

Happiness is Sharing a Vocabulary: A Study of Transliteration Methods

Haeji Jung¹, Jinju Kim^{2,5}, Kyungjin Kim³, Youjeong Roh⁴, David R. Mortensen⁵

¹University of British Columbia, Canada

²Sungkyunkwan University, Republic of Korea

³Seoul National University, Republic of Korea

⁴Electronics and Telecommunications Research Institute, Republic of Korea

⁵Carnegie Mellon University, USA

Correspondence: haeji.j@cs.ubc.ca, dmortens@cs.cmu.edu

Abstract

Transliteration has emerged as a promising means to bridge the gap between various languages in multilingual NLP, showing promising results especially for languages using non-Latin scripts. We investigate the degree to which shared script, overlapping token vocabularies, and shared phonology contribute to performance of multilingual models. To this end, we conduct controlled experiments using three kinds of transliteration (romanization, phonemic transcription, and substitution ciphers) as well as orthography. We evaluate each model on three downstream tasks—named entity recognition (NER), part-of-speech tagging (POS) and natural language inference (NLI)—and find that romanization significantly outperforms other input types in 7 out of 8 evaluation settings, largely consistent with our hypothesis that it is the most effective approach. We further analyze how each factor contributed to the success, and suggest that having longer (sub-word) tokens shared with pre-trained languages leads to better utilization of the model.

1 Introduction

Multilingual language modeling has drawn significant attention from researchers seeking to support diverse languages and promote fairness in AI. One crucial problem in improving multilingual language modeling is the *script barrier*, a phenomenon in which models struggle to share knowledge between languages written in different scripts, due to mismatched input representations.

Transliteration, which maps a given string to another sequence of characters (e.g., Cyrillic script → Latin script) has been explored as a potential solution to the problem of script barriers. In multilingual NLP, transliteration typically maps text to the Latin script or the International Phonetic Alphabet (IPA), giving various languages a shared input representation. Both representations encode linguistic information—specifically phonetic and

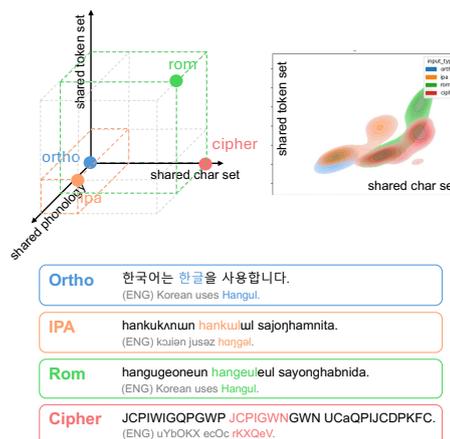


Figure 1: *Top left*: Conceptual visualization of the transliteration analysis schema, positioning input types (Ortho, IPA, Rom, Cipher) based on shared character set, token set, and phonology. *Top right*: KDE plot showing empirical distribution of overlap ratios for each quantifiable component. *Bottom*: Transliteration examples generated with each method.

phonological information—across languages. Here, we pose a question: *Is it the shared script itself or the linguistic information encoded in the scripts that helps the models adapt to other languages?*

To investigate this question, we define three key factors in transliteration—(i) shared character set, (ii) shared token set, and (iii) shared phonology of synonymous lexical items—that influence how a model processes and generalizes across languages. We then designate four input types varying in the degree to which these factors are present: Orthography (Ortho), IPA, Romanized (Rom), and Substitution Ciphersed text (Cipher). In the context of a script barrier—where two languages are written in different scripts—each input type can be characterized in terms of the three factors as illustrated in Figure 1 and detailed in Section 3.

We then conduct controlled experiments to understand how each factor of transliteration contributes to success. First, we pre-train Transformer-

based multilingual language models from scratch with each of these four input types, and then fine-tune them on target languages—including both seen and unseen languages—to examine how well each input type adapts to new languages. The languages used for pre-training are selected based on our defined language similarity scores as detailed in Section 4.1. We constructed four language sets to account for various scenarios in which language similarity or script familiarity may vary.

We hypothesize that Rom yields the best performance in handling diverse languages with diverse scripts, as it improves input representations across all three dimensions. Based on this assumption, IPA is expected to follow, as it enhances two out of three dimensions (shared phonology and tokens) while ciphered text only shares the character set and lacks other shared representations. Throughout the paper, we evaluate our hypothesis by comparing downstream task performance specifically on unseen languages and analyzing each method in terms of the defined factors, with an emphasis on shared tokens. We summarize our contribution as follows:

- We define three factors—shared character set, shared token set, and shared phonology—to explain the effect of transliteration on multilingual inputs.
- We conduct controlled pre-training across four language sets and input types, evaluating two downstream tasks to assess how each factor contributes to the success of transliteration.
- We analyze vocabulary overlap by token length, revealing how different transliteration methods yield distinct overlap patterns that contribute to cross-lingual adaptation.

2 Related Work

2.1 Script Barrier and Transliteration

Script barrier refers to the phenomenon in which multilingual language models struggle to share knowledge between languages written in different scripts (standard orthographies), due to mismatched input representations. Even massively multilingual models face persistent challenges in handling languages whose scripts were unseen or severely underrepresented during pre-training (de Vries et al., 2022; Fujinuma et al., 2022; Muller et al., 2021; Pfeiffer et al., 2021).

To overcome this limitation, *transliteration* has been explored as an efficient means to improve

cross-lingual transfer. Among transliteration methods, **romanization**—the conversion of non-Latin scripts into the Latin alphabet—has gained popularity due to the dominance of Latin scripts in pre-trained language models (Muller et al., 2021; Purkayastha et al., 2023; Husain et al., 2024). Another popular method is **grapheme-to-phoneme** (G2P) conversion, which converts original text (i.e., graphemes) into phonemes represented in International Phonetic Alphabet (IPA). G2P has been applied either by replacing input representations (Bharadwaj et al., 2016; Sohn et al., 2024) or by integrating phonemic information as an auxiliary signal (Nguyen et al., 2023, 2025).

While these approaches demonstrate transliteration’s practical benefits, relatively little is known about *why* it works. Moosa et al. (2023) examined how transliteration benefits multilingual models by analyzing representational similarity and tokenizer fertility. However, their study mainly focuses on closely related Indic languages, leaving open questions about how transliteration contributes to cross-lingual transfer in a broader, typologically diverse setting. Similarly, Jung et al. (2024) investigated phonemic representations in terms of representational similarity across languages, but without clarifying how transliterated inputs lead to improved cross-lingual performance.

Our work advances these studies by systematically isolating and analyzing the underlying factors—shared character set, shared token set, and shared phonology—that explain when and why transliteration is effective.

2.2 Vocabulary Overlap in Cross-lingual Transfer

Lexical overlap has been explored in various works in multilingual NLP. Prior studies have shown that shared lexical or subword units across languages allow models to reuse learned representations, thereby improving transferability (Pires et al., 2019; Chang et al., 2024; Limisiewicz et al., 2023). In contrast, a high proportion of unknown tokens—tokens that the model fails to segment due to their absence from the pre-training corpus—has been shown to hinder transfer and degrade downstream performance (Pfeiffer et al., 2021).

However, Philippy et al. (2023) put former literature together, acknowledging conflicting findings regarding how lexical overlap benefits cross-lingual transfer. They emphasized the need for task-specific analyses and a more comprehensive under-

| | Shared Char. Set | Shared Token Set | Shared Phonology |
|--------|------------------|------------------|------------------|
| Ortho | - | - | - |
| IPA | ± | ± | + |
| Rom | + | + | ± |
| Cipher | + | - | - |

Table 1: Input types characterized by the three factors for transliteration.

standing of this factor.

In this work, we examine vocabulary overlap to understand its contribution to the success of transliteration. By analyzing overlap by length, we provide deeper insights into how vocabulary overlap shapes multilingual adaptation.

3 Input Types

Input types are selected to vary according to three different dimensions:

- **Shared Character Set.** Transliteration usually enforces a shared character set across languages. For example, applying romanization to English (e.g., “hello” → “hello”) and Korean (e.g., “안녕하세요” → “annyeong-haseyo”) produces only Latin characters, which significantly reduces the number of unique characters and patterns that a tokenizer must capture.
- **Shared Token Set.** Transliteration also yield shared subword tokens across languages. Here, we specifically distinguish *tokens* from *characters*. By tokens we refer to subword tokens longer than one (more than one character). With English and Korean example again, “Canada” → “canada” and “캐나다” → “kaenada” produce “ada” as their shared token. Since a sequence of characters is more likely to contain semantic meaning than a single character, this distinction is designed to decompose the effect of sharing the surface form and sharing their associated meanings.
- **Shared Phonology.** Widely used transliteration methods (e.g., G2P and romanization) encode phonological information in their representations. The main motivation for this is that they are more likely to have similar representations for lexical items that sound similar. We consider the extent to which they capture the phonology of each language as one

of the main factors behind the effectiveness of transliteration, since it enables the recognition of cognates and borrowed vocabulary shared across languages. For example, G2P conversion captures similarities in pronunciation between borrowed words: English “smartphone” → “smɑ:rtfɒn” and Korean “스마트폰” → “sumat^huɪp^hɒn”.

To explore these different dimensions of transliteration, we employ four distinct input types: Orthography (Ortho), IPA, Romanized text (Rom), and Substitution Ciphred text (Cipher). How each input type is associated with each factors is summarized in Table 1. The following subsections provide a detailed explanation of each input type and the process of converting written text data (Ortho) into the other input types.

3.1 Ortho: Standard Orthography

Different languages use different scripts as their standard orthographies. For example, Korean uses Hangul (한글) and English is written in Latin alphabet. Because models often suffer from significant performance degradation when processing unfamiliar scripts (i.e., *script barrier*), this standard input type serves as a baseline for evaluating the impact of different transliteration methods.

3.2 IPA: G2P Conversion

Based on Latin scripts, IPA symbols are designed to represent pronunciations of human language as phonemes. It is therefore considered to be the most accurate method at representing shared phonology. While transliteration into IPA enables some degree of character set sharing, differences in phonemic inventories and phonotactic structures across languages cause each language to use its own distinct set of characters and subword tokens. For example, difference in their phonemic inventories cause ‘hotel’ to be converted into /otel/ in Spanish and /howtɛl/ in English.

To convert orthographic data into IPA symbols, we use Epitran (Mortensen et al., 2018). It is a rule-based G2P tool that supports more than a hundred languages, and is widely adopted in transliteration-based methods (Bharadwaj et al., 2016; Chaudhary et al., 2018; Leong and Whitenack, 2022; Sohn et al., 2024; Zhu et al., 2024). We chose a rule-based conversion framework to ensure both consistency and practicality, as such rules are well-established for many languages and the conversion process is much faster.

3.3 Rom: Romanization

Romanization is the process of converting various scripts into Latin letters, enforcing a stricter limit in character sets used across languages. Additionally, unlike G2P, which often adds language-specificity to languages originally using Latin scripts, romanization preserves the written form of Latin script. This allows the resulting text to retain the advantages of a shared script, increasing the likelihood of sharing more subword tokens. Since Latin scripts encode sound—though not as precisely as IPA—romanization produces phonologically informed representations for each language.

For romanization, we employ Uroman (Hermjakob et al., 2018) which supports conversion of any UTF-8-encoded script into Latin script. The tool has been widely used in prior work on romanization and shown to be robust (Amrhein and Sennrich, 2020; Liu et al., 2024; Purkayastha et al., 2023).

3.4 Cipher: Substitution Cipher

A substitution cipher is a method from cryptography where units of plaintext are replaced with ciphertext according to a predefined rule or key. We apply substitution cipher to the romanized text of each language—with different rules for each—to remove encoded phonological information. While this allows multilingual text to share the same character space as Rom, it no longer encodes cross-lingual phonological patterns and prevents the sharing of meaningful subword tokens across languages.

We employ a Caesar cipher, a simple substitution encryption technique that shifts each letter in the text by a fixed number of positions in the Latin alphabet. For each language, we assign an integer that determines the shift from the current position of each letter. For example, if English is assigned the integer 4, the word ‘apple’ would be represented as ‘ettpi’, with each letter replaced by the one that is four positions ahead in the alphabet. Details are provided in Appendix A.4.

4 Experimental Setup

4.1 Language Selection

To examine how different input types impact multilingual adaptation, we selected languages to form four language sets: (i) typologically similar languages using the same script (sim-same), (ii) similar languages using diverse scripts (sim-div), (iii)

| Similarity | Script Diversity | Language | Language Family | Script |
|------------|------------------|----------|-----------------|--------|
| similar | same | swe | Indo-European | Latn |
| | | por | Indo-European | Latn |
| | | lij | Indo-European | Latn |
| | | cat | Indo-European | Latn |
| | | ron | Indo-European | Latn |
| | | spa | Indo-European | Latn |
| | | sqi | Indo-European | Latn |
| similar | div | fra | Indo-European | Latn |
| | | ben | Indo-European | Beng |
| | | hin | Indo-European | Deva |
| | | hrv | Indo-European | Latn |
| | | ori | Indo-European | Odia |
| | | rus | Indo-European | Cyrl |
| | | srp | Indo-European | Cyrl |
| dissimilar | same | ilo | Austronesian | Latn |
| | | sna | Niger-Congo | Latn |
| | | lav | Indo-European | Latn |
| | | uzb | Turkic | Latn |
| | | deu | Indo-European | Latn |
| | | fin | Uralic | Latn |
| | | som | Afroasiatic | Latn |
| swa | Niger-Congo | Latn | | |
| dissimilar | div | amh | Afroasiatic | Ethi |
| | | ben | Indo-European | Beng |
| | | tel | Dravidian | Telu |
| | | fra | Indo-European | Latn |
| | | tha | Tai-Kadai | Thai |
| | | kat | Caucasian | Geor |
| | | kor | Koreanic | Hang |
| mya | Sino-Tibetan | Mymr | | |

Table 2: Languages selected for each language set.

dissimilar languages using the same script (dissim-same), and (iv) dissimilar languages using diverse scripts (dissim-div). Similar to Chang et al. (2024), we utilized lang2vec (Littell et al., 2017)¹ to compute language similarity. We extracted syntactic, geographic, and genetic feature vectors from lang2vec to obtain cosine similarities, and also defined lexical similarity based on word overlap ratio between training corpora of each language². We assigned eight languages to each set (see Table 2) by sampling languages based on the aggregated scores. More details are provided in Appendix A.1.

4.2 Vocabulary Overlap Ratio

As discussed in Section 3, different transliteration methods are likely to produce different patterns of token overlap. We therefore focus on analyzing overlaps between target language and pretraining languages to understand how each transliteration

¹Utilizing <https://github.com/antonisa/lang2vec>

²Words are segmented by white spaces.

method benefits. Throughout the analysis, we compute the overlap ratio as follows:

$$\text{OverlapRatio}(l_t, L_s) = \max_{l \in L_s} \frac{|S_l \cap S_{l_t}|}{|S_{l_t}|} \quad (1)$$

where L_s is a set of pre-trained languages, l_t the target language, and S_l the set of tokens from the language l . We suppose all sets S that are used to compute overlap ratio is based on the same tokenizer. We take overlap ratio with that of the

We further break down the overlaps by length to examine which lengths contribute most to the results through shared tokens. The overlap ratio for tokens of length m is computed as

$$\frac{|\{x \in S_{l_s} \cap S_{l_t} \mid \text{len}(x) = m\}|}{|S_{l_t}|} \quad (2)$$

where $l_s = \arg \max_{l \in L_s} \frac{|S_l \cap S_{l_t}|}{|S_{l_t}|}$.

We compute the overlap ratio using the pre-trained language that exhibits the greatest overlap. Counting token overlaps across all pre-trained languages would, in most cases, cover nearly all tokens from the target language, since the tokenizer is trained on those languages. For a more meaningful analysis, we therefore focus on the single pre-trained language that the target language is most likely to leverage for our main analysis. Analyses based on overlaps with all pre-trained languages are provided in the Appendix A.7.

4.3 Datasets

For pre-training, we utilize sampled version of a preprocessed Wikipedia corpus from Hugging Face.³ For downstream tasks, we utilized WikiAnn (Pan et al., 2017; Rahimi et al., 2019) dataset for NER, Universal Dependencies (v2.17) (Nivre et al., 2020) for POS, and XNLI (Conneau et al., 2018) for sentence classification (NLI) task. More details on preprocessing and dataset statistics can be found in Appendix A.12. In order to train the model with different input types, we converted all datasets into each input type using the aforementioned tools.

4.4 Model Training

To investigate the impact of different input types, we pre-train 16 models from scratch using four input types and four language sets. We avoid using publicly available pre-trained models to ensure a

³<https://huggingface.co/datasets/wikimedia/wikipedia>

controlled experimental setup, as most such models are trained on orthographic data, which would confound the analysis of transliteration effects.

For our controlled experiments, we first trained a SentencePiece tokenizer (BPE) for each model with fixed vocabulary size of 30K for all tokenizers. For model architecture, we employed Transformer encoder architecture as in XLM-R (Conneau et al., 2020)—a widely known model for its compelling multilingual performances in various tasks. As with XLM-R, we follow the training regime used for RoBERTa (Liu et al., 2019) with masked language modeling on each multilingual corpus. After pre-training, we fine-tune each model on target language dataset to evaluate its downstream performance. For details on the model configurations and training, refer to Appendix A.2 and Appendix A.3.

5 Results

Table 3 presents average scores across target languages for the two downstream tasks—F1 scores for NER and average accuracies on XNLI—for seen and unseen languages⁴. P -values obtained from paired t -tests on F1 scores across input types can be found in Appendix A.5.

NER Results We observe that Rom outperforms other input types on average for unseen languages, which simulate the script barrier. Even for seen languages, Rom performs comparably to Ortho. To further verify the significance of the results, we conduct paired t -tests to examine the statistical significance of differences between input types. We find that Rom yields a significant advantage over all other input types for unseen languages ($p < 0.05$), whereas Ortho performs worse than all other input types ($p < 0.05$). Interestingly, IPA and Cipher do not differ significantly, despite Cipher encoding no shared linguistic information across languages.

POS Results Results on the POS task exhibit similar trends to those observed for NER, with Rom outperforming all other input types even on seen languages, and Ortho significantly underperforming on unseen languages. We additionally report the results of paired t -tests in Appendix A.5.

XNLI Results For XNLI, we evaluated on 11 languages, most of which were unseen during pre-training. Similar to NER, transliteration does not improve performance on seen languages, with no

⁴Unseen languages refer to languages not included in pre-training of each model.

| Trained Lang. Set | Input Type | NER | | | POS | | | NLI | | |
|-------------------|------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | | All | Seen | Unseen | All | Seen | Unseen | All | Seen | Unseen |
| sim-same | # Lang. | 28 | 8 | 20 | 19 | 7 | 12 | 11 | 2 | 9 |
| | Ortho | 0.7141 | 0.8466 | 0.6611 | 0.8403 | 0.9513 | 0.7755 | 0.5847 | 0.6768 | 0.5642 |
| | IPA | 0.7168 | 0.8085 | 0.6801 | 0.8959 | 0.9386 | 0.8709 | 0.6057 | 0.6605 | 0.5936 |
| | Rom | 0.7589 | 0.8395 | 0.7267 | 0.9088 | 0.9499 | 0.8848 | 0.6276 | 0.6755 | 0.6170 |
| | Cipher | 0.7210 | 0.8173 | 0.6824 | 0.9042 | 0.9463 | 0.8796 | 0.6133 | 0.6658 | 0.6017 |
| sim-div | # Lang. | 28 | 8 | 20 | 19 | 5 | 14 | 11 | 4 | 7 |
| | Ortho | 0.6917 | 0.8409 | 0.6321 | 0.8150 | 0.9501 | 0.7667 | 0.6051 | 0.6242 | 0.5942 |
| | IPA | 0.7202 | 0.8239 | 0.6787 | 0.8942 | 0.9446 | 0.8762 | 0.6171 | 0.6192 | 0.6160 |
| | Rom | 0.7522 | 0.8451 | 0.7151 | 0.9020 | 0.9511 | 0.8844 | 0.6248 | 0.6273 | 0.6233 |
| | Cipher | 0.7200 | 0.8270 | 0.6772 | 0.8953 | 0.9473 | 0.8767 | 0.6128 | 0.6104 | 0.6142 |
| dissim-same | # Lang. | 28 | 7 | 21 | 19 | 4 | 15 | 11 | 2 | 9 |
| | Ortho | 0.6935 | 0.7860 | 0.6626 | 0.8189 | 0.8950 | 0.7986 | 0.6000 | 0.6243 | 0.5947 |
| | IPA | 0.7534 | 0.7732 | 0.7468 | 0.8737 | 0.8673 | 0.8754 | 0.6093 | 0.6307 | 0.6045 |
| | Rom | 0.7455 | 0.7981 | 0.7280 | 0.8912 | 0.8950 | 0.8902 | 0.6155 | 0.6195 | 0.6146 |
| | Cipher | 0.7592 | 0.7725 | 0.7547 | 0.8832 | 0.8913 | 0.8810 | 0.5995 | 0.6296 | 0.5928 |
| dissim-div | # Lang. | 28 | 8 | 20 | 19 | 5 | 14 | 11 | 2 | 9 |
| | Ortho | 0.7436 | 0.7402 | 0.7450 | 0.8399 | 0.8819 | 0.8249 | 0.5670 | 0.6481 | 0.5490 |
| | IPA | 0.7524 | 0.7524 | 0.7524 | 0.8991 | 0.8925 | 0.9014 | 0.6220 | 0.6039 | 0.6260 |
| | Rom | 0.7748 | 0.7538 | 0.7832 | 0.9059 | 0.8928 | 0.9106 | 0.6327 | 0.6172 | 0.6361 |
| | Cipher | 0.7502 | 0.7518 | 0.7496 | 0.8987 | 0.8896 | 0.9019 | 0.6268 | 0.6176 | 0.6288 |

Table 3: Downstream task performances for each model—F1 scores for NER and POS, and accuracy scores for NLI. Average scores for all, seen, and unseen languages are reported for each task. **Bold** indicates best performing input type per language set.

statistically significant differences between input types. However, for unseen languages, all transliterated input types outperform Ortho, with Rom consistently showing the strongest performance across all language sets. Paired *t*-tests examining the differences between input types support these findings (See Appendix A.5).

6 Analysis and Discussion

In this section, we analyze how each factor of transliteration—shared characters, shared (subword) tokens, and shared phonology—is associated with the downstream performance. Given that transliteration benefits unseen languages rather than seen ones, our analysis focuses on unseen languages unless otherwise noted.

6.1 Shared Character Set: Overcoming UNK Tokens

Figure 2a shows that the script barrier is closely related to the proportion of unknown (UNK) tokens, which we compute as the number of UNK tokens divided by the total number of tokens in the evaluation data. A UNK token occurs when the tokenizer fails to segment a sequence of characters using

its learned vocabulary, and high UNK ratios arise primarily in languages written in unseen scripts, where the tokenizer has no prior exposure to the characters.

Consequently, applying transliteration and sharing characters across languages greatly reduces UNK ratio in unseen languages (See Figure 3), leading to clear performance gains over Ortho as shown in Figure 2b. To this end, **shared character set** serves as the initial means by which transliteration overcomes the script barrier, reducing the proportion of UNK tokens. Results from Cipher further support this interpretation: despite containing no semantic or linguistic information shared across languages, its simple character sharing yields notable performance gains for languages written in unseen scripts compared to Ortho.

6.2 Sharing Longer Token Matters

While reducing the UNK ratio substantially mitigates the script barrier for languages written in unseen scripts, the correlation between UNK ratio and downstream performance becomes less pronounced across different transliteration methods—IPA, Rom, and Cipher. To further investigate what

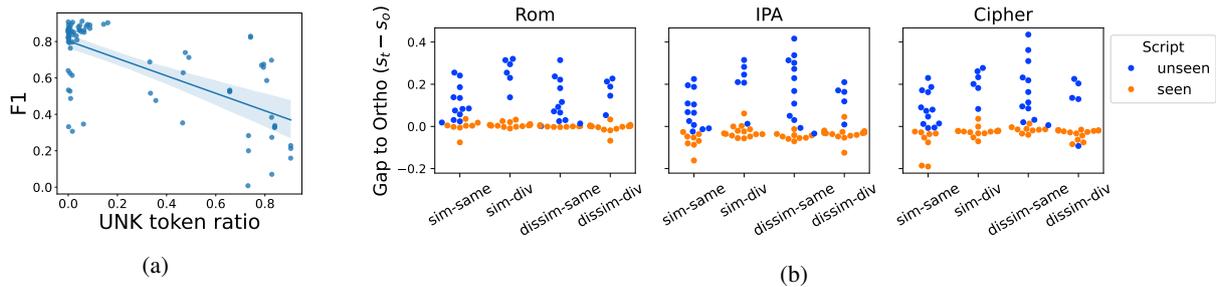


Figure 2: (a) Negative correlation between UNK token ratio and F1 score. (b) Performance gains of transliterated input types compared to orthography-based models, where s_t and s_o denote scores with transliterated and orthographic inputs, respectively. Performance gains appear primarily in target languages whose original scripts are unseen during pre-training.

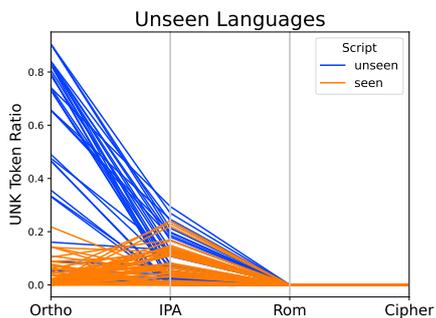


Figure 3: Unknown (UNK) token ratio for unseen languages across different input types.

drives these remaining differences, we analyze token overlaps across languages under each transliteration method, examining how lengths of overlapping tokens relate to downstream performance.

Effect of Longer Shared Tokens Figure 4 shows the Pearson correlation between token-overlap ratios of different lengths and downstream performance. The overlap ratio by length is computed as described in Equation (2). As shown in Figure 4, overlaps of relatively shorter tokens (including a single character) correlate negatively with performance, whereas overlaps of longer tokens show positive correlations. Notably, negative correlation between character-level overlaps and the performances indicates that while sharing characters alleviates the script barrier by reducing UNK tokens, excessive character-level overlap can be detrimental to downstream performance. One possible explanation is that shorter tokens, which tend to vary more in meaning across contexts, may confuse the model when shared across languages, whereas longer tokens are more likely to provide more stable and consistent semantic cues across languages.

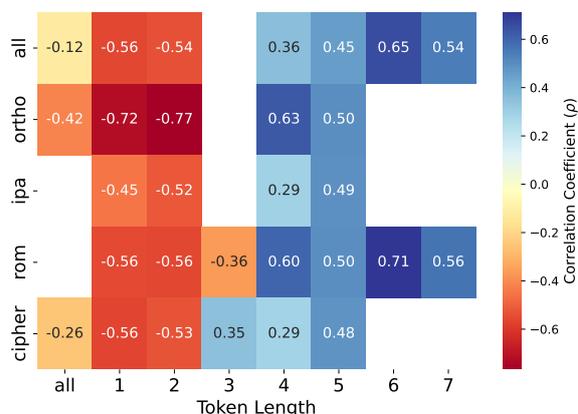


Figure 4: Correlation between token overlap ratio by length and downstream performance. Correlations with $p > 0.05$ are masked.

Romanization and Token Length We next examine how this observation relates to the notable success of romanization. Our hypothesis is that romanization produces longer tokens, which in turn positively affect performance due to their better consistency across languages. To test this hypothesis, we examine the distribution of unique token lengths produced by each model and find that Rom generates the largest proportion of longer tokens. Figure 5 shows the distribution of token lengths for models trained on sim-div languages, with results for other language sets provided in the Appendix A.8. We attribute this to the shared phonology introduced by transliteration, which we discuss further in Section 6.3.

Vocabulary Coverage To better understand the impact of longer shared tokens, we additionally analyze what we define as *vocabulary coverage* and how each token length contribute to it. We define vocabulary coverage as the ratio of unique tokens produced to the total vocabulary size of the tok-

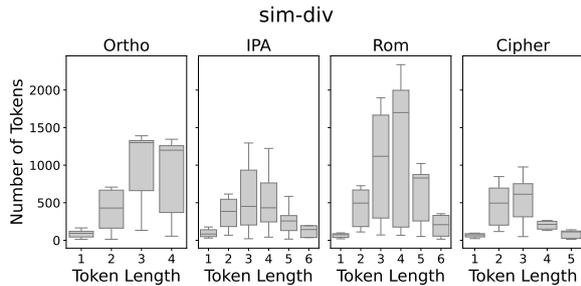


Figure 5: Number of unique tokens by length for unseen target languages, using models trained on sim-div languages.

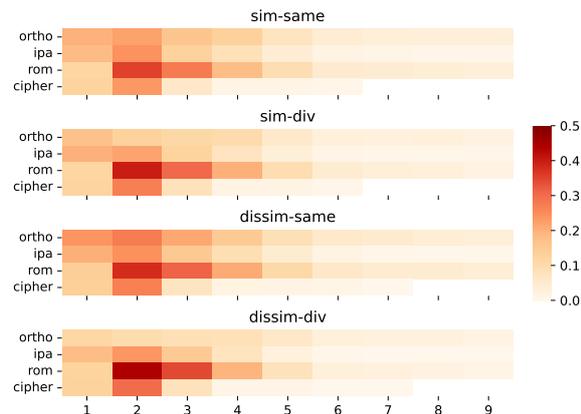


Figure 6: Average vocabulary coverage of each model on unseen languages. The x-axis indicates token length. Romanization shows higher coverage for tokens of length 2-4, whereas other input types exhibit more limited utilization of their vocabulary space.

enizer. A higher coverage indicates that a larger portion of the model’s embedding space is being used, allowing more effective utilization of its capacity through a greater share of token embeddings.

Figure 6 shows the average vocabulary coverage across unseen languages. It clearly illustrates that Rom achieves higher coverage than other input types, especially for tokens longer than two characters. Figure 7 further presents how the contribution of each token length to overall vocabulary coverage varies across input types. Rom clearly produces a greater variety of longer tokens than any other input type. This broader token usage ensures greater vocabulary coverage and thereby better leverages the model’s capacity.

Longer Tokens and Vocabulary Coverage

Building on the observation of how tokens of different lengths contribute to vocabulary coverage, we further interpret the advantage of overlapping longer tokens in terms of their combinatorial poten-

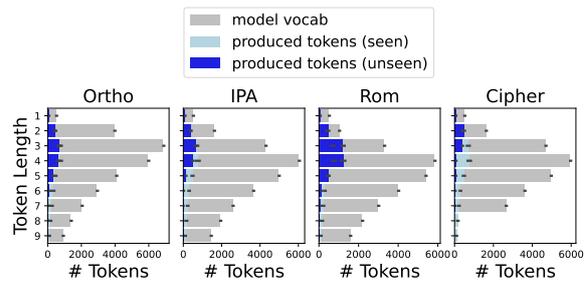


Figure 7: Number of tokens produced by each input type for seen and unseen languages. Rom utilizes a larger proportion of the model’s vocabulary, whereas other input types show more limited token usage.

| Input Type | Fertility | | Vocab. Coverage | |
|------------|-----------|----------------------|-----------------|-----------------------|
| | Coef. | p | Coef. | p |
| Ortho | -0.2547 | 0.0008 | 0.5236 | 4.6×10^{-12} |
| IPA | -0.0176 | 0.8162 | 0.3248 | 1.8×10^{-5} |
| Rom | -0.0797 | 0.2927 | 0.3705 | 9.9×10^{-7} |
| Cipher | -0.1145 | 0.1303 | 0.3405 | 6.8×10^{-6} |
| All | -0.1628 | 1.2×10^{-5} | 0.4240 | 5.0×10^{-30} |

Table 4: Spearman’s correlation coefficients between downstream performance and two tokenizer-related factors: fertility and vocabulary coverage.

tial. While single-character overlaps are bounded by the size of the alphabet or the number of unique characters available, longer tokens allow exponentially more possible combinations. This provides a greater potential to span a larger portion of the vocabulary, thereby enabling more effective utilization of the model’s capacity.

We also examine whether tokenizer fertility (Ács, 2019; Rust et al., 2021) explains the gains from transliteration. Fertility, defined as the average number of subword tokens a tokenizer generates per word, is commonly used to assess tokenizer quality. Prior work (Rust et al., 2021) suggests that lower fertility is associated with better multilingual performance. Using our trained tokenizers, We compute fertility scores on the NER datasets and analyze their correlation with downstream performance. As shown in Table 4, fertility only partially explains the performance gains and does not fully account for differences across input types. In contrast, vocabulary coverage exhibits a stronger and more consistent correlation—not only across input types, but also across languages within each input type. These results show that the benefit of sharing longer tokens for unseen languages cannot be explained solely by conventional tokenizer quality.

6.3 Shared Phonology: A Path to Longer Tokens

Here, we examine the role of shared phonology introduced by transliteration through the behavior of Cipher. Cipher is designed to isolate the effect of shared phonology from Rom. It uses the same character set as in Rom, but does not encode phonological—or any linguistic—information shared across languages.

Without shared phonology, Cipher struggles to produce longer tokens as shown in Figure 5, and consequently performs worse than Rom. Nevertheless, by eliminating unknown tokens, it performs comparably to IPA and surpasses Ortho on unseen languages. In contrast, although IPA has a relatively high proportion of UNK tokens compared to other transliteration methods, it produces longer shared tokens in unseen languages. These observations suggest that the shared character set alone is insufficient for effective transliteration; rather, a consistent form-meaning mapping across languages is crucial to enable models to form longer tokens, typically through shared phonology.

7 Conclusion

Transliteration has been used to bring together diverse languages using diverse scripts, but its benefits have not been well understood. We examined three aspects of transliteration—shared characters, shared subword tokens, and shared phonology—and analyzed their roles in adapting to languages and scripts unseen during pre-training, by comparing four different input types: Ortho, IPA, Rom, and Cipher. We find that Rom outperforms all other input types, with a greater number of longer tokens shared across languages. While shared characters largely facilitate the adaptation by preventing unknown tokens, we suggest that sharing longer tokens is crucial for effective transliteration that overcomes the script barrier. Our findings suggest that transliteration is effective not simply because of language similarity with pre-trained languages, but because it reshapes token distributions in ways that make multilingual models more adaptable and capable.

8 Limitation

The results reported here are suggestive, but there are two major limitations which prevent us from generalizing them too broadly. First, we only tested one type of transformer model with only subword

tokenization scheme. It is possible, for example, that we would have obtained much different results if we had trained character- or byte-level models. Also, we only tested one romanizer and one G2P transducer. It is possible that the results were influenced by the performance of each tool.

9 Ethics Statement

We believe that this research raises no significant ethical concerns or violations of the code of ethics mandated by the Association for Computational Linguistics. The data used in this study, all of which are publicly available, were collected in accordance with legal and institutional protocols, to the best of our knowledge. Furthermore, our use of these resources is compatible with the uses intended by the creators.

10 Acknowledgement

This research was supported by Korean Institute for Advancement of Technology(KIAT) grant funded by the Korea Government(MOTIE)(RS-2024-00435997, Human Resource Development Program for Industrial Innovation(Global)). This work was also supported by the Institute of Information & Communications Technology Planning&Evaluation(IITP) grant funded by the Korea government(MSIT) (No. RS-2024-00437337, Development of technology for linking and utilizing security information and event management with privacy-preserving internal security data)

References

- 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.
- Akash Bharadwaj, David Mortensen, Chris Dyer, and Jaime Carbonell. 2016. [Phonologically aware neural model for named entity recognition in low resource transfer settings](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1462–1472, Austin, Texas. Association for Computational Linguistics.
- Tyler A. Chang, Catherine Arnett, Zhuowen Tu, and Ben Bergen. 2024. [When is multilinguality a curse? language modeling for 250 high- and low-resource languages](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 4074–4096, Miami, Florida, USA. Association for Computational Linguistics.

- Aditi Chaudhary, Chunting Zhou, Lori Levin, Graham Neubig, David R. Mortensen, and Jaime Carbonell. 2018. [Adapting word embeddings to new languages with morphological and phonological subword representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3285–3295, Brussels, Belgium. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Wietse de Vries, Martijn Wieling, and Malvina Nissim. 2022. [Make the best of cross-lingual transfer: Evidence from POS tagging with over 100 languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7676–7685, Dublin, Ireland. Association for Computational Linguistics.
- Yoshinari Fujinuma, Jordan Boyd-Graber, and Katharina Kann. 2022. [Match the script, adapt if multilingual: Analyzing the effect of multilingual pretraining on cross-lingual transferability](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1500–1512, Dublin, Ireland. Association for Computational Linguistics.
- Ulf Hermjakob, Jonathan May, and Kevin Knight. 2018. [Out-of-the-box universal Romanization tool uroman](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 13–18, Melbourne, Australia. Association for Computational Linguistics.
- Jaavid Husain, Raj Dabre, Aswanth M, Jay Gala, Thanmay Jayakumar, Ratish Puduppully, and Anoop Kunchukuttan. 2024. [RomanSetu: Efficiently unlocking multilingual capabilities of large language models via Romanization](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15593–15615, Bangkok, Thailand. Association for Computational Linguistics.
- Haeji Jung, Changdae Oh, Joeon Kang, Jimin Sohn, Kyungwoo Song, Jinkyu Kim, and David R Mortensen. 2024. [Mitigating the linguistic gap with phonemic representations for robust cross-lingual transfer](#). In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*, pages 200–211, Miami, Florida, USA. Association for Computational Linguistics.
- David Kahn. 1996. *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet*, revised edition. Scribner, New York.
- Colin Leong and Daniel Whitenack. 2022. [Phone-ing it in: Towards flexible multi-modal language model training by phonetic representations of data](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5306–5315, Dublin, Ireland. Association for Computational Linguistics.
- Tomasz Limisiewicz, Jiří Balhar, and David Mareček. 2023. [Tokenization impacts multilingual language modeling: Assessing vocabulary allocation and overlap across languages](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5661–5681, Toronto, Canada. Association for Computational Linguistics.
- Patrick Littell, David R. Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. 2017. [URIEL and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 8–14, Valencia, Spain. Association for Computational Linguistics.
- Yihong Liu, Chunlan Ma, Haotian Ye, and Hinrich Schuetze. 2024. [TransliCo: A contrastive learning framework to address the script barrier in multilingual pretrained language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2476–2499, Bangkok, Thailand. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *ArXiv*, abs/1907.11692.
- 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.
- David R. Mortensen, Siddharth Dalmia, and Patrick Littell. 2018. [Epitran: Precision G2P for many languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- 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.
- Hoang Nguyen, Chenwei Zhang, Tao Zhang, Eugene Rohrbach, and Philip Yu. 2023. [Enhancing cross-lingual transfer via phonemic transcription integration](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9163–9175, Toronto, Canada. Association for Computational Linguistics.
- Hoang H Nguyen, Khyati Mahajan, Vikas Yadav, Julian Salazar, Philip S. Yu, Masoud Hashemi, and Rishabh Maheshwary. 2025. [Prompting with phonemes: Enhancing LLMs’ multilinguality for non-Latin script languages](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 11975–11994, Albuquerque, New Mexico. Association for Computational Linguistics.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. [Universal Dependencies v2: An evergrowing multilingual treebank collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2021. [UNks everywhere: Adapting multilingual language models to new scripts](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10186–10203, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Fred Philippy, Siwen Guo, and Shohreh Haddadan. 2023. [Towards a common understanding of contributing factors for cross-lingual transfer in multilingual language models: A review](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5877–5891, Toronto, Canada. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. 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.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. [Massively multilingual transfer for NER](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164, Florence, Italy. Association for Computational Linguistics.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? on the monolingual performance of multilingual language models](#). 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 3118–3135, Online. Association for Computational Linguistics.
- Jimin Sohn, Haeji Jung, Alex Cheng, Jooeon Kang, Yilin Du, and David R Mortensen. 2024. [Zero-shot cross-lingual NER using phonemic representations for low-resource languages](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13595–13602, Miami, Florida, USA. Association for Computational Linguistics.
- Jian Zhu, Changbing Yang, Farhan Samir, and Jahu-rul Islam. 2024. [The taste of IPA: Towards open-vocabulary keyword spotting and forced alignment in any language](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 750–772, Mexico City, Mexico. Association for Computational Linguistics.
- Judit Ács. 2019. [Exploring bert’s vocabulary](#). Blog Post.

A Appendix

A.1 Language Selection

To examine the impact on multilingual adaptation that differences in input types have, we selected four language sets : (i) similar languages using the same script (sim-same), (ii) similar languages using diverse scripts (sim-div), (iii) dissimilar languages using the same script (dissim-same), and (iv) dissimilar languages using diverse scripts (dissim-div). These sets were used to train multilingual models with varying linguistic similarities and scripts. For each set, we assigned eight languages based on a computed similarity score as shown in Table 2.

Similar to Chang et al. (2024), we utilized lang2vec (Littell et al., 2017)⁵ to compute language similarity. Specifically, we extracted syntactic, geographic, and genetic features from lang2vec and computed cosine similarities, denoted as s_{syn} , s_{geo} , and s_{gen} in Eq. 3. We also defined lexical similarity s_{lex} , which is obtained by calculating the word overlap ratio between training corpora of each language⁶. Finally, we aggregated all similarity scores (i.e., syntactic, geographic, genetic, and lexical) to derive the overall similarity score between two languages:

$$\text{sim}_s(x, y) = s_{syn}(x, y) + s_{geo}(x, y) + s_{gen}(x, y) + s_{lex}(x, y). \quad (3)$$

With initial set of languages L that are supported by Wikipedia corpus and Epitran, we use average pairwise similarity scores to compute similarity score for a set of languages and obtain an optimal set L_s^* , where $s \in \{\text{sim-same, sim-div}\}$:

$$L_s^* = \arg \max_{\substack{L_s \subset L \\ |L_s|=8}} \left(\frac{1}{|L_s|(|L_s|-1)} \sum_{x \in L_s} \sum_{\substack{y \in L_s \\ y \neq x}} \text{sim}_s(x, y) + \alpha \cdot \left(\mathbb{1}_{s \in \{\text{sim-div}\}} |SC_{L_s}| - \mathbb{1}_{s \in \{\text{dissim-div}\}} |SC_{L_s}| \right) \right) \quad (4)$$

As for an optimal set L_d^* , where $d \in \{\text{dissim-same, dissim-div}\}$:

$$L_d^* = \arg \min_{\substack{L_d \subset L \\ |L_d|=8}} \left(\frac{1}{|L_d|(|L_d|-1)} \sum_{x \in L_d} \sum_{\substack{y \in L_d \\ y \neq x}} \text{sim}_s(x, y) + \alpha \cdot \left(\mathbb{1}_{d \in \{\text{sim-div}\}} |SC_{L_d}| - \mathbb{1}_{d \in \{\text{dissim-div}\}} |SC_{L_d}| \right) \right). \quad (5)$$

To select languages for the sets with same script (i.e., sim-same and dissim-same), we limited the search space to languages that use the Latin script to maximize the number of languages available for similarity-based sampling.

For sets with diverse scripts (i.e., -div), we additionally consider how many different scripts are involved in each set.

A.2 Model Configuration

Table 5 summarizes the key configuration details of our RoBERTa-based model. Number of parameters per model is 109,082,112.

⁵Utilizing <https://github.com/antonisa/lang2vec>

⁶Words are segmented by white spaces.

| Parameter | Value |
|----------------------------|--------|
| Vocabulary Size | 30,000 |
| Hidden Size | 768 |
| Hidden Layers | 12 |
| Attention Heads | 12 |
| Intermediate Size | 3072 |
| Activation Function | GELU |
| Dropout (Hidden/Attention) | 0.1 |
| Max Position Embeddings | 514 |

Table 5: Model Configuration

A.3 Training Setup

To investigate the impact of different input types, we pre-trained and fine-tuned a total of 16 models across four distinct input types and language sets. In addition, we trained a SentencePiece BPE tokenizer for each model, fixing the vocabulary size to 30K. Table 6 summarizes the key hyperparameters used in our experiments for both the pretraining phase and the downstream NER task.

Hyperparameter Sweep We conducted grid search to find learning rates that converges or achieves the best results. For pre-training, the search space was $\{1e-5, 2e-5, 3e-5, 5e-5, 1e-4, 2e-4, 3e-4\}$ and for NER, it was $\{3e-5, 5e-5, 1e-4\}$.

| Parameter | Pretraining | NER | XNLI | POS |
|-----------------------------|-------------|------------------|------------------|------------------|
| FP16 Training | True | True | True | True |
| Max Seq. Length | 512 | 512 | 512 | 512 |
| Batch Size (per device) | 64 | 64 | 64 | 128 |
| Gradient Accumulation Steps | 1 | - | - | - |
| Warmup Steps | 50 | - | - | - |
| Learning Rate | 1e-4 | 5e-5 | 3e-5 | 5e-5 |
| Weight Decay | 0.01 | 0.01 | 0.01 | 0.01 |
| LR Scheduler Type | Linear | - | - | - |
| MLM Probability | 0.15 | - | - | - |
| Epochs | 300 | 20 | 5 | 20 |
| Log Interval | - | 1 | 1 | 1 |
| | 4 | 2 | 1 | 1 |
| GPU Resources | NVIDIA L40S | NVIDIA RTX A6000 | NVIDIA RTX A6000 | NVIDIA RTX A6000 |

Table 6: Training Configurations

A.4 Substitution Cipher (Cipher)

A substitution cipher is a method from cryptography where units of plaintext are replaced with ciphertext according to a predefined rule or key. We apply substitution cipher to the Romanized text to remove encoded phonological information.

Specifically, we use the Caesar cipher (Kahn, 1996), a simple substitution encryption technique that shifts each letter in the text by a fixed number of positions in the Latin alphabet. For each

language, we assign an integer that determines the shift from the current position of each letter. For example, if English is assigned the integer 4, the word ‘apple’ would be represented as ‘ettpi’, with each letter replaced by the one four positions ahead in the alphabet.

A.5 *P*-values of Paired *t*-tests

Table 3 presents the scores of NER and XNLI tasks for different input types across various language settings. To assess the significance of the observed differences, we performed paired *t*-tests. To be specific, we run a paired, two-sided *t*-test, to allow for differences in either direction. For each pair of input types (e.g., (Rom, Ortho), (Ortho, Cipher), etc.), we take the scores for all languages available in the task. The number of observations for each test varies depending on the languages available for each task. For example, in NER, the *t*-test for unseen languages (*p*-values in the upper triangle in Figure 8 (a)) uses 81 observations, and for seen languages (*p*-values in the lower triangle of the same figure), 31. We select observations in the same manner in other tasks as well. Figure 8 displays the corresponding *p*-values derived from these tests. The results show that transliteration (IPA, Rom, Cipher) significantly benefits Ortho for unseen languages. For NER, Rom and Ortho perform on par for seen languages, while there is no significant difference across any input types for XNLI.

A.6 Multi-seed Results

We conducted experiments with multiple random seeds for NER and POS to improve the reliability of the results. For each training configuration, we trained the models three times and averaged the F1 scores. The results show patterns consistent with our main observations (see Table 7).

A.7 Vocabulary Overlap Analysis

Figure 9 shows heatmap of correlation coefficients between overlap ratio and NER F1 score, where overlap ratio is computed differently from Section 4.2. Figure 10 shows the correlation analysis for XNLI task, and Figure 11 for POS task.

A.8 Number of Unique Tokens

Figure 12 shows the distribution of token lengths for models trained on each language set.

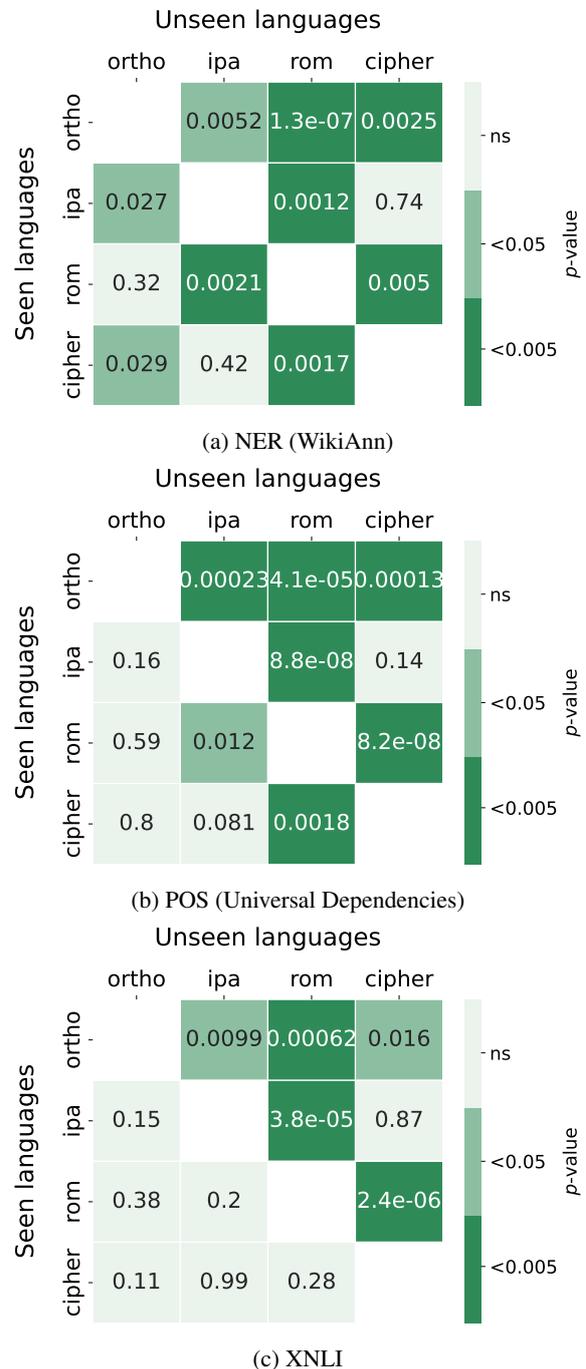
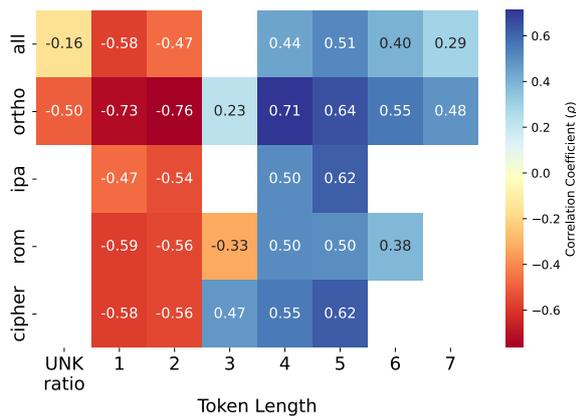
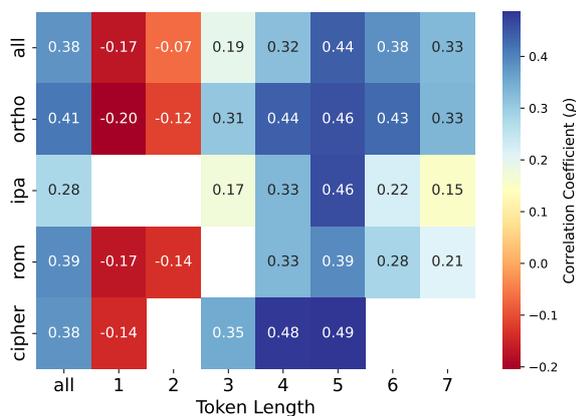


Figure 8: *P*-values for paired *t*-tests on performances for all tasks. Upper triangular elements show results across unseen languages and lower triangular elements show those across seen languages.



(a)



(b)

Figure 9: Spearman's correlation coefficients between overlap ratio and NER F1 score. (a) Computed by type ratio. Leftmost column indicates correlations with UNK token ratio. (b) Computed by overlap ratio across all source languages. Leftmost column indicates correlations with total number of overlaps regardless of token length.

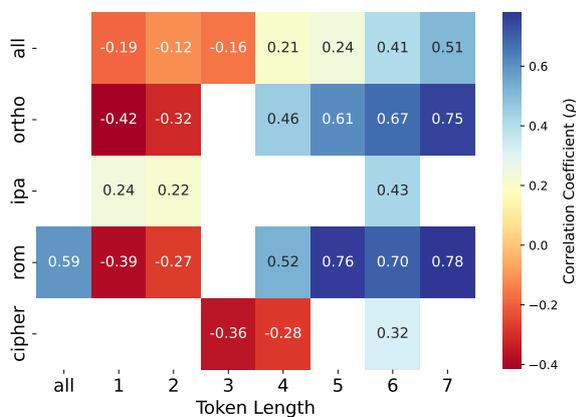


Figure 10: Spearman's correlation coefficients between overlap ratio and XNLI Accuracy. Overlap ratio computed as in Section 4.2.

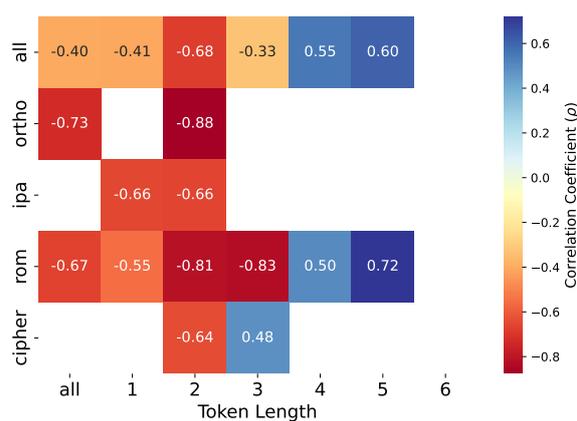
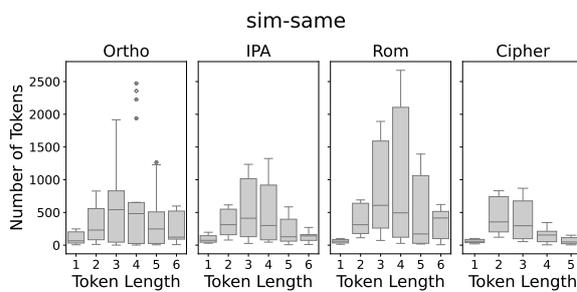
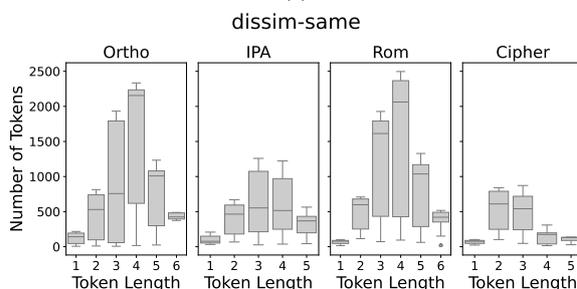


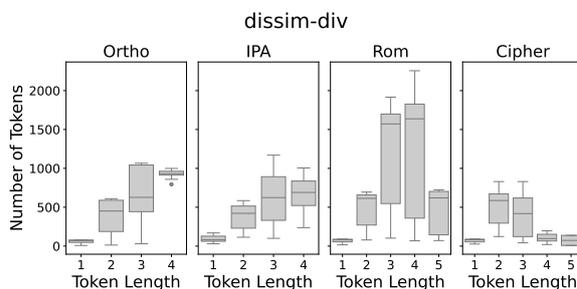
Figure 11: Spearman's correlation coefficients between overlap ratio and POS F1 score. Overlap ratio computed as in Section 4.2.



(a)



(b)



(c)

Figure 12: Number of unique tokens by length for unseen target languages (WikiAnn dataset), using models trained on (a) sim-same (b) dissim-same (c) dissim-div languages.

| Trained Lang. Set | Input Type | NER | | | POS | | |
|-------------------|------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | | All | Seen | Unseen | All | Seen | Unseen |
| sim-same | # Lang. | 28 | 8 | 20 | 19 | 7 | 12 |
| | Ortho | 0.6954 | 0.8277 | 0.6425 | 0.8336 | 0.9471 | 0.7673 |
| | IPA | 0.7046 | 0.7984 | 0.6670 | 0.8936 | 0.9368 | 0.8684 |
| | Rom | 0.7438 | 0.8340 | 0.7077 | 0.9055 | 0.9457 | 0.8820 |
| | Cipher | 0.7084 | 0.8089 | 0.6682 | 0.9009 | 0.9450 | 0.8751 |
| sim-div | # Lang. | 28 | 8 | 20 | 19 | 5 | 14 |
| | Ortho | 0.6817 | 0.8281 | 0.6232 | 0.8107 | 0.9499 | 0.7610 |
| | IPA | 0.7047 | 0.8131 | 0.6614 | 0.8913 | 0.9443 | 0.8724 |
| | Rom | 0.7437 | 0.8379 | 0.7061 | 0.8981 | 0.9484 | 0.8802 |
| | Cipher | 0.7041 | 0.8176 | 0.6587 | 0.8900 | 0.9462 | 0.8700 |
| dissim-same | # Lang. | 28 | 7 | 21 | 19 | 4 | 15 |
| | Ortho | 0.6895 | 0.7902 | 0.6559 | 0.8180 | 0.8946 | 0.7976 |
| | IPA | 0.7109 | 0.7654 | 0.6928 | 0.8726 | 0.8666 | 0.8742 |
| | Rom | 0.7338 | 0.7873 | 0.7159 | 0.8911 | 0.8951 | 0.8901 |
| | Cipher | 0.7234 | 0.7614 | 0.7113 | 0.8836 | 0.8908 | 0.8817 |
| dissim-div | # Lang. | 28 | 8 | 20 | 19 | 5 | 14 |
| | Ortho | 0.7312 | 0.7239 | 0.7341 | 0.8406 | 0.8837 | 0.8252 |
| | IPA | 0.7414 | 0.7394 | 0.7422 | 0.8996 | 0.8921 | 0.9023 |
| | Rom | 0.7627 | 0.7339 | 0.7742 | 0.9061 | 0.8909 | 0.9115 |
| | Cipher | 0.7406 | 0.7403 | 0.7407 | 0.8994 | 0.8905 | 0.9026 |

Table 7: Multi-seeded results for NER and POS. Results are averaged over three runs with different random seeds. **Bold** indicates best performing input type per language set.

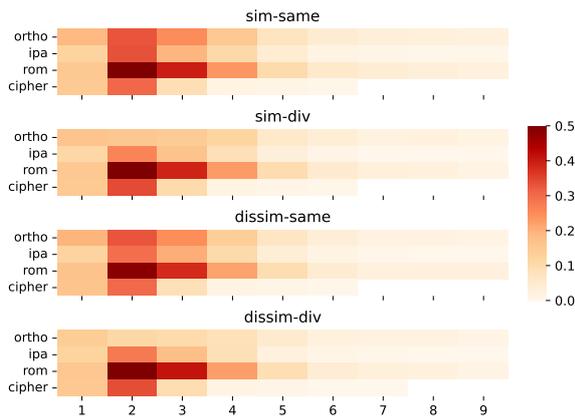


Figure 13: Vocabulary coverage over unseen languages on XNLI dataset computed as in Section 6.2.

A.9 Vocabulary Coverage

Here we provide more plots for vocabulary coverage. Figures 13 and 14 shows vocabulary coverage by length over unseen languages using XNLI dataset and UD Treebank, respectively. The pattern is similar to that of NER (WikiAnn) dataset, where Rom exhibits evidently high coverage compared to other input types. Figure 15 shows vocabulary coverage by length over seen languages using WikiAnn dataset. While Rom still shows the most coverage, Ortho also takes large portion when the languages share same script (i.e., sim-same, dissim-same).

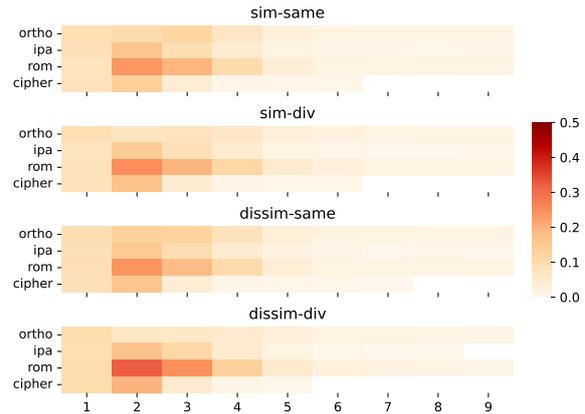


Figure 14: Vocabulary coverage over unseen languages on UD Treebank computed as in Section 6.2.

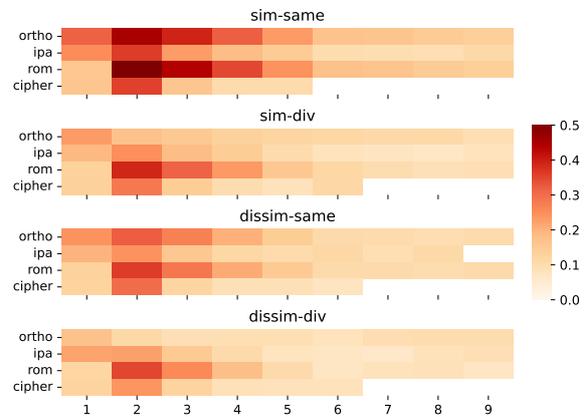


Figure 15: Vocabulary coverage over seen languages on WikiAnn dataset computed as in Section 6.2.

| Trained Lang. Set | Input Type | All | Seen | Unseen |
|-------------------|------------|--------|--------|--------|
| sim-same | Ortho | 0.5472 | 0.7717 | 0.4575 |
| | IPA | 0.5777 | 0.7024 | 0.5278 |
| | Rom | 0.7092 | 0.7908 | 0.6766 |
| | Cipher | 0.6162 | 0.6694 | 0.5949 |
| sim-div | Ortho | 0.6781 | 0.7894 | 0.6336 |
| | IPA | 0.5765 | 0.6708 | 0.5387 |
| | Rom | 0.7042 | 0.8204 | 0.6578 |
| | Cipher | 0.7433 | 0.8489 | 0.7011 |

Table 8: Average F1 scores for NER task using Unigram.

| Trained Lang. Set | Input Type | 1 | 2 | 3 | 4 | 5 |
|-------------------|------------|--------|--------|--------|--------|--------|
| sim-same | IPA | 0.0302 | 0.2635 | 0.1636 | 0.1014 | 0.0570 |
| | Cipher | 0.1491 | 0.3450 | 0.0714 | 0.0147 | 0.0202 |
| sim-div | IPA | 0.0304 | 0.1917 | 0.1592 | 0.0871 | 0.0483 |
| | Cipher | 0.1415 | 0.3543 | 0.0907 | 0.0185 | |

Table 9: Vocabulary coverage of Unigram tokenizer.

A.10 Tokenization Algorithm: Unigram

We conducted additional pre-training experiments using the Unigram tokenizer. We focused on two language sets, sim-same and sim-div, to limit computational cost while preserving diversity in script configurations.

We ran two-sided paired t -test, as we did in the main analyses, to examine the statistical significance of the results. To summarize the statistically significant findings ($p < 0.05$):

- Models trained on sim-same languages set,
 - Rom was significantly better than IPA ($p < 0.005$)
 - Rom was significantly better than Ortho ($p < 0.001$)
- Models trained on sim-div language set,
 - Rom was significantly better than IPA ($p < 0.005$)
 - Cipher was significantly better than IPA ($p < 0.0005$)

We note that Cipher behaves differently under the Unigram tokenizer compared to BPE, in that it often outperforms IPA. To investigate this, we analyzed the vocabulary coverage by token length as in Figure 6, and found that Cipher exhibits notably higher coverage than IPA, consuming more tokens for unseen languages (See Table 9).

We believe this difference arises from how Unigram constructs its vocabulary: rather than relying on greedy merges, it chooses subwords based on probability, resulting in a different token-length distribution in the vocabulary. While our main findings based on token length distributions still hold, this ablation illustrates that different tokenization algorithms may emphasize different aspects of each factor we defined.

A.11 External Tools for Transliteration

In this study, we used Epitran and Uroman as transliteration tools to unify script and facilitate multilingual processing. These tools are widely used for converting text into standardized phonemic or Romanized forms, which aids in cross-lingual learning and transferability. Below, we describe their functionalities and implementation details.

Epitran(Mortensen et al., 2018) is a tool for grapheme-to-phoneme (G2P) conversion, capable of converting text into the International Phonetic Alphabet (IPA) representations. It can be downloaded from the link below

<https://github.com/dmort27/epitran>

Uroman(Hermjakob et al., 2018) is a universal transliteration tool that converts text from various scripts into a Romanized format. It can be downloaded from the link below <https://github.com/isi-nlp/uroman>

A.12 Datasets

In Table 10, the specific number of datasets per corresponding language is provided. For pre-training, we utilized sampled version of preprocessed Wikipedia corpus from Huggingface⁷.

We limited each language with its number of words around 10M⁸. For those languages with less number of tokens than 10M, we kept all the documents and oversampled during training, to match the model’s exposure to all languages. For downstream task, we utilized WikiAnn (Pan et al., 2017; Rahimi et al., 2019) dataset for named entity recognition, Universal Dependencies Treebank (Nivre et al., 2020) for POS, and XNLI dataset for natural language inference (NLI) task. In order to train the model with different input types, we converted all datasets into each corresponding input type.

Wikipedia corpora used for pre-training are licensed under the GNU Free Documentation License (GFDL) and the Creative Commons Attribution-Share-Alike 3.0 License. License type for WikiAnn dataset is ODC-BY.

A.13 Detailed Experimental Results

Tables 11, 12, 13, 14 summarize the performance results (F1 scores) across different language sets under various evaluation settings. In our experiments, "Seen" refers to languages included in both pre-training and fine-tuning, "Unseen" to those entirely absent during training, and "Zero-Shot" to languages evaluated without task-specific fine-tuning. The language sets differ in terms of typological similarity and script usage. Detailed results for each setting are provided in the respective tables.

⁷<https://huggingface.co/datasets/wikimedia/wikipedia>

⁸For each language, we randomly shuffled the order of the documents, and iterated over each document, counting the words segmented by whitespaces. We stop adding the documents when adding the number of words of the last document exceeds 10M.

| Lang | Dataset | # Train | # Validate | # Test | Lang | Dataset | # Train | # Validate | # Test |
|------|-------------|---------|------------|--------|------|-------------|---------|------------|--------|
| am | WikiAnn | 100 | 100 | 100 | my | WikiAnn | 100 | 100 | 100 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | - | - | - | | UD Treebank | - | - | - |
| bn | WikiAnn | 10000 | 1000 | 1000 | or | WikiAnn | 100 | 100 | 100 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | - | - | - | | UD Treebank | - | - | - |
| ca | WikiAnn | 20000 | 10000 | 10000 | pt | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | 13123 | 1709 | 1846 | | UD Treebank | 9616 | 1204 | 1200 |
| de | WikiAnn | 20000 | 10000 | 10000 | ro | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | 392702 | 2490 | 5010 | | XNLI | - | - | - |
| | UD Treebank | 13813 | 799 | 977 | | UD Treebank | 8043 | 752 | 729 |
| es | WikiAnn | 20000 | 10000 | 10000 | ru | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | 392702 | 2490 | 5010 | | XNLI | 392702 | 2490 | 5010 |
| | UD Treebank | 14187 | 1400 | 427 | | UD Treebank | 3850 | 579 | 601 |
| fi | WikiAnn | 20000 | 10000 | 10000 | so | WikiAnn | 100 | 100 | 100 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | 12217 | 1364 | 1555 | | UD Treebank | - | - | - |
| fr | WikiAnn | 20000 | 10000 | 10000 | sq | WikiAnn | 5000 | 1000 | 1000 |
| | XNLI | 392702 | 2490 | 5010 | | XNLI | - | - | - |
| | UD Treebank | 14450 | 1476 | 416 | | UD Treebank | 160 | 20 | 20 |
| hi | WikiAnn | 5000 | 1000 | 1000 | sr | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | 392702 | 2490 | 5010 | | XNLI | - | - | - |
| | UD Treebank | 13306 | 1659 | 1684 | | UD Treebank | - | - | - |
| hr | WikiAnn | 20000 | 10000 | 10000 | sv | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | 6914 | 960 | 1136 | | UD Treebank | 3457 | 1118 | 1121 |
| ilo | WikiAnn | 100 | 100 | 100 | sw | WikiAnn | 1000 | 1000 | 1000 |
| | XNLI | - | - | - | | XNLI | 392702 | 2490 | 5010 |
| | UD Treebank | - | - | - | | UD Treebank | - | - | - |
| ka | WikiAnn | 10000 | 10000 | 10000 | te | WikiAnn | 1000 | 1000 | 1000 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | 870 | 130 | 818 | | UD Treebank | 1051 | 131 | 146 |
| ko | WikiAnn | 20000 | 10000 | 10000 | th | WikiAnn | 20000 | 10000 | 10000 |
| | XNLI | - | - | - | | XNLI | 392702 | 2490 | 5010 |
| | UD Treebank | 4400 | 950 | 989 | | UD Treebank | 2902 | 362 | 363 |
| lij | WikiAnn | 100 | 100 | 100 | ur | WikiAnn | 20000 | 1000 | 1000 |
| | XNLI | - | - | - | | XNLI | 392702 | 2490 | 5010 |
| | UD Treebank | - | - | - | | UD Treebank | 4043 | 552 | 535 |
| lv | WikiAnn | 10000 | 10000 | 10000 | uz | WikiAnn | 1000 | 1000 | 1000 |
| | XNLI | - | - | - | | XNLI | - | - | - |
| | UD Treebank | 15058 | 2110 | 2412 | | UD Treebank | 483 | - | 198 |

Table 10: Statistics of the transliterated datasets for downstream tasks. All datasets exist in four parallel versions: original orthographic, phonemic IPA, Romanized, and cipher-transcribed. A dash (-) indicates unavailable values. Some languages, such as English, Turkish, and Vietnamese, are used only in the XNLI dataset, and the number of items is equal across all XNLI datasets.

| | | Ortho | IPA | Rom | Cipher |
|--------|--------|--------|--------|--------|--------|
| Seen | ca | 0.8997 | 0.8725 | 0.8993 | 0.8803 |
| | es | 0.8773 | 0.8584 | 0.8788 | 0.8657 |
| | fr | 0.8628 | 0.8252 | 0.8639 | 0.8384 |
| | lij | 0.5064 | 0.4052 | 0.4615 | 0.4082 |
| | pt | 0.8798 | 0.8605 | 0.8796 | 0.8674 |
| | ro | 0.9129 | 0.8855 | 0.9103 | 0.8956 |
| | sq | 0.9120 | 0.8738 | 0.8979 | 0.8785 |
| | sv | 0.9215 | 0.8872 | 0.9247 | 0.9046 |
| Unseen | am | 0.2000 | 0.3089 | 0.3383 | 0.3623 |
| | bn | 0.8230 | 0.8907 | 0.9081 | 0.8969 |
| | de | 0.8204 | 0.7400 | 0.8236 | 0.7676 |
| | fi | 0.8573 | 0.8050 | 0.8609 | 0.8237 |
| | hi | 0.7395 | 0.8043 | 0.8225 | 0.7861 |
| | hr | 0.8682 | 0.8318 | 0.8727 | 0.8403 |
| | ilo | 0.6400 | 0.5714 | 0.6757 | 0.4498 |
| | ka | 0.6878 | 0.7920 | 0.8227 | 0.7780 |
| | ko | 0.5329 | 0.7578 | 0.7883 | 0.7626 |
| | lv | 0.8940 | 0.8463 | 0.8919 | 0.8695 |
| | my | 0.2286 | 0.2541 | 0.2857 | 0.2232 |
| | or | 0.2738 | 0.2647 | 0.3492 | 0.3533 |
| | ru | 0.8083 | 0.7842 | 0.8268 | 0.8010 |
| | sn | - | - | - | - |
| | so | 0.6256 | 0.4641 | 0.5500 | 0.4397 |
| | sr | 0.8574 | 0.8442 | 0.8879 | 0.8691 |
| | sw | 0.8250 | 0.7381 | 0.8195 | 0.7494 |
| | te | 0.3384 | 0.5336 | 0.5797 | 0.5252 |
| | th | 0.4762 | 0.6637 | 0.6622 | 0.6477 |
| | ur | 0.9032 | 0.9101 | 0.9273 | 0.9172 |
| uz | 0.8266 | 0.7962 | 0.8402 | 0.7862 | |

Table 11: F1 scores of each model trained on **sim-same** language set. **Seen**: languages used during both pre-training. **Unseen**: languages not encountered during pre-training. Results are provided for four different input types: Orthographic (Ortho), IPA, Romanized (Rom), and Cipher (Cipher).

| | | Ortho | IPA | Rom | Cipher |
|--------|--------|--------|--------|--------|--------|
| Seen | bn | 0.9380 | 0.9375 | 0.9466 | 0.9377 |
| | fr | 0.8436 | 0.8255 | 0.8430 | 0.8378 |
| | hi | 0.8524 | 0.8577 | 0.8394 | 0.8314 |
| | hr | 0.8741 | 0.8527 | 0.8767 | 0.8605 |
| | or | 0.5483 | 0.4981 | 0.5873 | 0.4962 |
| | ru | 0.8395 | 0.8286 | 0.8375 | 0.8304 |
| | sr | 0.8918 | 0.8484 | 0.8969 | 0.8900 |
| | ur | 0.9396 | 0.9424 | 0.9333 | 0.9318 |
| Unseen | am | 0.0079 | 0.2902 | 0.3282 | 0.2695 |
| | ca | 0.8769 | 0.8399 | 0.8801 | 0.8443 |
| | de | 0.7934 | 0.7381 | 0.8047 | 0.7608 |
| | es | 0.8511 | 0.8144 | 0.8573 | 0.8265 |
| | fi | 0.8427 | 0.7993 | 0.8460 | 0.8201 |
| | ilo | 0.5333 | 0.5356 | 0.5537 | 0.4627 |
| | ka | 0.5860 | 0.7961 | 0.8162 | 0.7872 |
| | ko | 0.5244 | 0.7318 | 0.7792 | 0.7577 |
| | lij | 0.3071 | 0.3684 | 0.2975 | 0.3064 |
| | lv | 0.8826 | 0.8468 | 0.8891 | 0.8605 |
| | my | 0.1596 | 0.1721 | 0.2975 | 0.2424 |
| | pt | 0.8535 | 0.8206 | 0.8547 | 0.8312 |
| | ro | 0.8889 | 0.8754 | 0.8963 | 0.8695 |
| | sn | - | - | - | - |
| | so | 0.4874 | 0.4870 | 0.5128 | 0.5236 |
| | sq | 0.8557 | 0.8319 | 0.8604 | 0.8315 |
| | sv | 0.9059 | 0.8583 | 0.9043 | 0.8850 |
| | sw | 0.7634 | 0.7429 | 0.7955 | 0.7359 |
| | te | 0.3297 | 0.5753 | 0.6440 | 0.5119 |
| | th | 0.3531 | 0.6680 | 0.6479 | 0.6302 |
| uz | 0.8384 | 0.7819 | 0.8360 | 0.7863 | |

Table 12: F1 scores of each model trained on **sim-div** language set. **Seen**: languages used during both pre-training. **Unseen**: languages not encountered during pre-training. Results are provided for four different input types: Orthographic (Ortho), IPA, Romanized (Rom), and Cipher (Cipher).

| | | Ortho | IPA | Rom | Cipher |
|--------|--------|--------|--------|--------|--------|
| Seen | de | 0.8184 | 0.7924 | 0.8248 | 0.8095 |
| | fi | 0.8618 | 0.8264 | 0.8638 | 0.8436 |
| | ilo | 0.6757 | 0.7123 | 0.6368 | 0.6549 |
| | lv | 0.8995 | 0.8736 | 0.9006 | 0.8998 |
| | sn | - | - | - | - |
| | so | 0.5551 | 0.5887 | 0.6577 | 0.5556 |
| | sw | 0.8291 | 0.7981 | 0.8421 | 0.8125 |
| | uz | 0.8621 | 0.8210 | 0.8608 | 0.8314 |
| Unseen | am | 0.2833 | 0.5560 | 0.2845 | 0.5018 |
| | bn | 0.8269 | 0.8791 | 0.9005 | 0.9430 |
| | ca | 0.8733 | 0.8255 | 0.8750 | 0.8542 |
| | es | 0.8518 | 0.8103 | 0.8583 | 0.8377 |
| | fr | 0.8312 | 0.7607 | 0.8294 | 0.8447 |
| | hi | 0.7128 | 0.8210 | 0.8055 | 0.7981 |
| | hr | 0.8531 | 0.8404 | 0.8532 | 0.8495 |
| | ka | 0.6289 | 0.8577 | 0.8103 | 0.8606 |
| | ko | 0.5282 | 0.8297 | 0.7652 | 0.8381 |
| | lij | 0.3319 | 0.2893 | 0.3333 | 0.2979 |
| | my | 0.2128 | 0.5263 | 0.2785 | 0.5750 |
| | or | 0.0708 | 0.4082 | 0.3851 | 0.2339 |
| | pt | 0.8566 | 0.8015 | 0.8558 | 0.8449 |
| | ro | 0.8906 | 0.8548 | 0.8880 | 0.8768 |
| | ru | 0.7992 | 0.7922 | 0.8132 | 0.8051 |
| | sq | 0.8658 | 0.8120 | 0.8627 | 0.8259 |
| | sr | 0.8540 | 0.8201 | 0.8790 | 0.8739 |
| | sv | 0.9075 | 0.8484 | 0.9076 | 0.8919 |
| | te | 0.3278 | 0.7441 | 0.5494 | 0.7632 |
| | th | 0.5162 | 0.6841 | 0.6320 | 0.6110 |
| ur | 0.8906 | 0.9208 | 0.9205 | 0.9220 | |

Table 13: F1 scores of each model trained on **dissim-same** language set. **Seen**: languages used during both pre-training. **Unseen**: languages not encountered during pre-training. Results are provided for four different input types: Orthographic (Ortho), IPA, Romanized (Rom), and Cipher (Cipher).

| | | Ortho | IPA | Rom | Cipher |
|--------|--------|--------|--------|--------|--------|
| Seen | am | 0.4941 | 0.5403 | 0.5760 | 0.5364 |
| | bn | 0.9579 | 0.9488 | 0.9552 | 0.9479 |
| | fr | 0.8528 | 0.8265 | 0.8487 | 0.8432 |
| | ka | 0.8647 | 0.8607 | 0.8619 | 0.8598 |
| | ko | 0.7699 | 0.8347 | 0.8382 | 0.8333 |
| | my | 0.5259 | 0.5738 | 0.5164 | 0.5477 |
| | te | 0.7532 | 0.7529 | 0.7528 | 0.7734 |
| | th | 0.7031 | 0.6813 | 0.6810 | 0.6727 |
| Unseen | ca | 0.8797 | 0.8503 | 0.8803 | 0.8513 |
| | de | 0.8088 | 0.7555 | 0.8134 | 0.7855 |
| | es | 0.8615 | 0.8315 | 0.8687 | 0.8352 |
| | fi | 0.8504 | 0.8188 | 0.8532 | 0.8311 |
| | hi | 0.6585 | 0.8223 | 0.8472 | 0.7939 |
| | hr | 0.8642 | 0.8381 | 0.8652 | 0.8428 |
| | ilo | 0.5272 | 0.5726 | 0.5122 | 0.4516 |
| | lij | 0.3465 | 0.3243 | 0.3793 | 0.2833 |
| | lv | 0.8948 | 0.8544 | 0.891 | 0.8762 |
| | or | 0.3840 | 0.3931 | 0.4373 | 0.2913 |
| | pt | 0.8609 | 0.8245 | 0.8630 | 0.8437 |
| | ro | 0.8940 | 0.8746 | 0.8979 | 0.8788 |
| | ru | 0.6753 | 0.7941 | 0.8207 | 0.8049 |
| | sn | - | - | - | - |
| | so | 0.6140 | 0.4893 | 0.5462 | 0.5299 |
| | sq | 0.8720 | 0.8395 | 0.8533 | 0.8389 |
| | sr | 0.6697 | 0.8405 | 0.8826 | 0.8735 |
| | sv | 0.9117 | 0.8641 | 0.9119 | 0.8920 |
| | sw | 0.7968 | 0.7527 | 0.7855 | 0.7536 |
| | ur | 0.6974 | 0.9072 | 0.9243 | 0.9226 |
| uz | 0.8317 | 0.8004 | 0.8300 | 0.8121 | |

Table 14: F1 scores of each model trained on **dissim-div** language set. **Seen**: languages used during both pre-training. **Unseen**: languages not encountered during pre-training. Results are provided for four different input types: Orthographic (Ortho), IPA, Romanized (Rom), and Cipher (Cipher).

| | sim-same | | | | dissim-same | | | | sim-div | | | | dissim-div | | | |
|----|----------|--------|--------|--------|-------------|--------|--------|--------|---------|--------|--------|--------|------------|--------|--------|--------|
| | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher |
| de | 0.6627 | 0.6389 | 0.6719 | 0.6529 | 0.6527 | 0.6513 | 0.6577 | 0.6613 | 0.6425 | 0.6391 | 0.6403 | 0.6483 | 0.6463 | 0.6469 | 0.6673 | 0.6543 |
| en | 0.6822 | 0.6601 | 0.6864 | 0.6667 | 0.6647 | 0.6439 | 0.6814 | 0.6469 | 0.6739 | 0.6569 | 0.6834 | 0.6649 | 0.6930 | 0.6764 | 0.7072 | 0.7030 |
| es | 0.6820 | 0.6691 | 0.6856 | 0.6725 | 0.6601 | 0.6563 | 0.6721 | 0.6487 | 0.6587 | 0.6579 | 0.6651 | 0.6571 | 0.6553 | 0.6603 | 0.6699 | 0.6615 |
| fr | 0.6715 | 0.6519 | 0.6653 | 0.6591 | 0.6377 | 0.6355 | 0.6407 | 0.6353 | 0.6695 | 0.6433 | 0.6615 | 0.6567 | 0.6752 | 0.6649 | 0.6735 | 0.6691 |
| hi | 0.4611 | 0.5790 | 0.6112 | 0.5890 | 0.5230 | 0.5802 | 0.6098 | 0.5747 | 0.6068 | 0.6036 | 0.6246 | 0.6174 | 0.3747 | 0.5970 | 0.6086 | 0.5962 |
| ru | 0.6382 | 0.6181 | 0.6279 | 0.6094 | 0.6299 | 0.6146 | 0.6335 | 0.6228 | 0.6487 | 0.6531 | 0.6487 | 0.6455 | 0.3770 | 0.6279 | 0.6319 | 0.6407 |
| sw | 0.5780 | 0.5894 | 0.6020 | 0.5727 | 0.5958 | 0.6100 | 0.5812 | 0.5978 | 0.5743 | 0.5749 | 0.5816 | 0.5868 | 0.5848 | 0.5848 | 0.5972 | 0.5890 |
| th | 0.4032 | 0.4774 | 0.5148 | 0.5116 | 0.4323 | 0.5176 | 0.4752 | 0.4539 | 0.3826 | 0.5208 | 0.5118 | 0.5112 | 0.6210 | 0.5429 | 0.5609 | 0.5661 |
| tr | 0.6299 | 0.6224 | 0.6397 | 0.6399 | 0.6343 | 0.6228 | 0.6176 | 0.6102 | 0.6142 | 0.6301 | 0.6437 | 0.6218 | 0.6337 | 0.6457 | 0.6295 | 0.6102 |
| ur | 0.3970 | 0.5387 | 0.5615 | 0.5621 | 0.5271 | 0.5439 | 0.5609 | 0.5591 | 0.5717 | 0.5766 | 0.5745 | 0.5218 | 0.3501 | 0.5545 | 0.5701 | 0.5661 |
| vi | 0.6257 | 0.6182 | 0.6375 | 0.6108 | 0.6429 | 0.6257 | 0.6405 | 0.5836 | 0.6132 | 0.6321 | 0.6371 | 0.6096 | 0.6263 | 0.6403 | 0.6431 | 0.6385 |

Table 15: Accuracy of each model on sentence classification (i.e., XNLI) task.

| | sim-same | | | | dissim-same | | | | sim-div | | | | dissim-div | | | |
|----|----------|--------|--------|--------|-------------|--------|--------|--------|---------|--------|--------|--------|------------|--------|--------|--------|
| | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher | Ortho | IPA | Rom | Cipher |
| ca | 0.9860 | 0.9680 | 0.9848 | 0.9830 | 0.9770 | 0.9514 | 0.9770 | 0.9743 | 0.9811 | 0.9575 | 0.9798 | 0.9763 | 0.9815 | 0.9591 | 0.9804 | 0.9770 |
| de | 0.9480 | 0.9258 | 0.9475 | 0.9419 | 0.9465 | 0.9314 | 0.9475 | 0.9454 | 0.9445 | 0.9242 | 0.9452 | 0.9400 | 0.9453 | 0.9249 | 0.9454 | 0.9431 |
| es | 0.9615 | 0.9569 | 0.9607 | 0.9584 | 0.9504 | 0.9390 | 0.9503 | 0.9468 | 0.9535 | 0.9464 | 0.9532 | 0.9506 | 0.9564 | 0.9479 | 0.9563 | 0.9519 |
| fi | 0.9370 | 0.9194 | 0.9391 | 0.9335 | 0.9336 | 0.9218 | 0.9323 | 0.9302 | 0.9349 | 0.9215 | 0.9338 | 0.9320 | 0.9373 | 0.9260 | 0.9357 | 0.9375 |
| fr | 0.9768 | 0.9636 | 0.9763 | 0.9747 | 0.9685 | 0.9516 | 0.9644 | 0.9629 | 0.9735 | 0.9646 | 0.9739 | 0.9715 | 0.9757 | 0.9684 | 0.9738 | 0.9736 |
| hi | 0.8047 | 0.9432 | 0.9464 | 0.9448 | 0.7682 | 0.9376 | 0.9435 | 0.9365 | 0.9592 | 0.9593 | 0.9617 | 0.9585 | 0.5317 | 0.9476 | 0.9533 | 0.9464 |
| hr | 0.9597 | 0.9476 | 0.9586 | 0.9528 | 0.9538 | 0.9362 | 0.9515 | 0.9427 | 0.9648 | 0.9594 | 0.9641 | 0.9606 | 0.9564 | 0.9488 | 0.9576 | 0.9542 |
| ka | 0.6951 | 0.7912 | 0.8039 | 0.7916 | 0.5980 | 0.7658 | 0.7785 | 0.7687 | 0.4667 | 0.8021 | 0.8066 | 0.8008 | 0.8662 | 0.8567 | 0.8590 | 0.8568 |
| ko | 0.4201 | 0.8551 | 0.8836 | 0.8687 | 0.4122 | 0.8400 | 0.8567 | 0.8601 | 0.4128 | 0.8665 | 0.8742 | 0.8712 | 0.8246 | 0.9047 | 0.9085 | 0.9158 |
| lv | 0.9559 | 0.9411 | 0.9472 | 0.9457 | 0.9537 | 0.9423 | 0.9454 | 0.9441 | 0.9507 | 0.9430 | 0.9451 | 0.9417 | 0.9500 | 0.9442 | 0.9476 | 0.9432 |
| pt | 0.9719 | 0.9515 | 0.9700 | 0.9679 | 0.9584 | 0.9234 | 0.9571 | 0.9528 | 0.9629 | 0.9321 | 0.9595 | 0.9561 | 0.9641 | 0.9349 | 0.9617 | 0.9579 |
| ro | 0.9664 | 0.9607 | 0.9647 | 0.9619 | 0.9483 | 0.9447 | 0.9489 | 0.9434 | 0.9526 | 0.9540 | 0.9556 | 0.9477 | 0.9526 | 0.9539 | 0.9541 | 0.9509 |
| ru | 0.9359 | 0.9126 | 0.9393 | 0.9314 | 0.9302 | 0.8988 | 0.9319 | 0.9147 | 0.9530 | 0.9415 | 0.9512 | 0.9497 | 0.6027 | 0.9169 | 0.9365 | 0.9309 |
| sq | 0.8468 | 0.8317 | 0.8456 | 0.8322 | 0.6723 | 0.6809 | 0.6410 | 0.6229 | 0.7227 | 0.7288 | 0.7334 | 0.6989 | 0.6938 | 0.7291 | 0.6980 | 0.6534 |
| sv | 0.9498 | 0.9379 | 0.9474 | 0.9457 | 0.9177 | 0.8926 | 0.9178 | 0.9050 | 0.9271 | 0.9135 | 0.9302 | 0.9200 | 0.9264 | 0.9138 | 0.9281 | 0.9219 |
| te | 0.6428 | 0.8697 | 0.8712 | 0.8874 | 0.6346 | 0.8530 | 0.8773 | 0.8596 | 0.6273 | 0.8673 | 0.8778 | 0.8649 | 0.9032 | 0.9180 | 0.9011 | 0.8885 |
| th | 0.4540 | 0.7830 | 0.7885 | 0.7817 | 0.4407 | 0.7676 | 0.7848 | 0.7563 | 0.2018 | 0.7927 | 0.7938 | 0.7810 | 0.8398 | 0.8146 | 0.8214 | 0.8134 |
| ur | 0.8434 | 0.8732 | 0.8795 | 0.8776 | 0.8492 | 0.8487 | 0.8725 | 0.8689 | 0.9001 | 0.8984 | 0.9044 | 0.8961 | 0.4645 | 0.8741 | 0.8879 | 0.8797 |
| uz | 0.7095 | 0.6894 | 0.7124 | 0.6985 | 0.7463 | 0.6737 | 0.7547 | 0.7455 | 0.6954 | 0.7172 | 0.6939 | 0.6924 | 0.6860 | 0.6987 | 0.7054 | 0.6791 |

Table 16: F1 scores of each model on the POS task under different language-set conditions.