

# ROMANSETU: Efficiently unlocking multilingual capabilities of Large Language Models via Romanization

Jaavid Aktar Husain<sup>1,2†§</sup> Raj Dabre<sup>1,5,8,9§</sup> Aswanth Kumar<sup>3</sup> Jay Gala<sup>4‡</sup>  
Thanmay Jayakumar<sup>1</sup> Ratish Puduppully<sup>6</sup> Anoop Kunchukuttan<sup>1,7,8§</sup>

<sup>1</sup>Nilekani Centre at AI4Bharat <sup>2</sup>IIT D&M Kancheepuram <sup>3</sup>Flipkart

<sup>4</sup>Mohamed bin Zayed University of Artificial Intelligence

<sup>5</sup>National Institute of Information and Communications Technology, Kyoto, Japan

<sup>6</sup>Institute for Infocomm Research (I<sup>2</sup>R), A\*STAR, Singapore

<sup>7</sup>Microsoft, India <sup>8</sup>Indian Institute of Technology Madras, India

<sup>9</sup>Indian Institute of Technology Bombay, India

## Abstract

This study addresses the challenge of extending Large Language Models (LLMs) to non-English languages that use non-Roman scripts. We propose an approach that utilizes the romanized form of text as an interface for LLMs, hypothesizing that its frequent informal use and shared tokens with English enhance cross-lingual alignment. Our approach involves the continual pretraining of an English LLM like Llama 2 (Touvron et al., 2023) on romanized text of non-English, non-Roman script languages, followed by instruction tuning on romanized data. The results indicate that romanized text not only reduces token fertility by 2x-4x but also matches or outperforms native script representation across various NLU, NLG, and MT tasks. Moreover, the embeddings computed on romanized text exhibit closer alignment with their English translations than those from the native script. Our approach presents a promising direction for leveraging the power of English LLMs in languages traditionally underrepresented in NLP. Our code is available on <https://github.com/AI4Bharat/romansetu>.

## 1 Introduction

Large Language Models (LLMs) demonstrate remarkable proficiency across a broad spectrum of natural language processing (NLP) tasks, as evidenced by various studies (Liu et al., 2023; Chung et al., 2022; Chowdhery et al., 2023; Wei et al., 2024; Goyal et al., 2022b). They excel not only

in tasks for which they were explicitly trained but also in those for which they were not trained. This achievement is mainly due to the availability of corpora (Wenzek et al., 2020; Abadji et al., 2021; Suárez et al., 2019) as well as the advancements in LLMs that leverage these datasets for pretraining (Touvron et al., 2023; Workshop et al., 2022; Chowdhery et al., 2023). Despite their proficiency in English, these models typically demonstrate reduced effectiveness when applied to non-English languages, highlighting a significant challenge in extending their benefits to non-English languages.

The English-heavy LLMs (Touvron et al., 2023; Jiang et al., 2023; Zhang et al., 2022) still have some representation coverage from other languages due to data leakage while creating the pre-training dataset, particularly for languages that use the same script as English *i.e.*, the Roman script. This script sharing enables cross-lingual transfer and bestows some of these LLM capabilities to these languages. For languages using non-Latin scripts, the data representation is very limited to non-existent. Tokenization of text in the languages exhibits high fertility (Ács, 2020) and byte-level representation (Artetxe et al., 2020) due to inadequate representation of these languages in the tokenizer vocabulary. Hence, these LLMs perform poorly on most of these languages, and inefficient tokenization also leads to high inference latency as well as the inability to process long sequences. This disparity in performance raises a critical question: *how can we extend the capabilities of LLMs to the languages written in non-Latin scripts?*

A widely explored solution is the extension of the tokenizer vocabulary to incorporate new languages and continual pre-training on native language data (Cui et al., 2023; Nguyen et al., 2023; Minixhofer et al., 2022). This approach is computationally demanding since the models need to

<sup>†</sup> Work done during internship at AI4Bharat

<sup>‡</sup> Work done during employment at AI4Bharat

<sup>§</sup> Corresponding Authors: Jaavid Aktar Husain ([jaavidaktar@gmail.com](mailto:jaavidaktar@gmail.com)), Raj Dabre ([raj.dabre@nict.go.jp](mailto:raj.dabre@nict.go.jp)) and Anoop Kunchukuttan ([ankunchu@microsoft.com](mailto:ankunchu@microsoft.com))

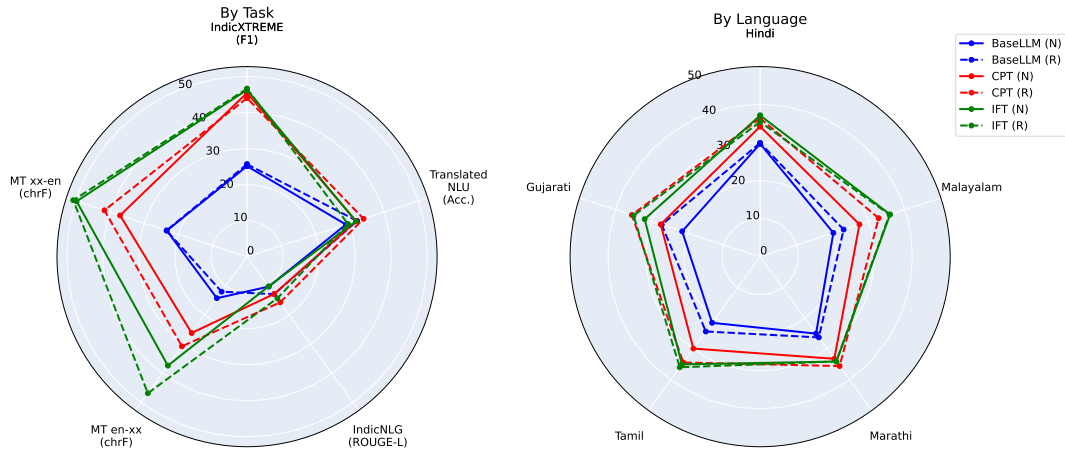


Figure 1: Performance of the BaseLLM (Llama 2 Base), our continually pretrained model (CPT), and our instruction finetuned model (IFT), in both native (N) and romanized (R) script settings. The average scores for different tasks (left) and languages (right) are compared in each radar chart. For CPT, we show 3-shot results where available, or else 1-shot results. For IFT, we use zero-shot results.

be continually pre-trained long enough to effectively learn the new embeddings and align the representations of English and the new language. Furthermore, this approach requires the availability of large volumes of text corpora.

In this work, we explore an alternative approach for more efficient knowledge transfer. Instead of utilizing the native script, we use the romanized form of text as the interface to the LLMs in our approach RomanSetu.<sup>1</sup> The adoption of romanized representation is justified for several reasons. In many languages, romanized text has been frequently used in informal settings and on social media in recent history. This usage creates the potential for the inclusion of some romanized data representation in the pre-training corpus. Additionally, code-mixing with English is a common occurrence and the romanized form shares tokens with English. This leads us to hypothesize that the romanized form is better aligned with English than the native script, thereby facilitating a more effective transfer from English.

In our approach, we first continually pre-train an English LLM like Llama 2 (Touvron et al., 2023) on romanized text for a language generated via transliteration. Subsequently, we perform instruction tuning on this continually pretrained model. We experiment with 5 languages from two language families and evaluate our approach on several NLU, NLG and MT benchmarks. Our experi-

ments and analysis indicate that:

- The fertility of romanized text is 2x-4x smaller than native text, making the romanized form far more efficient than the native script.
- The embeddings of the romanized text are closer to the embeddings of their English translations, compared to the embeddings of their native script equivalents, suggesting the former’s suitability for cross-lingual transfer.
- Results across several NLU, NLG, and MT tasks show that multilingual instruction finetuning on romanized data produces competitive or superior results than instruction finetuning on native script data, highlighting efficiency and better cross-lingual transfer. Figure 1 summarizes these results of various models on romanized and native inputs.
- Specifically, for generation tasks we see significant improvement due to utilization of romanization. To the best of our knowledge, ours is the first work that shows that romanization can help natural language generation tasks.
- Romanized representation can enable cross-lingual transfer in decoder-only English-heavy LLMs. Previous work has primarily focused on the use of romanization in multilingual, encoder-only models.

<sup>1</sup>The word *setu* in Sanskrit means *bridge*, referring to the Roman script serving as a bridge between English and other languages.

## 2 Related Work

Transliteration refers to the conversion of text written in one script to another. Romanization is a specific instance of transliteration, where the target script is the Roman/Latin script. Romanization has special significance since Roman script is by far the most widely adopted script in the world and many language models are primarily trained for English, which is written in the Roman script. Transliteration is typically used to represent different languages in the same script to enable cross-lingual transfer.

In NMT, [Amrhein and Sennrich \(2020\)](#) show that transliteration shows improvements for low-resource languages with different scripts by transferring from related high-resource languages that use different scripts. [Goyal et al. \(2020\)](#) show that transliteration helps even when only a contact relationship exists between the languages involved. [Song et al. \(2020\)](#) show that transliteration during the pretraining stage for NMT also helps cross-lingual transfer.

Transliteration has also been used for cross-lingual transfer in the context of pretrained language models. Some works ([Khemchandani et al., 2021](#); [Dhamecha et al., 2021](#); [Moosa et al., 2023](#); [Purkayastha et al., 2023](#)) employ transliteration to a common script during the pretraining phase to enable cross-lingual transfer. Other studies ([Dabre et al., 2022](#); [Muller et al., 2021](#); [Chau and Smith, 2021](#)) adopt transliteration during the fine-tuning. Transliteration could be done to a common non-Latin script ([Khemchandani et al., 2021](#); [Dhamecha et al., 2021](#); [Dabre et al., 2022](#); [Doddapaneni et al., 2023](#)) or to the Latin script ([Muller et al., 2021](#); [Moosa et al., 2023](#); [Purkayastha et al., 2023](#)).

While transliteration has been explored for language modeling as described above, our work differs from previous work in some important aspects:

- Previous work explored cross-lingual transfer using transliteration with multilingual language models. We focus our attention on English LLMs and try to achieve cross-lingual transfer via romanization using English LLMs. This is a challenging scenario since the language can be unrelated to English, and very little (if any) native or romanized data in the language under study might be seen during pretraining of the English LLM. At the same time, this is a very practical need since most best

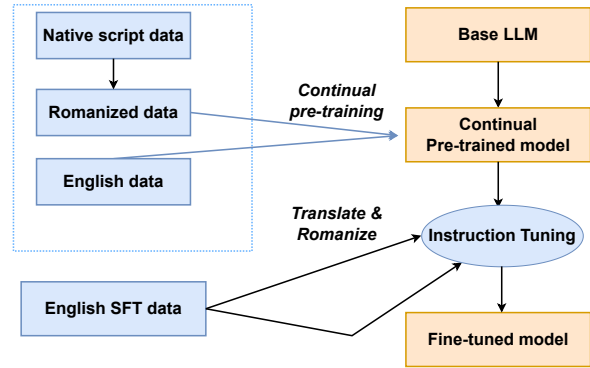


Figure 2: Overview of our proposed approach.

performing LLMs are English-heavy ([Touvron et al., 2023](#); [Jiang et al., 2023](#)) and hence cross-lingual transfer via romanization is an important direction to explore.

- We investigate the utility of transliteration in decoder-only language models, which is currently the standard architecture for LLMs. In contrast, all previous work explored cross-lingual transfer with transliteration in the context of encoder-only models (the exception is [Dabre et al. \(2022\)](#) which use encoder-decoder models).
- While previous work has explored the use of romanization for cross-lingual transfer in language understanding tasks, we show that romanization can be beneficial for generation tasks as well.

To the best of our knowledge, no prior research has investigated leveraging of romanization for mostly English LLMs in the context of cross-lingual transfer to non-English languages.

## 3 Utilizing Romanized Data to make LLMs Multilingual

Our proposed approach aims to enhance the capabilities of mostly English LLMs for non-English languages by leveraging romanized data. We first continually pretrain the LLM with both romanized and English data to create a base LM that is romanization-aware. Subsequently, we extend the process by instruction tuning the continually pretrained model. The framework of our approach is illustrated in Figure 2, encompassing the stages of data romanization, pretraining, and instruction tuning.

### 3.1 Romanization Scheme

A variety of romanization schemes are available, each driven by multiple considerations. One key factor is the resemblance of the romanized representation to the way people typically write romanized text. This is particularly advantageous if the pretraining data includes romanized text, as it might aid in aligning with English. Another important aspect is the fertility achieved by the romanization scheme when considering the original LLM’s tokenizer. A third consideration is whether the transliteration scheme is lossy or lossless. A lossless scheme is preferable when the objective is to convert the output back to the native script. Typically, deterministic transliteration schemes are lossless, whereas natural transliteration schemes tend to be lossy.

In this work, we have focused on Indic languages. Hence, we evaluated two romanization schemes for Indic languages: (a) the extended ITRANS scheme from the IndicNLP library (Kunchukuttan, 2020), which defines a fixed, reversible mapping between Indic scripts and Roman characters, and (b) the IndicXlit scheme (Madhani et al., 2023), which generates romanizations as commonly used by Hindi speakers in informal contexts and is learned from parallel transliteration corpora. These mappings are inherently lossy. The IndicXlit romanization demonstrates lower fertility compared to ITRANS romanization (See Table 1). Preliminary prompting experiments with the Llama 2 model also indicated that IndicXlit outperforms ITRANS in romanized Hindi to English translation. Consequently, we selected IndicXlit as our romanization scheme for this project. However, it is important to note that the transliterations are not reversible. With continued pretraining, ITRANS transliterations might eventually achieve similar task performance to IndicXlit while ensuring script reversibility. This direction might also be relevant for other languages which do not have a transliteration model available, while romanization mappings are available for almost all scripts. We leave this exploration for future work.

### 3.2 Continual Pretraining

While the base model is trained on roman script, most of the data it has been exposed to is in English. Hence, the base model cannot be used as-is for processing romanized inputs, particularly for generation into non-English languages. Hence, the base model needs to be updated. To make the base

model romanization-aware, we continue to pretrain the base LLM with romanized document-level data. To prevent any catastrophic forgetting of English capabilities, we also incorporate an equal amount of English data into the pretraining mix. We perform full-finetuning of the models. We do not expand or change the model vocabulary, hence the romanized inputs can take advantage of the token embeddings from the base model while adapting to the new languages.

### 3.3 Instruction Finetuning (IFT)

Non-English languages possess very limited finetuning data. A common method to overcome this limitation is to translate English IFT datasets into the languages. We could choose to translate both input and outputs (Wei et al., 2023), or translate only the inputs and ask a powerful, proprietary LLM like ChatGPT to generate translations in the native language (Li et al., 2023). We chose the former approach since the quality of LLM responses (even for the best commercial models) might not be of good quality (Ahuja et al., 2023), while high-quality open-source translation models (Costa-jussà et al., 2024; Gala et al., 2023) are available making translation of both inputs and outputs feasible at scale for many languages. The translated IFT datasets are then romanized using IndicXlit.

## 4 Experimental Settings

We conducted comprehensive experiments spanning multiple languages, benchmarks and models. This section elaborates on the experimental settings. In all our experiments, we experiment with Llama 2 7B (Touvron et al., 2023) as the base model under various settings.

### 4.1 Languages

We experimented with the following Indic languages: Hindi, Marathi, Gujarati, Tamil and Malayalam. These languages span 2 language families and 4 different scripts and are mid to low-resource languages. The languages considered in this study belong to the Indo-Aryan branch of Indo-European (Hindi, Marathi, Gujarati) and Dravidian (Tamil, Malayalam) families. In these languages, the typical word order is SOV. These languages have an inflectional morphology, with the Indo-Aryan leaning somewhat towards fusional, while Dravidian languages tend to have an agglutinative morphology. It is important to note that the script systems used for these languages are abugida, effectively



faithfully capturing vowels and consonants. We see the choice of Indian languages as a strength since it spans multiple language families, multiple scripts, and diverse linguistic characteristics and resource levels. Past work has also used Indic languages as a case study for romanized representation (Moosa et al., 2023; Dhamecha et al., 2021).

## 4.2 Datasets

**Continual Pretraining:** For continual pretraining, we sourced approximately 500 million words of document-level data from web-crawled corpora (Doddapaneni et al., 2023) for each language under consideration along with English. To generate the romanized dataset, we transliterated the native script dataset using the *IndicXlit* model (Madhani et al., 2023), a state-of-the-art open-source transliteration model for Indian languages. Both the native script dataset and its romanized counterpart were then used for continual pretraining in various configurations explained later.

**Instruction Fine-tuning:** These languages have very little native instruction tuning data for diverse tasks. Following Wei et al. (2023), we rely on translating high-quality English-supervised instruction-tuning datasets into the languages under consideration. We use IndicTrans2 (Gala et al., 2023), the state-of-the-art open-source MT model for Indian languages compared to commercial offerings for translation. We sampled examples from various English instruction tuning datasets to ensure a diverse mix of tasks from the Flan collection (Longpre et al., 2023) (65k) and Dolly (Conover et al., 2023) (15k). These are translated into all the languages. The instruction tuning dataset includes a translation subset as well to learn the translation task as well as to drive alignment between English and other languages. We sample 20k high-quality manually translated examples in each direction from the BPCC-Human subset (Gala et al., 2023). Thus, the final IFT dataset has 120k examples per language. The instruction tuning datasets are further romanized using IndicXlit.

**Evaluation Data:** We evaluate our models on a variety of NLU and NLG tasks. These include native language test sets from the IndicXTREME (Doddapaneni et al., 2023) and Indic NLG Suite (Kumar et al., 2022) benchmarks. For the translation task, we use the FLORES-200 devtest (Goyal et al., 2022a; Costa-jussà et al., 2024) for evaluation. Following recent work (Lai et al., 2023; Üstün et al., 2024; OpenAI et al., 2023), we eval-

uate the knowledge and reasoning capabilities of the model on the machine-translated versions of the benchmarks such as MMLU (Hendrycks et al., 2021a), ARC (Clark et al., 2018a), BoolQ (Clark et al., 2019a) and CommonSenseQA (Talmor et al., 2019). Note that as the CommonSenseQA benchmark comprises of a blind test set, we report the evaluations on the dev set. The translations were obtained using IndicTrans2 and the romanized inputs were created using IndicXlit.

## 4.3 Training and Finetuning Details

The models are fully finetuned during continual pretraining as well as supervised fine-tuning. Both the continued pretraining and instruction fine-tuning are multilingual i.e. a single model is trained jointly with data for all languages (including English to preserve English performance). Note that no vocabulary expansion is done. For pre-training, the loss is computed for all the tokens in the sequence. During instruction fine-tuning, the loss was computed only for the output tokens. Detailed information about the hyperparameters used for both continual pretraining and instruction fine-tuning can be found in Appendix A. We adapt *open-instruct*<sup>2</sup> (Wang et al., 2023) for our continual pretraining and fine-tuning experiments.

## 4.4 Trained Models

We begin with Llama 2 as the base model. Then, we create two variants that have undergone continual pre-training: one with native script data (CPT-N) and the other with romanized data (CPT-R). For each of these continually pre-trained models, we further perform instruction fine-tuning (IFT). We fine-tune CPT-N using native script IFT data to obtain IFT-N, and CPT-R is fine-tuned with romanized script IFT data to produce IFT-R.

## 4.5 Evaluation Metrics

The evaluation metric for machine translation in our experiments is chrF (Papineni et al., 2002) computed using the SacreBLEU toolkit (Post, 2018).<sup>3</sup> For other NLG tasks, we report Rouge-L (Lin, 2004) scores. *We primarily report evaluations on the native script outputs. In the case of romanized outputs, we transliterate them back to native script*

<sup>2</sup><https://github.com/allenai/open-instruct/tree/main>

<sup>3</sup>nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.3.1

using IndicXlit combined with a unigram word-level language model (Madhani et al., 2023). Combining word-level results with a unigram language model is efficient and preserves most gains of incorporating a complete n-gram language model (Kirov et al., 2024). In this native script evaluation, the roman script model outputs can also have transliteration errors, which put them at a disadvantage compared to the native script model. To purely compare task performance, we also report results on romanized outputs. For this, we romanize the native script outputs using IndicXlit. Since the final output is an artifact of the same romanization model, this evaluation compares only the task performance.

#### 4.6 Models

We use the 7B-parameter Llama 2 base model as our starting point for all experiments. We evaluate the following models: First, the base Llama 2 model (base) is assessed using both native script and romanized inputs. We then evaluate the continually pre-trained models with their respective script data; thus, CPT-N is evaluated with native script data, and CPT-R is evaluated with romanized data. Additionally, the instruction-tuned models are evaluated with their respective script data; therefore, IFT-N is assessed with native script data, and IFT-R is assessed with romanized data.

#### 4.7 Prompting and Decoding

For the base and CPT models, we do 1-shot and 3-shot prompting. For IFT models, we experiment with 0-shot prompting. The prompt templates for various tasks are shown in Appendix B. For generation tasks, we use greedy decoding. For classification tasks, we completed the prompts with various possible options as per class labels and computed the output probabilities. We choose the class that maximizes the probability (Liu et al., 2021).

### 5 Results

In this section, we present the results of various experiments comparing romanized and native representations and discuss the results. We aim to address the following research questions:

- Is there an efficiency gain through romanization, and if so, which approach to romanization is the most efficient?
- Do romanized inputs lead to embeddings that

Language	N	R	R-IndicNLP
Gujarati	18.44	<b>3.39</b>	4.16
Hindi	7.36	<b>2.98</b>	3.98
Malayalam	12.85	<b>5.04</b>	5.56
Marathi	8.91	<b>3.64</b>	4.84
Tamil	12.11	<b>4.89</b>	5.35

Table 1: Fertility scores on the FLORES-200 dev set (Goyal et al., 2022a; Costa-jussà et al., 2024) for native (N), romanized with IndicXLIT (R) and romanized with IndicNLP library (R-IndicNLP). Bold indicates lowest fertility.

are better aligned to English representations compared to native inputs?

- Does romanization improve downstream task performance across multiple NLP tasks?

**Efficiency Gains with Romanization:** To address this question, we compute fertility scores on the FLORES-200 dev set (Goyal et al., 2022a; Costa-jussà et al., 2024). Fertility refers to the number of tokens per word generated by the Llama 2 tokenizer (Touvron et al., 2023). We compare the native script against two forms of romanized text: transliteration using IndicXlit (Madhani et al., 2023) and rule-based extended ITRANS transliteration scheme implemented in the IndicNLP library (Kunchukuttan, 2020). Table 1 presents the fertility scores on the Llama 2 model for the languages we experimented with. We observed that the fertility (and hence sequence length) in the native script is at least twice that of the IndicXlit romanized text. In languages like Gujarati, the native script fertility is very large (greater than 4 times the romanized fertility). Hence, we can expect significant efficiency gains when processing romanized text over native text in terms of reduced memory consumption, reduced generation time and increased maximum sequence length limit.

Furthermore, romanization using IndicNLP (Kunchukuttan, 2020) has slightly higher fertility compared to IndicXlit. While we have used IndicXlit-based transliterations in this work for reasons explained earlier, it would be interesting to explore deterministic and reversible romanizations like IndicNLP at the cost of slightly increased inefficiency. The advantage of reversible mapping would be avoiding transliteration errors when mapping the outputs back to the native script.

**Cross-lingual Representation Similarity:** We compare the sentence embeddings generated by the

Language	E - N	E - R	N - R
Gujarati	0.39	0.47	0.51
Hindi	0.40	0.50	0.34
Malayalam	0.40	0.46	0.52
Marathi	0.44	0.48	0.58
Tamil	0.44	0.43	0.53

Table 2: Cosine similarities are computed using the last token representations of the last layer from Llama 2. In this context, E - N represents the cosine similarity between English and the native script, E - R represents the similarity between English and the romanized script and N - R represents the similarity between the native and the romanized script.

native and romanized inputs as well as English on the base model. Specifically, we compare cosine similarities between sentence embeddings generated by these inputs. Sentence embeddings are computed using the last token representation of the final hidden layer. In Table 2, we see that the Roman script representations are closer to English representations compared to the native script representations, making them a better candidate for cross-lingual transfer from English.

**Machine Translation:** We discuss the results of machine translation, a canonical multilingual language generation task. Table 3 (Rows XX-En translation and En-XX translation) shows results on machine translation both in and out of English.

For translation into English, the base model translation quality is roughly equivalent for both native and roman scripts. We see that continual pretraining helps improve translation quality significantly for both representations. Romanized scripts show significantly more improvement on the CPT models compared to native scripts. We hypothesize that romanization is better able to align cross-lingual representations between English and other languages to improve translation performance. Romanized IFT models also perform better than their native counterparts, albeit the gap is now narrower compared to CPT models.

For translation out of English, the base model performs very poorly which is understandable since the model might have seen only negligible non-English data. CPT significantly improves translation quality on both representations, with romanized representation showing better performance. The IFT-R models are able to take good advantage of the romanized representations, and they show an average improvement of 8 chrF points over the

	Script	BaseLLM		CPT		IFT
		1-shot	3-shot	1-shot	3-shot	0-shot
XX-En translation	N	22.46	23.42	38.54	37.07	49.78
	R	<b>22.52</b>	<b>23.52</b>	<b>42.53</b>	<b>41.64</b>	<b>50.75</b>
En-XX translation	N	13.95	<b>14.25</b>	25.55	26.19	37.40
	R	<b>14.20</b>	12.02	<b>29.55</b>	<b>30.77</b>	<b>46.87</b>
XLSum	N	6.88	-	7.59	-	7.77
	R	<b>10.16</b>	-	<b>12.44</b>	-	<b>12.56</b>
IndicHeadline	N	13.66	-	18.04	-	12.61
	R	<b>15.56</b>	-	<b>18.92</b>	-	<b>16.03</b>

Table 3: Results on NLG Benchmarks - averaged across languages (RougeL score for XLSum and IndicHeadline; chrF for MT. En-XX refers to out-of-English and XX-En refers to into-English directions. N and R indicate whether the model was trained and decoded on native or romanized script. Evaluation is always in the native script.

Language	Native-Script Eval		Roman-Script Eval	
	N	R	N	R
Gujarati	23.9	<b>26.0</b>	45.4	<b>53.5</b>
Hindi	44.1	<b>49.1</b>	45.9	<b>56.2</b>
Malayalam	39.4	<b>42.8</b>	49.0	<b>53.8</b>
Marathi	40.3	<b>44.8</b>	44.0	<b>50.7</b>
Tamil	39.3	<b>41.5</b>	50.1	<b>53.7</b>

Table 4: Native and Roman script evaluation results via chrF scores for En-XX translation on IFT model in 0-shot setting. N and R indicate whether the model was trained and decoded on native or romanized script.

corresponding native script IFT-N models. Thus, we show that even generation into romanized script can take advantage of the English capabilities of the LLM.

Table 4 shows the effect of the the choice of script when evaluating translation out of English to each language. We also report evaluation on Roman script outputs for both models and see that the romanized IFT models report even larger gains in this evaluation setting. This suggests that the true gains from romanization might be larger and it is not reflected in native script evaluation on account of errors in Roman to native transliteration that is performed for the evaluation.

**Language Generation Tasks:** Table 3 shows the results on other generation tasks *viz.* summarization and headline generation. Similar to MT, we observe that models utilizing romanized script consistently outperform those using native script by ~4 RougeL points across both continually pretraining and instruction fine-tuning setups with different shots. On summarization using the XLSum dataset (Hasan et al., 2021), we find that CPT and IFT models demonstrate competitive performance in

	Script	BaseLLM		CPT		IFT
		1-shot	3-shot	1-shot	3-shot	0-shot
IndicSentiment	N	45	<b>44.32</b>	70.58	87.32	<b>88.88</b>
	R	<b>48.16</b>	43.7	<b>88.46</b>	<b>92.82</b>	82.44
IndicCOPA	N	<b>21.08</b>	<b>25.43</b>	<b>20.47</b>	<b>35</b>	37.3
	R	17.53	8.12	18.23	18.78	<b>45.45</b>
IndicXNLI	N	16.98	21.1	18.68	35.76	<b>42.3</b>
	R	<b>27.24</b>	<b>31.16</b>	<b>36.8</b>	<b>37.28</b>	38.38
IndicQA (with context)	N	9.77	-	23.83	-	16.79 <sup>†</sup>
	R	<b>19.74</b>	-	<b>27.25</b>	-	<b>20.33</b> <sup>†</sup>

Table 5: Results on IndicXTREME NLU Testsets (average F1 score across languages). <sup>†</sup> indicates 1-shot as scores for 0-shot were close to zero and non-interpretable so we made an exception for this task. N and R indicate whether the model was trained and decoded on native or romanized script. Evaluation is always in the native script.

	Script	BaseLLM		CPT		IFT
		1-shot	3-shot	1-shot	3-shot	0-shot
MMLU	N	26.87	27.43	28.70	28.98	24.59
	R	<b>26.94</b>	<b>27.48</b>	<b>30.59</b>	<b>31.14</b>	<b>26.28</b>
BoolQ	N	57.90	44.45	59.09	47.48	<b>60.04</b>
	R	<b>60.75</b>	<b>61.03</b>	<b>61.85</b>	<b>57.10</b>	46.95
Arc Easy	N	25.82	26.35	<b>28.77</b>	<b>28.25</b>	<b>26.01</b>
	R	<b>26.38</b>	<b>26.44</b>	26.69	27.10	25.34
Arc Challenge	N	<b>26.54</b>	<b>25.70</b>	<b>27.56</b>	<b>27.20</b>	<b>24.20</b>
	R	24.61	25.56	26.40	26.66	22.80
CommonSenseQA	N	<b>21.01</b>	<b>20.45</b>	<b>26.55</b>	27.42	24.16
	R	20.07	20.18	25.55	<b>28.18</b>	<b>25.95</b>
IndicQA (without context)	N	1.48	-	5.56	-	1.13 <sup>†</sup>
	R	<b>12.96</b>	-	<b>17.62</b>	-	<b>16.87</b> <sup>†</sup>

Table 6: Results on Translated NLU Testsets (average Accuracy score across languages with the exception of IndicQA where use F1 score). <sup>†</sup> indicates scores reported for 1-shot as scores for 0-shot were close to zero and non-interpretable so we made an exception for this task. N and R indicate whether the model was trained and decoded on native or romanized script. Evaluation is always in the native script.

both native and romanized scripts. A bit surprisingly, we find that IFT models underperform by 3-5 RougeL points compared to CPT models on IndicHeadline. We hypothesize that the format of the example was not seen in the IFT dataset, resulting in poorer zero-shot performance compared to the one-shot performance of the CPT model. Overall, we find romanization to benefit not only in terms of efficiency but also in terms of downstream performance.

**Language Understanding Tasks:** Table 5 shows the results on different NLU tasks part of IndicXTREME benchmark (Doddapaneni et al., 2023), including IndicSentiment, IndicCOPA, and IndicXNLI. We find that romanized script models are competitive to those that utilize native script.

Specifically, in tasks such as IndicSentiment and IndicXNLI, romanized models consistently outperform their native counterparts across various shots. However, native models generally perform superior compared to romanized variants for the IndicCOPA task, except for the IFT model. Furthermore, we observe notable performance improvements for IndicQA without context (a reading comprehension task) across all models with romanized script when we provide the relevant passage in the context. This indicates that models with the romanized script tend to benefit from knowledge transfer from English and are also efficient in terms of sequence length.

**Knowledge and Reasoning tasks:** Table 6 shows the results on translated versions of various knowledge and reasoning benchmarks such as MMLU (Hendrycks et al., 2021b), BoolQ (Clark et al., 2019b), ARC (Clark et al., 2018b) and CommonsenseQA (Talmor et al., 2019). For most of the tasks, the results of the romanized script are competitive with the native script. At the same time, they are efficient to process, as discussed earlier. For IndicQA without context (open domain question answering task), we observe that the romanized model shows ~10 F1 point gain over the native script model. This demonstrates the overall effectiveness of utilizing romanization for non-Latin script languages.

**Discussion:** From the analysis and results presented earlier, we can summarize the following. Romanized representations are significantly more efficient compared to native script representations. They are better aligned to English representations compared to native script representations, hence they are better for cross-lingual transfer. Across a range of tasks and languages, romanized representations provide competitive or better performance over their native script counterparts (detailed results per language and task are presented in Appendix D. Specifically, for a generation task like machine translation, we see significant gains while translating out of English with romanized representations. In Appendix C, we show some examples of romanized and native script translations, illustrating the gains from romanization. These results indicate that romanization is a promising alternative to extending the capabilities of English LLMs to other languages using non-Roman scripts.

We primarily base our study on English-only models due to their state-of-the-art performance



across a diverse array of tasks and consistent improvements in capabilities. In contrast, multilingual models like BLOOM significantly lag behind English-only models, even on English tasks, and their progress has been slower. This disparity is due to the limited amount of English data and the lack of fine-tuning data for non-English languages used in training multilingual models. Therefore, improving English-only models is a practical way to enhance multilingual LLMs. We fine-tuned the BLOOM 7B model with the instruction dataset in both native and roman script. We observe that the roman script model outperforms the native script model, consistent with the findings on the Llama 2 base model. We provide the evaluation results for BLOOM 7B model across different tasks in Appendix E.

## 6 Conclusion

In this study, we proposed the use of romanization to extend the performance of LLMs primarily trained in English to other languages. Our approach successfully unlocks LLM capabilities for non-English languages by using romanization to bridge English and non-English language representations. We have empirically demonstrated the effectiveness of this strategy through experiments involving few-shot prompting, continual pretraining, and instruction fine-tuning on a variety of tasks. Additionally, using romanized data increases inference speed and maximum processable sequence length, while reducing memory requirements by 2x to 4x depending on the language.

## 7 Future Work

Looking forward, we also plan to explore reversible, deterministic transliteration so that transliteration errors in output post-processing are eliminated. In addition, it would be interesting to explore if romanized representation can act as a bridge to efficiently improve native script performance also when English LLMs are extended to incorporate non-Roman tokens.

## 8 Limitations

In this work, we chose to use natural romanizations since they tended to have lower fertility and might have been seen in the base LM pretraining. Transliterating the outputs back to the native script can be lossy. We use a high-quality transliteration

model to minimize the transliteration errors. However, such high-quality transliterations might not be available for all languages. Using a deterministic and reversible transliteration scheme can alleviate this problem.

The experiments in this work have been conducted on Indian languages. Indian languages are representative of the scenarios we study and we cover multiple language families and scripts with varying resource status. Previous studies have also used Indic languages as a case study for multilinguality and multi-script scenarios (Moosa et al., 2023; Khemchandani et al., 2021). Hence, we believe the findings should generalize to other language families - but further experimentation can confirm this hypothesis.

We have done only a limited amount of pretraining data in our experiments due to compute constraints. Resource constraints led us to limit our experiments to a 7B model, but experiments with larger models can yield more insights. Larger models and pre-training on larger datasets might benefit the multilingual models (Kaplan et al., 2020; Fernandes et al., 2023). We expect the broad conclusions to be the same - that romanized models are more efficient and better/comparable in task performance to native-script models.

Some of the evaluations have been done on translated benchmarks. While this is not ideal, most low-resource languages lack testsets across diverse tasks. Hence, multiple works have relied on translated benchmarks to measure multilingual performance - these can give indicative trends on performance across languages. More efforts are needed to create multilingual benchmarks for diverse tasks.

## 9 Ethics Statement

We use romanization as a way to align English with other languages using non-Latin scripts. Given the current state of LLMs, this seems like a practical direction to extend the capabilities of the best LLMs to other languages. The intention is not to supplant the use of a native script (which is widely adopted and has a rich literary tradition) with romanized script but to use romanization as a way to efficiently bring the benefits of LLM technology to low-resource languages written in non-Latin scripts. Further advancements are needed to extend this line of research to improve native script performance efficiently.

This work does not involve any new data collec-

tion and does not employ any annotators for data collection. We utilize publicly available datasets for the experiments reported in this work. Some of these datasets originate from web crawls, and we do not explicitly attempt to identify any biases within these datasets, using them in their original form.

### Acknowledgments

We would like to thank EkStep Foundation and Nilekani Philanthropies for their generous grant towards research and building datasets, models, tools and other resources for Indian languages.

### References

- Julien Abadji, Pedro Javier Ortiz Suárez, Laurent Romary, and Benoît Sagot. 2021. Ungoliant: An optimized pipeline for the generation of a very large-scale multilingual web corpus. In *CMLC 2021-9th Workshop on Challenges in the Management of Large Corpora*.
- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2023. [MEGA: Multilingual evaluation of generative AI](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4232–4267, Singapore. Association for Computational Linguistics.
- Chantal Amrhein and Rico Sennrich. 2020. [On Romanization for model transfer between scripts in neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2461–2469, Online. Association for Computational Linguistics.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Ethan C. Chau and Noah A. Smith. 2021. [Specializing multilingual language models: An empirical study](#). In *Proceedings of the 1st Workshop on Multilingual Representation Learning*, pages 51–61, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2023. [Palm: Scaling language modeling with pathways](#). *Journal of Machine Learning Research*, 24(240):1–113.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019a. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019b. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018a. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv: 1803.05457*.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018b. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#).
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht,

- Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, Jeff Wang, and NLLB Team. 2024. [Scaling neural machine translation to 200 languages](#). *Nature*.
- Yiming Cui, Ziqing Yang, and Xin Yao. 2023. Efficient and effective text encoding for chinese llama and alpaca. *arXiv preprint arXiv:2304.08177*.
- Raj Dabre, Himani Shrotriya, Anoop Kunchukuttan, Ratish Puduppully, Mitesh Khapra, and Pratyush Kumar. 2022. [IndicBART: A pre-trained model for indic natural language generation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1849–1863, Dublin, Ireland. Association for Computational Linguistics.
- Tejas Dhamecha, Rudra Murthy, Samarth Bharadwaj, Karthik Sankaranarayanan, and Pushpak Bhatlacharyya. 2021. [Role of Language Relatedness in Multilingual Fine-tuning of Language Models: A Case Study in Indo-Aryan Languages](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8584–8595, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sumanth Doddapaneni, Rahul Aralikkatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan, and Pratyush Kumar. 2023. [Towards leaving no Indic language behind: Building monolingual corpora, benchmark and models for Indic languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12402–12426, Toronto, Canada. Association for Computational Linguistics.
- Patrick Fernandes, Behrooz Ghorbani, Xavier Garcia, Markus Freitag, and Orhan Firat. 2023. [Scaling laws for multilingual neural machine translation](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10053–10071. PMLR.
- Jay Gala, Pranjal A Chitale, A K Raghavan, Varun Gumma, Sumanth Doddapaneni, Aswanth Kumar M, Janki Atul Nawale, Anupama Sujatha, Ratish Puduppully, Vivek Raghavan, Pratyush Kumar, Mitesh M Khapra, Raj Dabre, and Anoop Kunchukuttan. 2023. [Indictrans2: Towards high-quality and accessible machine translation models for all 22 scheduled indian languages](#). *Transactions on Machine Learning Research*.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’ Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022a. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022b. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Vikrant Goyal, Anoop Kunchukuttan, Rahul Kejriwal, Siddharth Jain, and Amit Bhagwat. 2020. [Contact relatedness can help improve multilingual NMT: Microsoft STCI-MT @ WMT20](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 202–206, Online. Association for Computational Linguistics.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XL-sum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring massive multitask language understanding](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- 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*.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#).
- Yash Khemchandani, Sarvesh Mehtani, Vaidehi Patil, Abhijeet Awasthi, Partha Talukdar, and Sunita Sarawagi. 2021. [Exploiting language relatedness for low web-resource language model adaptation: An Indic languages study](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1312–1323, Online. Association for Computational Linguistics.
- Christo Kirov, Cibu Johny, Anna Katanova, Alexander Gutkin, and Brian Roark. 2024. [Context-aware Transliteration of Romanized South Asian Languages](#). *Computational Linguistics*, pages 1–61.



- Aman Kumar, Himani Shrotriya, Prachi Sahu, Amogh Mishra, Raj Dabre, Ratish Puduppully, Anoop Kunchukuttan, Mitesh M. Khapra, and Pratyush Kumar. 2022. **IndicNLG benchmark: Multilingual datasets for diverse NLG tasks in Indic languages**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5363–5394, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Anoop Kunchukuttan. 2020. The IndicNLP Library. [https://github.com/anoopkunchukuttan/indic\\_nlp\\_library/blob/master/docs/indicnlp.pdf](https://github.com/anoopkunchukuttan/indic_nlp_library/blob/master/docs/indicnlp.pdf).
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. **Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327, Singapore. Association for Computational Linguistics.
- Haonan Li, Fajri Koto, Minghao Wu, Alham Fikri Aji, and Timothy Baldwin. 2023. **Bactrian-x: Multilingual replicable instruction-following models with low-rank adaptation**.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. **Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing**.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. **Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing**. *ACM Computing Surveys*, 55(9):1–35.
- S. Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. **The flan collection: Designing data and methods for effective instruction tuning**. *International Conference on Machine Learning*.
- Yash Madhani, Sushane Parthan, Priyanka Bedekar, Gokul Nc, Ruchi Khapra, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Khapra. 2023. **Aksharant: Open Indic-language transliteration datasets and models for the next billion users**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 40–57, Singapore. Association for Computational Linguistics.
- Benjamin Minixhofer, Fabian Paischer, and Navid Reksabsaz. 2022. **WECHSEL: Effective initialization of subword embeddings for cross-lingual transfer of monolingual language models**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3992–4006, Seattle, United States. Association for Computational Linguistics.
- Ibraheem Muhammad Moosa, Mahmud Elahi Akhter, and Ashfia Binte Habib. 2023. **Does transliteration help multilingual language modeling?** In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 670–685, Dubrovnik, Croatia. Association for Computational Linguistics.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. **When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462, Online. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Wenxuan Zhang, Xin Li, Mahani Aljunied, Qingyu Tan, Liying Cheng, Guanzheng Chen, Yue Deng, Sen Yang, Chaoqun Liu, et al. 2023. **Seallms—large language models for southeast asia**. *arXiv preprint arXiv:2312.00738*.
- OpenAI, :, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, et al. 2023. **Gpt-4 technical report**.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Matt Post. 2018. **A call for clarity in reporting BLEU scores**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Sukannya Purkayastha, Sebastian Ruder, Jonas Pfeiffer, Iryna Gurevych, and Ivan Vulić. 2023. **Romanization-based large-scale adaptation of multilingual language models**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7996–8005, Singapore. Association for Computational Linguistics.
- Haiyue Song, Raj Dabre, Zhuoyuan Mao, Fei Cheng, Sadao Kurohashi, and Eiichiro Sumita. 2020. **Pre-training via leveraging assisting languages for neural machine translation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 279–285, Online. Association for Computational Linguistics.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. **Asynchronous pipeline for processing huge corpora on medium to low resource infrastructures**. In *7th Workshop on the Challenges in the*



- Management of Large Corpora (CMLC-7)*. Leibniz-Institut für Deutsche Sprache.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [How far can camels go? exploring the state of instruction tuning on open resources](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2024. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Xiangpeng Wei, Haoran Wei, Huan Lin, Tianhao Li, Pei Zhang, Xingzhang Ren, Mei Li, Yu Wan, Zhiwei Cao, Binbin Xie, Tianxiang Hu, Shangjie Li, Binyuan Hui, Bowen Yu, Dayiheng Liu, Baosong Yang, Fei Huang, and Jun Xie. 2023. PolyLM: An open source polyglot large language model. *arXiv preprint arXiv:2307.06018*.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2020. [CCNet: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4003–4012, Marseille, France. European Language Resources Association.
- BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luciani, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Judit Ács. 2020. Exploring BERT’s vocabulary. <http://juditacs.github.io/2019/02/19/bert-tokenization-stats.html>.
- Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction finetuned open-access multilingual language model](#).

## A Model Training Details

Tables 7 and 8 show the hyperparameters used for training the CPT and IFT models.

Hyperparameter	Value
Batch Size (tokens)	1M
Learning Rate	5e-5
Number of Epochs	1
Maximum Sequence Length	2,048

Table 7: The range of hyperparameters used for continual pretraining.

Hyperparameter	Value
Batch Size (examples)	128
Learning Rate	5e-5
Number of Epochs	1
Maximum Sequence Length	2,048 <sup>4</sup>

Table 8: The range of hyperparameters used for supervised fine-tuning.

## B Prompt Templates

We list the various prompt templates for evaluations across different test sets (see Figures 3 to 12).

## C Translation Examples

This section shows some examples of translations using the romanized and native-script IFT models (See Figures 13 to 16). We can see in these selected examples that romanized representation performs better.

## D Language Specific Results

We report the language-wise evaluation results for 5 Indic languages considered for this study across a diverse array of tasks ranging from text generation, text understanding, and text reasoning in Tables 9 to 22.

## E Results on BLOOM 7B

We report evaluation results for 5 Indic languages for the tasks such as machine translation, headline generation and summarization on instruction fine-tuned BLOOM 7B in Tables 23 to 25.

---

<sup>4</sup>Although this is the maximum permissible length, most examples fall far below this length during fine-tuning.

```
Predict the sentiment of the review. The possible choices for the sentiment are: 'positive' and 'negative'.

Review: {{ text }}
Sentiment: {{ label }}

...

Review: {{ text }}
Sentiment:
```

Figure 3: Prompt template for IndicSentiment

```
I am hesitating between two options. Help me choose the more likely cause or effect.

{{ premise }} {{ connector }}
A. {{ choice1 }}
B. {{ choice2 }}
Answer: {{ label }}

...

{{ premise }} {{ connector }}
A. {{ choice1 }}
B. {{ choice2 }}
Answer:
```

Figure 4: Prompt template for IndicCOPA

```
Answer whether the hypothesis is more likely to be true (entailment), false (contradiction), or unknown (neutral) based on the given premise.

Premise: {{ premise }}
Hypothesis: {{ hypothesis }}
Answer: {{ label }}

...

Premise: {{ premise }}
Hypothesis: {{ hypothesis }}
Answer:
```

Figure 5: Prompt template for IndicXNLI

Answer the following question based on the information in the given passage.

Passage: {{ passage }}  
Question: {{ question }}  
Answer: {{ answer }}

...  
...

Passage: {{ passage }}  
Question: {{ question }}  
Answer:

Figure 6: Prompt template for IndicQA.

Summarize the following {{ language }} article(s) as accurately as possible in few sentences.

{{ language }} article: {{ article }}  
{{ language }} summary: {{ summary }}

...  
...

{{ language }} article: {{ article }}  
{{ language }} summary:

Figure 7: Prompt template for XLSum.

Generate a headline for the following article(s) as accurately as possible.

{{ language }} article: {{ article }}  
{{ language }} headline: {{ headline }}

...  
...

{{ language }} article: {{ article }}  
{{ language }} headline:

Figure 8: Prompt template for IndicHeadline.



The following are multiple choice questions (with answers) about {{ subject }}.

```
{{ question }}  
A. {{ choice1 }}  
B. {{ choice2 }}  
C. {{ choice3 }}  
D. {{ choice4 }}  
Answer: {{ label }}
```

...  
...

{{ subject }}

```
{{ question }}  
A. {{ choice1 }}  
B. {{ choice2 }}  
C. {{ choice3 }}  
D. {{ choice4 }}  
Answer:
```

Figure 9: Prompt template for MMLU.

The following are binary yes/no choice questions (with answers).

```
Passage: {{ passage }}  
Question: {{ question }}  
Answer: {{ label }}
```

...  
...

```
Passage: {{ passage }}  
Question: {{ question }}  
Answer:
```

Figure 10: Prompt template for BoolQ.

The following are multiple choice questions (with answers) about science.

```
{{ question }}  
A. {{ choice1 }}  
B. {{ choice2 }}  
C. {{ choice3 }}  
D. {{ choice4 }}  
E. {{ choice5 }}  
Answer: {{ label }}
```

...  
...

```
{{ question }}  
A. {{ choice1 }}  
B. {{ choice2 }}  
C. {{ choice3 }}  
D. {{ choice4 }}  
E. {{ choice5 }}  
Answer: {{ label }}
```

Figure 11: Prompt template for ARC.

The following are multiple choice questions (with answers) requiring common sense.

```
{{ question }}
A. {{ choice1 }}
B. {{ choice2 }}
C. {{ choice3 }}
D. {{ choice4 }}
E. {{ choice5 }}
Answer: {{ label }}

...

{{ question }}
A. {{ choice1 }}
B. {{ choice2 }}
C. {{ choice3 }}
D. {{ choice4 }}
E. {{ choice5 }}
Answer: {{ label }}
```

Figure 12: Prompt template for CommonsenseQA.

**English:** The movie, featuring Ryan Gosling and Emma Stone, received nominations in all major categories

**Generated Native Output:** फिल्म, जिसमें रायन गोजलिंग और एमा स्टोन के साथ है, सभी महत्वपूर्ण प्रकार के प्रति प्रतिनिधित्व प्राप्त हुए हैं।

Figure 13: Example of *zero-shot* translation for English to Hindi machine translation task to generate text in native script using a model fine-tuned on native script data.

**English:** The movie, featuring Ryan Gosling and Emma Stone, received nominations in all major categories

**Generated Romanized Output:** ryan gosling our emma stone abhineet film quo sabhi pramukh shreniyon main namankan praapt hue.

**Transliterated Native Output:** रयान गोजलिंग और एम्मा स्टोन अभिनीत फिल्म को सभी प्रमुख श्रेणियों में नामांकन प्राप्त हुए।

Figure 14: Example of *zero-shot* translation for English to Hindi machine translation task to generate text in romanized script using a model fine-tuned on romanized data. The romanized output is then transliterated to native script.

**Hindi:** जिस फिल्म में रायन गोजलिंग और एम्मा स्टोन ने मुख्य भूमिकाएँ निभाई हैं, उसे सभी प्रमुख श्रेणियों में नॉमिनेशन किया गया है।

**Generated Output:** The film was nominated for the Academy Award for Best Picture, Best Director, Best Original Screenplay, Best Original Score, Best Sound, Best Film Editing, Best Cinematography, Best Art Direction, Best Costume Design, Best Makeup, Best Visual Effects, Best Sound Mixing, Best Sound Editing, Best Original Song, Best Original Song Score and Best Original Song.

Figure 15: Example of *zero-shot* translation for the Hindi to English machine translation task using native script input for a model fine-tuned on native script data.

**Hindi:** jis film main rayan gozlling our amma stone nay mukhya bhoomikaaen nibhai hain, use sabhi pramukh shreniyon main nomination kiya gaya hai.

**Generated Output:** The film stars Rayan Gosling and Amma Stone in the lead roles, with nominations in all major categories.

Figure 16: Example of *zero-shot* translation for the Hindi to English machine translation task using romanized script input for a model fine-tuned on romanized data.

Model	Shots	Input text	Hindi		Gujarati		Tamil		Marathi		Malayalam	
			N	R	N	R	N	R	N	R	N	R
Llama 2	1	Native	22.95	29.21	8.15	11.9	12.12	14.62	15.32	20.22	11.2	14.98
Llama 2	3	Native	23.73	29.91	7.38	11.48	12.4	15.49	16.19	21.4	11.55	15.31
Llama 2	1	Romanized	17.25	18.86	14.63	16.16	12.8	14.87	13.42	15.85	12.88	15.48
Llama 2	3	Romanized	15.01	18.25	10.8	14.79	9.3	11.39	12.09	16.61	12.88	15.53
CPT - N	1	Native	32.07	37.89	15.94	18.79	25.81	28.87	29.47	34.24	24.48	28.36
CPT - N	3	Native	32.9	38.51	16.3	19.29	26.2	29.12	29.9	34.75	25.64	29.45
CPT - R	1	Romanized	30.93	33.66	29.73	35.35	30.11	32.71	28.27	33.71	28.72	32.47
CPT - R	3	Romanized	31.51	34.99	30.55	36.8	31.5	35.04	29.42	35.08	30.85	34.86
IFT (Llama 2) - N	0	Native	43.28	48.43	23.54	25.82	38.99	41.13	40.13	44.72	38.85	42.12
IFT (CPT - N) - N	0	Native	44.08	49.12	23.86	25.97	39.3	41.51	40.31	44.8	39.44	42.78
IFT (Llama 2) - R	0	Romanized	45.15	55.23	44.8	52.98	49.13	52.79	42.78	49.6	47.71	52.63
IFT (CPT - R) - R	0	Romanized	<b>45.94</b>	<b>56.19</b>	<b>45.37</b>	<b>53.5</b>	<b>50.06</b>	<b>53.71</b>	<b>44</b>	<b>50.73</b>	<b>48.96</b>	<b>53.76</b>

Table 9: chrF scores for En-XX translation task on FLORES-200 devtest (Goyal et al., 2022a) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	Shots	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	40.49	13.63	16.06	24.18	17.93
Llama 2	3	Native	40.38	17.58	17.22	24.15	17.78
Llama 2	1	Romanized	29.24	22.62	17.06	22.56	21.1
Llama 2	3	Romanized	28.91	23.28	19.86	23.6	21.96
CPT - N	1	Native	51.37	24.77	33.2	47.04	36.34
CPT - N	3	Native	50.31	25.73	28.56	43.37	37.4
CPT - R	1	Romanized	45.43	45.08	38.34	44.28	39.52
CPT - R	3	Romanized	45.4	44.82	33.72	43.56	40.71
IFT (Llama 2) - N	0	Native	51.7	47.26	42.27	48.8	47.83
IFT (CPT - N) - N	0	Native	53.67	49.51	44.44	<b>51.29</b>	50
IFT (Llama 2) - R	0	Romanized	51.84	50.17	41.67	48.51	47.14
IFT (CPT - R) - R	0	Romanized	<b>53.89</b>	<b>53.3</b>	<b>45.09</b>	51.06	<b>50.43</b>

Table 10: chrF scores for XX-En translation task on FLORES-200 devtest (Goyal et al., 2022a) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	Shots	Input text	Hindi		Gujarati		Tamil		Marathi	
			N	R	N	R	N	R	N	R
Llama 2	1	Native	13.32	9.59	3.2	2.79	4.13	3.47	6.85	6.11
Llama 2	1	Romanized	13.43	15.96	8.95	7.65	9.05	6.69	9.22	7.45
CPT - N	1	Native	<b>16.32</b>	12.18	5.18	4.58	4.3	3.91	4.55	4.61
CPT - R	1	Romanized	13.05	14.77	<b>13.8</b>	<b>12.98</b>	10.86	<b>9.68</b>	<b>12.06</b>	<b>11.34</b>
IFT (Llama 2) - N	0	Native	11.71	8.66	4.38	3.18	6.62	5.62	7.41	6.61
IFT (CPT - N) - N	0	Native	12.56	9.21	4.3	3.52	6.51	5.66	7.71	6.64
IFT (Llama 2) - R	0	Romanized	14.73	16.87	11.84	10.56	10.83	8.59	11	9.53
IFT (CPT - R) - R	0	Romanized	15.05	<b>17.92</b>	12.36	11.16	<b>11.28</b>	9.32	11.56	9.87

Table 11: RougeL scores on XLSum (Hasan et al., 2021) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi		Gujarati		Tamil		Marathi		Malayalam	
			N	R	N	R	N	R	N	R	N	R
Llama 2	1	Native	20.91	15.01	6.17	5.55	19.14	17.8	9.26	8.45	12.84	11.45
Llama 2	1	Romanized	17.58	20.55	12.18	13.72	23.5	26.16	10.55	11.32	13.97	18.83
CPT - N	1	Native	<b>22.74</b>	16.52	9.07	7.97	23.59	22.79	<b>19.21</b>	<b>18.72</b>	15.57	14.74
CPT - R	1	Romanized	20.19	<b>24.44</b>	12.83	14.34	<b>29.13</b>	<b>32.03</b>	15.14	17.39	<b>17.29</b>	<b>22.38</b>
IFT (Llama 2) - N	0	Native	18.38	13.43	6.36	5.64	15.88	15.35	9.96	9.84	10.51	9.85
IFT (CPT - N) - N	0	Native	19.47	14.24	6.65	5.98	15.14	14.75	10.77	10.7	11.04	10.68
IFT (Llama 2) - R	0	Romanized	17.76	21.46	12.88	15.43	20.89	22.27	11.51	12.62	13.45	17.45
IFT (CPT - R) - R	0	Romanized	18.74	21.91	<b>13.04</b>	<b>15.64</b>	21.61	23.82	12.2	13.59	14.56	19.27

Table 12: RougeL scores on IndicHeadline (Kumar et al., 2022) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	28.67	25.59	26.06	27.58	26.44
Llama 2	3	Native	29.11	26.57	26.75	27.85	26.86
Llama 2	1	Romanized	27.52	26.57	26.34	27.14	27.11
Llama 2	3	Romanized	28.38	27.62	26.65	27.24	27.52
CPT - N	1	Native	29.61	27.34	28.63	29.58	28.35
CPT - N	3	Native	30.08	27.62	28.66	29.91	28.61
CPT - R	1	Romanized	30.8	30.72	<b>30.38</b>	31.33	29.7
CPT - R	3	Romanized	<b>31.55</b>	<b>31.05</b>	30.18	<b>32.15</b>	<b>30.76</b>
IFT (Llama 2) - N	0	Native	24.47	24.08	24.09	23.3	24.36
IFT (CPT - N) - N	0	Native	24.72	24.05	24.78	24.62	24.81
IFT (Llama 2) - R	0	Romanized	24.75	25.96	26.03	25.54	26.48
IFT (CPT - R) - R	0	Romanized	25.58	26.4	26.73	26.1	26.57

Table 13: Accuracy on translated MMLU (Hendrycks et al., 2021a) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	57.58	58.78	56.88	61.07	55.2
Llama 2	3	Native	45.08	57.58	39.3	40.92	39.39
Llama 2	1	Romanized	62.23	60.03	58.72	60.61	62.14
Llama 2	3	Romanized	61.93	<b>61.77</b>	60.09	58.9	<b>62.48</b>
CPT - N	1	Native	60.7	56.21	57.71	61.41	59.45
CPT - N	3	Native	50.67	60.64	45.78	39.36	40.95
CPT - R	1	Romanized	62.17	61.5	<b>61.44</b>	<b>61.93</b>	62.23
CPT - R	3	Romanized	<b>64.77</b>	60.92	55.41	47.86	56.54
IFT (Llama 2) - N	0	Native	60.12	61.65	60.06	61.01	60.18
IFT (CPT - N) - N	0	Native	60.34	61.5	58.32	60.43	59.6
IFT (Llama 2) - R	0	Romanized	61.62	56.91	58.01	58.69	47.68
IFT (CPT - R) - R	0	Romanized	60.7	45.02	38.01	53.24	37.77

Table 14: Accuracy on translated BoolQ (Clark et al., 2019a) across 5 Indic languages in both native (N) and romanized (R) script with different shots.



Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	27.27	25.67	25.63	25.13	25.42
Llama 2	3	Native	27.36	26.09	26.18	26.6	25.51
Llama 2	1	Romanized	27.23	25.29	26.05	26.98	26.35
Llama 2	3	Romanized	27.27	<b>26.64</b>	25.29	26.85	26.14
CPT - N	1	Native	<b>30.98</b>	26.22	<b>28.87</b>	<b>30.56</b>	27.23
CPT - N	3	Native	30.68	26.43	27.4	29.12	<b>27.61</b>
CPT - R	1	Romanized	27.23	25.88	26.26	26.89	27.19
CPT - R	3	Romanized	27.61	26.56	26.39	27.74	27.23
IFT (Llama 2) - N	0	Native	25.42	25.88	25.25	25.42	26.09
IFT (CPT - N) - N	0	Native	26.35	25.84	25.67	25.93	26.26
IFT (Llama 2) - R	0	Romanized	25.08	25.72	25.04	25.34	25.76
IFT (CPT - R) - R	0	Romanized	25.97	25.34	24.24	25.42	25.72

Table 15: Accuracy on translated ARC-Easy (Clark et al., 2018b) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	26.45	27.05	26.71	26.71	25.77
Llama 2	3	Native	27.22	26.19	25.17	25.6	24.32
Llama 2	1	Romanized	24.4	25.68	24.32	24.32	24.32
Llama 2	3	Romanized	24.83	27.47	25.09	25.09	25.34
CPT - N	1	Native	27.65	26.62	<b>27.3</b>	27.3	<b>28.92</b>
CPT - N	3	Native	<b>29.69</b>	25.85	26.96	26.96	26.54
CPT - R	1	Romanized	25.34	26.79	24.57	<b>28.07</b>	27.22
CPT - R	3	Romanized	24.49	<b>27.56</b>	25.94	27.99	27.3
IFT (Llama 2) - N	0	Native	22.27	22.78	23.46	21.93	23.29
IFT (CPT - N) - N	0	Native	25.85	23.21	24.57	24.06	23.29
IFT (Llama 2) - R	0	Romanized	21.67	22.87	23.04	22.87	21.67
IFT (CPT - R) - R	0	Romanized	22.78	22.95	22.7	22.35	23.21

Table 16: Accuracy on translated ARC-Challenge (Clark et al., 2018b) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	23.99	19	20.55	21.45	20.06
Llama 2	3	Native	23.75	18.75	20.22	20.22	19.32
Llama 2	1	Romanized	21.54	20.48	19.49	19.74	19.08
Llama 2	3	Romanized	21.54	19.57	19.98	19.82	19.98
CPT - N	1	Native	28.09	23.75	28.5	26.94	25.47
CPT - N	3	Native	<b>31.12</b>	22.44	<b>28.99</b>	<b>28.09</b>	26.45
CPT - R	1	Romanized	23.75	27.19	25.3	25.22	26.28
CPT - R	3	Romanized	25.71	<b>30.22</b>	28	27.1	<b>29.89</b>
IFT (Llama 2) - N	0	Native	20.64	19.25	20.07	20.97	20.31
IFT (CPT - N) - N	0	Native	23.34	23.18	23.59	25.14	25.55
IFT (Llama 2) - R	0	Romanized	20.23	21.54	21.29	22.93	22.44
IFT (CPT - R) - R	0	Romanized	23.91	26.7	25.55	26.04	27.52

Table 17: Accuracy on translated CommonsenseQA (Talmor et al., 2019) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	53.2	35.1	46	50.3	40.4
Llama 2	3	Native	77.8	28.4	34.5	49.2	31.7
Llama 2	1	Romanized	59.1	44.9	42.4	53.5	40.9
Llama 2	3	Romanized	67.5	38.9	33.8	44.3	34
CPT - N	1	Native	64.9	68	77.8	75.1	67.1
CPT - N	3	Native	89.7	86.6	87.8	89.6	82.9
CPT - R	1	Romanized	91.2	86.6	88.4	90	86.1
CPT - R	3	Romanized	<b>91.9</b>	<b>93.3</b>	<b>93.3</b>	<b>93.1</b>	<b>92.5</b>
IFT (Llama 2) - N	0	Native	90.4	88.1	87.6	89.5	89
IFT (CPT - N) - N	0	Native	87.2	88.7	89.6	90	88.9
IFT (Llama 2) - R	0	Romanized	75.3	68.8	71.2	77.6	77.3
IFT (CPT - R) - R	0	Romanized	83.7	73.5	87.7	81.9	85.4

Table 18: F1 score on IndicSentiment (Doddapaneni et al., 2023) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	41.1	41.1	1.6	41.1	1.6
Llama 2	3	Native	52.1	56.7	7.5	34.7	1.6
Llama 2	1	Romanized	39.2	36.2	1.6	26.6	1.6
Llama 2	3	Romanized	23.6	16.1	3.1	5.9	0
CPT - N	1	Native	40.5	40.7	1.6	38.4	1.6
CPT - N	3	Native	<b>61</b>	<b>65.7</b>	11.6	<b>60</b>	11.7
CPT - R	1	Romanized	38.6	37.4	0	31.8	1.6
CPT - R	3	Romanized	30.3	42.7	6.6	32.3	0.8
IFT (Llama 2) - N	0	Native	5.2	1.7	44.1	1.8	7.4
IFT (CPT - N) - N	0	Native	51.3	33.2	47.6	35.4	56.3
IFT (Llama 2) - R	0	Romanized	18.9	58.4	18.8	35.9	63.6
IFT (CPT - R) - R	0	Romanized	33.2	62.3	<b>63.4</b>	47.1	<b>66.7</b>

Table 19: F1 score on IndicCOPA (Doddapaneni et al., 2023) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
Llama 2	1	Native	17.6	16.7	16.8	17.1	16.7
Llama 2	3	Native	27.4	25.1	17.4	18.7	16.9
Llama 2	1	Romanized	27.5	28.5	26.5	26.4	27.3
Llama 2	3	Romanized	31.4	31.6	30.5	31.2	31.1
CPT - N	1	Native	19.9	19.3	19.4	17.7	17.1
CPT - N	3	Native	<b>44.6</b>	30.1	32.8	<b>45.2</b>	26.1
CPT - R	1	Romanized	34.1	37	36.6	37.3	39
CPT - R	3	Romanized	38.6	37.3	37.5	37.3	35.7
IFT (Llama 2) - N	0	Native	36.3	37.9	<b>43.1</b>	42.9	38.7
IFT (CPT - N) - N	0	Native	42.2	41.2	<b>43.1</b>	42.5	<b>42.5</b>
IFT (Llama 2) - R	0	Romanized	24	22.5	23	22.1	28.6
IFT (CPT - R) - R	0	Romanized	38.5	<b>38.9</b>	37.2	38.1	39.2

Table 20: F1 score on IndicXNLI (Doddapaneni et al., 2023) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi		Gujarati		Tamil		Marathi		Malayalam	
			N	R	N	R	N	R	N	R	N	R
Llama 2	1	Native	21.02	24.63	3.26	3.39	8.75	8.91	10.08	12.22	5.73	5.73
Llama 2	1	Romanized	29.11	32.8	19.95	21.97	11.34	13.32	21.11	21.6	17.18	20.42
CPT - N	1	Native	28.98	30.36	17.01	17.76	21.01	21.04	28.76	28.96	23.37	23.43
CPT - R	1	Romanized	<b>29.58</b>	32.8	<b>27.01</b>	<b>28.85</b>	<b>24.69</b>	<b>27.11</b>	<b>30.01</b>	<b>30.53</b>	<b>24.95</b>	<b>27.51</b>
IFT (Llama 2) - N	0	Native	19.75	22.79	11.66	14.3	17.56	17.81	18.39	21.96	11.15	11.4
IFT (CPT - N) - N	0	Native	21.85	25.68	13.54	16.21	18.49	18.68	18.02	21.43	12.07	12.49
IFT (Llama 2) - R	0	Romanized	23.54	33.79	19.15	24.3	16.44	25.54	21.26	25.79	13.2	20.31
IFT (CPT - R) - R	0	Romanized	24.44	<b>33.81</b>	21.41	26.9	17.65	25.27	23.64	28.41	14.53	22.01

Table 21: F1 score on IndicQA (with context) (Doddapaneni et al., 2023) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	n-shot	Input text	Hindi		Gujarati		Tamil		Marathi		Malayalam	
			N	R	N	R	N	R	N	R	N	R
Llama 2	1	Native	4.97	6.69	0.2	0.22	1.02	1.09	0.9	1.16	0.31	0.36
Llama 2	1	Romanized	6.96	7.33	<b>16.77</b>	16.82	12.26	12.26	16.45	16.38	12.34	12.77
CPT - N	1	Native	7.44	8.67	2.73	3.15	4.66	4.64	7.41	7.82	5.55	5.67
CPT - R	1	Romanized	<b>13.85</b>	<b>15.19</b>	13.46	<b>13.41</b>	<b>16.66</b>	<b>17.13</b>	<b>26.75</b>	<b>26.82</b>	<b>17.38</b>	<b>17.93</b>
IFT (Llama 2) - N	0	Native	2.36	2.64	0.66	0.72	1.48	1.53	1.18	1.28	0.89	0.9
IFT (CPT - N) - N	0	Native	2.73	3.03	0.85	0.95	1.15	1.16	1.27	1.47	0.82	0.9
IFT (Llama 2) - R	0	Romanized	2.21	2.93	0.65	0.76	1.37	1.42	1.75	1.75	1.01	1.33
IFT (CPT - R) - R	0	Romanized	3.32	3.69	1.03	0.96	1.4	1.67	1.81	1.9	1.05	1.21

Table 22: F1 score on IndicQA (without context) (Doddapaneni et al., 2023) across 5 Indic languages in both native (N) and romanized (R) script with different shots.

Model	Shots	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
IFT (BLOOM) - N	0	Native	51.03	48.39	<b>52.1</b>	45.63	<b>49.94</b>
IFT (BLOOM) - R	0	Romanized	<b>53.74</b>	<b>50.95</b>	50.09	<b>48.06</b>	48.93

Table 23: chrF scores for En-XX translation task on FLORES-200 devtest (Goyal et al., 2022a) across 5 Indic languages in both native and romanized script.

Model	Shots	Input text	Hindi	Gujarati	Tamil	Marathi	Malayalam
IFT (BLOOM) - N	0-shot	Native	7.44	4.36	7.66	0.8	3.7
IFT (BLOOM) - R	0-shot	Romanized	<b>20.96</b>	<b>15.87</b>	<b>25.4</b>	<b>13.46</b>	<b>19.42</b>

Table 24: RougeL scores on IndicHeadline (Kumar et al., 2022) across 5 Indic languages in both native and romanized script.

Model	Shots	Input text	Hindi	Gujarati	Tamil	Marathi
IFT (BLOOM) - N	0-shot	Native	1.73	2.84	2.56	2.73
IFT (BLOOM) - R	0-shot	Romanized	<b>17.16</b>	<b>10.74</b>	<b>8.63</b>	<b>9.36</b>

Table 25: RougeL scores on XLSum (Hasan et al., 2021) across 5 Indic languages in both native and romanized script.