\mathcal{S}$  do
   $s \leftarrow$  Random sample from  $\mathcal{D}$ 
   $l_{in}, l_{ot} \leftarrow$  Random sample from  $\mathcal{L}$ 
   $\mathcal{C} \leftarrow \mathcal{C} \cup (\mathcal{D}[l_{in}][s], \mathcal{D}[l_{ot}][s])$ 
   $l_t \leftarrow$  Random sample from  $\mathcal{L}$ 
   $\mathcal{T} \leftarrow \mathcal{T} \cup (t_p, \mathcal{D}[l_t][s], \mathcal{D}["English"][s])$ 
end for

```

---

Previous work shows that incorporating sample translation helps map English to other languages, allowing the model to generalize English knowledge in a broader space (Zhu et al., 2023b). For an extensive comparison, we also investigate how adding a separate translation task might enhance the multilingual abilities of LLMs, compared with using cross-lingual instruction tuning alone. More specifically, aside from the CrossIn data, we add a direct translation task of instructions from English to other languages. The influence on model performance of additional instruction translation is discussed in Section 5.3.

Algorithm 1 illustrates the complete algorithm to create CrossIn<sub>x2x</sub> with translation dataset, where  $\mathcal{S}$  is the desired number of samples to be added with the **Base**.  $\mathcal{C}$ ,  $\mathcal{T}$ ,  $l_{in}$  indicate **CrossIn**, **Trans** and the sampled language, respectively.

Models	Cross-XQuAD			Cross-MMLU			Cross-LogiQA		
	Acc	Consis	AC3	Acc	Consis	AC3	Acc	Consis	AC3
<b>General LLMs</b>									
<i>ChatGPT-3.5</i>	90.6	83.7	87.0	66.8	51.8	58.4	53.3	40.5	46.0
<i>LLaMA-2-7B-Chat</i> (Touvron et al., 2023b)	74.9	67.5	71.1	40.1	42.0	41.1	36.8	43.5	39.9
<i>Mistral-7B-Instruct-v0.2</i> (Jiang et al., 2023)	84.6	72.2	77.9	49.0	26.2	34.1	46.0	38.5	41.9
<i>LLaMA-7B</i> (Touvron et al., 2023a)	40.3	21.5	28.0	29.8	27.8	28.8	27.6	23.0	25.1
<i>m-LLaMA-7B</i> (Zhu et al., 2023b)	46.8	41.1	43.8	26.7	22.3	24.3	28.1	22.0	24.7
<b>Base Model: Gemma-2B</b> (Team et al., 2024)									
<i>Tuning w/ Alpaca</i>	42.0	49.7	45.5	36.0	59.8	45.0	28.3	63.8	39.2
<i>Tuning w/ Platypus</i>	60.8	55.8	58.2	36.5	29.7	32.7	36.4	47.9	41.3
CrossIn <sub>en2x</sub>	<b>60.1</b>	62.8	<b>61.5</b>	39.2	43.0	41.0	39.5	37.8	38.6
CrossIn <sub>x2en</sub>	54.2	<b>64.7</b>	59.0	<b>41.2</b>	<b>57.8</b>	<b>48.1</b>	36.8	<b>48.3</b>	41.8
CrossIn <sub>x2x</sub>	53.3	64.3	58.3	37.0	54.5	44.1	<b>39.6</b>	46.2	<b>42.6</b>
<b>Base Model: Mistral-7B-v0.1</b> (Jiang et al., 2023)									
<i>Tuning w/ Alpaca</i>	62.2	52.9	57.2	36.2	43.5	39.5	35.7	33.8	34.7
<i>Tuning w/ Platypus</i>	61.1	33.2	43.0	38.8	20.2	26.5	47.9	29.8	36.8
CrossIn <sub>en2x</sub>	74.9	64.0	69.0	41.0	41.5	41.2	44.6	40.1	42.2
CrossIn <sub>x2en</sub>	77.4	63.8	69.9	34.8	<b>47.2</b>	40.1	45.3	42.5	43.8
CrossIn <sub>x2x</sub>	<b>78.6</b>	<b>67.9</b>	<b>72.9</b>	<b>41.0</b>	42.3	<b>41.7</b>	<b>48.9</b>	<b>48.3</b>	<b>48.6</b>

Table 3: Experimental results on three cross-lingual consistency datasets: Cross-XQuAD, Cross-MMLU, Cross-LogiQA. Three metrics presented are Accuracy (ACC), Consistency (Consis), and AC3 as introduced in Section 3.

## 5 Experiments

### 5.1 Experimental Setting

In our experiments, we selected four languages: English, Chinese, Vietnamese, and Spanish across all three datasets. We utilized two representative open LLM as base model: *Mistral-7B-v0.1* (Jiang et al., 2023) and *Gemma-2B* (Team et al., 2024). For base models, we employed the Platypus (Lee et al., 2023) corpus as the **Base** dataset for instruction tuning, since previous work shows that it can enable models’ higher diverse and robust generalization capabilities than the Alpaca dataset.

For the **CrossIn** instruction tuning data, we utilize the Alpaca (Taori et al., 2023) corpus as the seed dataset. This dataset is expanded into a multilingual format to four languages using an off-the-shelf translation engine, producing a total of (52k×4) samples. From the enriched datasets, both the **CrossIn** and **Trans** parts can be formulated in a variant number of samples. While the Alpaca dataset lacks the complex problem-solving capabilities of the **Base** set from Platypus, it contains English instructions without complex elements like coding and math, which results in a higher translation quality. Meantime, this setup allows us to investigate whether a dataset of simple instructions can adequately support effective knowledge alignment across languages.

In model training, we leverage LoRA (Hu et al., 2022) with  $rank = 64$  as a parameter-efficient way to train LLMs. For fair comparison, we fine-tune base models with either the Platypus or Alpaca dataset with the same set of hyperparameters. Besides, following standard benchmarks, we also compared several representative general-purpose LLMs including *ChatGPT-3.5*, *LLaMA-2-7B-Chat*, *Mistral-7B-Instruct-v0.2*, *m-LLaMA-7B* and its base model, *LLaMA-7B*.

### 5.2 Main Results and Analysis

Table 3 shows the benchmark results of current general LLMs and models tuned with Alpaca, Platypus and different CrossIn variants. Our findings can be summarized as follows.

**English-centric LLMs do not perform well on our multi-lingual benchmark.** First, we evaluate the performance of representative LLMs using our benchmarks and observed that *ChatGPT-3.5* exhibits outstanding performance across all three test-sets, indicating strong multilingual capabilities and consistency. For open-source models, we observe that models after instruction tuning (e.g., *LLaMA-2-7B-Chat*, *Mistral-7B-Instruct-v0.2*) significantly outperform the non-tuned models (e.g., *LLaMA-7B*, *m-LLaMA-7B*) on all fronts, while their accuracy and cross-lingual consistency lag behind that of *ChatGPT-3.5*. Moreover, *m-LLaMA-7B* demon-

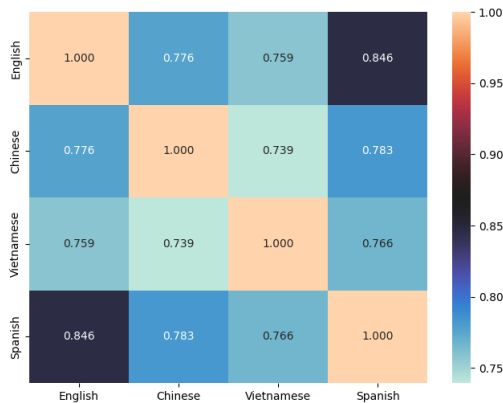


Figure 2: Consistency score between languages on Cross-XQuAD with CrossIn<sub>x2x</sub> method

strated some improvements over *LLaMA-7B* in the Cross-XQuAD dataset, but it only managed to achieve similar results on the Cross-MMLU and Cross-LogiQA. This suggests that a purely monolingual data mix may not be adequate for training models on complex multilingual tasks, highlighting the importance of our proposed approach.

**English-centric instruction tuning is limited.** We analyzed the performance of base models fine-tuned on different original instruction datasets (i.e., Alpaca and Platypus). Our findings indicate that models exhibit distinct characteristics depending on the instruction tuning corpus. Fine-tuning with Platypus results in higher accuracy, potentially due to the diversity of tasks in the dataset. Conversely, models fine-tuned with Alpaca shows a higher consistency across most benchmark datasets, albeit with marginally lower accuracy. These observations suggest that Alpaca may be less effective than Platypus in augmenting LLMs with task-solving and reasoning. In addition, fusing a wide range of knowledge in English could potentially lead to a forgetting of information in other languages, thus affect the consistency. This results show a trade-off between accuracy and consistency from fine-tuning on different English-centric instruction tuning corpora. We aim to bridge the gap of both datasets, thereby enhancing both accuracy and consistency.

**CrossIn is simple but effective.** We further review the results from our CrossIn instruction tuning method, which leverages the strengths of both the English-centric Platypus and the diverse Multilingual Alpaca datasets. By implementing the CrossIn augmentation, we successfully raised the AC3 score by 30% on the Cross-XQuAD bench-




Figure 3: Results of different cross-lingual instruction tuning methods compared with baseline.

mark and by about 12% on both the Cross-MMLU and Cross-LogiQA testsets. This improvement was achieved using the CrossIn<sub>x2x</sub> approach with the Mistral-7B-v0.1 as the foundational model. Enhancements were evident in the model’s accuracy and consistency across various languages, contributing to the higher AC3 scores. Our findings highlight the effectiveness of the CrossIn method in enriching a model’s performance on multilingual tasks. By starting with a task-diverse, strong instruction set from the Platypus dataset and integrating simpler, language-varied data from Alpaca, we crafted a cross-lingual knowledge base that significantly improve accuracy and consistency in multilingual understanding.

**Language discrepancy affects consistency.** We investigate the consistency scores across all pairs of languages. As shown in Figure 2, Spanish and English exhibit the highest consistency, potentially due to their linguistic similarities, among all other language pairs. On the other hand, Chinese and Vietnamese have the lowest correlation, which may be attributed to their completely distinct character sets. Apart from the linguistic discrepancies, this could also stem from language bias during the pre-training phase of language models. When considering the consistency score between English and other languages, Vietnamese, typically categorized as a low-resource language in pre-training, shows the least consistency with English. This points to the importance of diversifying the data used in training language models to ensure fair and effective language representation, particularly for languages that are typically categorized as low-resource.

### 5.3 Ablation Study

We conduct three comprehensive ablation studies to systematically assess the effects of various data formations, the integration of translation data, and




Figure 4: Comparison of AC3 score of adding translation data in cross-lingual instruction tuning.

the influence of different alignment dataset sizes on the performance of our models, aiming to identify key factors that enhance or inhibit their effectiveness.

**Data Formulation Comparison.** Figure 3 shows the AC3 scores from three tests when the language backbone is the Mistral-7B-v0.1. The results make it clear that methods designed for cross-lingual instructions work better than the basic method, which only uses English-centric instruction tuning data from Platypus or Alpaca. In particular, the CrossIn<sub>x2x</sub> method does much better than the CrossIn<sub>en2x</sub> and CrossIn<sub>x2en</sub> methods. This suggests that fully mixing multiple languages (CrossIn<sub>x2x</sub>) can make the most of what the Mistral-7B-v0.1 model offers by effectively using data from different languages. The mixed composition in training examples seems to help the model understand and apply knowledge from one language to another, leading to more accurate and consistent results.

**Efficacy of Translation Data.** Figure 4 compares the performance of the CrossIn<sub>x2x</sub> method with the CrossIn<sub>x2x\_T</sub> strategy, which adds translations to the Alpaca samples (as described in Algorithm 1). The experimental results indicate that additional translation pairs does not bring performance gains. We speculate that this is because tasks included in our benchmark focus on understanding and reasoning, and the cross-lingual instruction tuning approach stimulate both of them under a multilingual setting. Additionally, the translations used here may be too basic, especially compared to larger datasets like WikiMatrix. This suggests that improving multilingual knowledge alignment may be better achieved through a mixed-language approach at the sample level rather than by incorporating simple translation data.




Figure 5: Comparison of AC3 score by adding different numbers of CrossIn data. Base model: *Mistral-7B-v0.1*

**Essential Cross-Lingual Data Quantities.** Figure 5 shows the AC3 score of the LLMs with different quantity of cross-lingual alignment data. It can be shown that adding 5000 alignment data could already achieve a good result of cross-lingual consistency, there are not much improvement trend if we add more data. The observation that only a small amount of cross-lingual alignment data is required to achieve satisfactory consistency in LLMs can be attributed to its efficient learning mechanism. This characteristic allows the model to quickly assimilate and generalize from limited data, making it particularly adept at few-shot learning scenarios. Additionally, the model’s pretraining on diverse linguistic corpora might have already equipped it with a foundational understanding of various languages, thereby reducing the need for extensive alignment data to bridge linguistic gaps. This efficient use of data not only demonstrates the model’s robustness but also highlights its practicality in situations where data availability is constrained.

## 6 Conclusion

In this paper, we presented a study on improving cross-lingual knowledge alignment of multilingual large language models, and contributed to both evaluation benchmarks and methodologies. We built a machine comprehension dataset that is a robust resource for extensive multilingual evaluation, emphasizing cross-lingual consistency in compensation with previous datasets. Our cross-lingual instruction tuning method CrossIn brought significant improvements in knowledge accuracy and consistency across languages, highlighting the potential of efficient tuning to create more robust multilingual large language models.



## Limitations

Our approach depends on the availability of high-quality translation and cross-lingual data, which may not be accessible for all languages. Addressing these data availability challenges is essential for further research on enhancing multilingual consistency in large language models.

In this study, we did not examine the impact of our cross-lingual data formulation on the pretraining stage of large language models. Pre-training is crucial as it significantly shapes the model’s foundational knowledge and capabilities. Considering the larger scale of pretraining compared to fine-tuning, exploring whether our method could improve the efficiency and effectiveness of pretraining multilingual language models is a vital direction for future research. However, conducting such an ablation study on the pre-training stage is computationally demanding and may not be feasible with limited resources.

## Acknowledgment

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme, Ministry of Digital Development and Information (MDDI) and Infocomm Media Development Authority under National Large Language Models Funding Initiative, AI For Mother Tongue Language Learning and Public Sector Translational R&D Grant Funding Initiative (TRANSGrant). The aim of TRANSGrant is to tap on the research community to solve public sector challenges with innovative use of digital technologies. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore, Ministry of Digital Development and Information (MDDI) and Infocomm Media Development Authority.

## References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. [On the cross-lingual transferability of monolingual representations](#). *CoRR*, abs/1910.11856.

Terra Blevins and Luke Zettlemoyer. 2022. [Language](#)

[contamination helps explain the cross-lingual capabilities of english pretrained models](#).

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).

Viet Dac Lai, Chien Van Nguyen, Nghia Trung Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan A Rossi, and Thien Huu Nguyen. 2023. [Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback](#). *arXiv e-prints*, pages arXiv–2307.

Longxu Dou, Qian Liu, Guangtao Zeng, Jia Guo, Jiahui Zhou, Wei Lu, and Min Lin. 2024. [Sailor: Open language models for south-east asia](#).

Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023. [From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models](#).

Changjiang Gao, Hongda Hu, Peng Hu, Jiajun Chen, Jixing Li, and Shujian Huang. 2024. [Multilingual pre-training and instruction tuning improve cross-lingual knowledge alignment, but only shallowly](#).

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#).

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*.

Tannon Kew, Florian Schottmann, and Rico Sennrich. 2023. [Turning english-centric llms into polyglots: How much multilinguality is needed?](#)

Ariel N. Lee, Cole J. Hunter, and Nataniel Ruiz. 2023. [Platypus: Quick, cheap, and powerful refinement of llms](#).

Patrick Lewis, Barlas Oğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. [Mlqa: Evaluating cross-lingual extractive question answering](#). *arXiv preprint arXiv:1910.07475*.

- Minzhi Li, Taiwei Shi, Caleb Ziems, Min-Yen Kan, Nancy Chen, Zhengyuan Liu, and Diyi Yang. 2023. Coannotating: Uncertainty-guided work allocation between human and large language models for data annotation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1487–1505.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2023. [An empirical study of catastrophic forgetting in large language models during continual fine-tuning](#).
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. [A survey on bias and fairness in machine learning](#).
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning](#).
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2024. [A comprehensive overview of large language models](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Jirui Qi, Raquel Fernández, and Arianna Bisazza. 2023. [Cross-lingual consistency of factual knowledge in multilingual language models](#).
- Libo Qin, Qiguang Chen, Yuhang Zhou, Zhi Chen, Yinghui Li, Lizi Liao, Min Li, Wanxiang Che, and Philip S. Yu. 2024. [Multilingual large language model: A survey of resources, taxonomy and frontiers](#).
- Tal Schuster, Adam D. Lelkes, Haitian Sun, Jai Gupta, Jonathan Berant, William W. Cohen, and Donald Metzler. 2023. [Semqa: Semi-extractive multi-source question answering](#).
- Neha Sengupta, Sunil Kumar Sahu, Bokang Jia, Satheesh Katipomu, Haonan Li, Fajri Koto, Osama Mohammed Afzal, Samta Kamboj, Onkar Pandit, Rahul Pal, et al. 2023. Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models. *arXiv preprint arXiv:2308.16149*.
- Uri Shaham, Jonathan Herzig, Roei Aharoni, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2024. [Multilingual instruction tuning with just a pinch of multilinguality](#).
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrubti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#).
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, Ai Ti Aw, and Nancy F. Chen. 2023. [SeaEval for multilingual foundation models: From cross-lingual alignment to cultural reasoning](#).
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#).

BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.

Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.

Wenhao Zhu, Hongyi Liu, Qingxiu Dong, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2023a. [Multilingual machine translation with large language models: Empirical results and analysis](#).

Wenhao Zhu, Yunzhe Lv, Qingxiu Dong, Fei Yuan, Jingjing Xu, Shujian Huang, Lingpeng Kong, Jiajun Chen, and Lei Li. 2023b. [Extrapolating large language models to non-english by aligning languages](#).

## A Appendix

### A.1 Prompt for Building Cross-XQuAD Data

Initial Prompt:  
Given the following context: {context}.

Question: {question}  
Could you generate 4-multiple-choice with 3 incorrect options. Please try to make the incorrect options challenging by extracting them from the context as much as possible. No explanation. Just give me the 4 options. The options should start with (A), (B), (C), (D).  
The correct answer is: (A) {answer}

Figure 6: Prompt For Generating English Choice

Translation Prompt:  
Given the following context: {context}

Could you translate the above 4 choices into {language}. No explanation. Just give me the 4 options. The options should start with (A), (B), (C), (D).  
The correct answer is: (A) {answer}

Figure 7: Prompt to Translate English Choice

### A.2 Fine-tuning Parameters

Hyperparameter	Value
learning_rate	1e-4
batch_size	16
epochs	1
lora_rank	64
lora_alpha	128
lora_trainable	p_proj, k_proj, v_proj, o_proj, gate_proj, down_proj, up_proj
modules_to_save	embed_tokens, lm_head
lora_dropout	0.05
warmup_ratio	0.03
weight_decay	0
optimizer	Adam
bf16	True

Table 4: Fine-tuning Hyperparameters