



Cross-Lingual Transfer Learning for Phrase Break Prediction with Multilingual Language Model

Hoyeon Lee, Hyun-Wook Yoon, Jong-Hwan Kim, Jae-Min Kim

NAVER Cloud, South Korea

yeon.lee@navercorp.com

Abstract

Phrase break prediction is a crucial task for improving the prosody naturalness of a text-to-speech (TTS) system. However, most proposed phrase break prediction models are monolingual, trained exclusively on a large amount of labeled data. In this paper, we address this issue for low-resource languages with limited labeled data using cross-lingual transfer. We investigate the effectiveness of zero-shot and few-shot cross-lingual transfer for phrase break prediction using a pre-trained multilingual language model. We use manually collected datasets in four Indo-European languages: one high-resource language and three with limited resources. Our findings demonstrate that cross-lingual transfer learning can be a particularly effective approach, especially in the few-shot setting, for improving performance in low-resource languages. This suggests that cross-lingual transfer can be inexpensive and effective for developing TTS front-end in resource-poor languages.

Index Terms: cross-lingual transfer, multilingual, phrase break prediction, text-to-speech front-end, language model

1. Introduction

The text processing front-end, such as text normalization, grapheme-to-phoneme (G2P), and phrase break prediction, has become a core part of modern text-to-speech (TTS) systems. Many studies have shown that these text processing modules successfully leverage the naturalness of synthetic speech in various approaches, including traditional statistical methods [1–3] and deep learning-based methods [4–9]. Recently, with the great success of BERT [10] in various natural language processing (NLP) tasks, most of the proposed works have adopted mainstream pre-trained language models (PLMs) [10–13]. These works have reported stable and robust performance on several downstream tasks of the TTS front-end, such as text normalization [4], G2P conversion [5], and phrase break prediction [6–9]. While PLMs have enabled significant advances in a wide range of TTS front-end tasks, a sizeable set of labeled, task-specific data is essentially required.

However, a well-structured, large-scale database that thereby affects the quality of models carries a tremendous cost in linguistic annotation, as well as notorious tediousness, and subjectivity [14, 15]. Even building it by language or task is more challenging in a practical way. Current research on the TTS front-end is actively conducted on manually constructed, large-scale monolingual datasets or a handful of resource-rich languages, such as English, to address the data scarcity issue. One potential approach to address this challenge is to leverage relatively low-resources via *cross-lingual transfer* from a high-resource language. This approach can be effective in reducing the cost and time of data annotation for low-resource languages,

while improving their performance [12, 16–22].

In this paper, we empirically investigate the effectiveness of *zero-shot* and *few-shot cross-lingual transfer* for phrase break prediction in four Indo-European languages: English, French, Spanish, and Portuguese. We explore this question using DistilBERT [11, 12], which is one of the mainstream pre-trained multilingual language models to utilize multilingual representation space. Our manually annotated dataset consists of one high-resource language and three other relatively low-resource languages, which are approximately 12%-16% in size compared to the English dataset. We consider two different transfer learning settings and compare the performance of 1) *zero-shot, cross-lingual* – fine-tuning on one source language and testing on different target languages, and 2) *few-shot, cross-lingual* – fine-tuning on one source language and a few target language instances, and testing on the target language. Our contributions are summarized as follows:

- This is the first study to investigate the effectiveness of using a cross-lingual transfer learning framework based on a pre-trained multilingual language model for phrase break prediction in low-resource languages.
- To verify its efficacy in transfer settings, we conducted cross-lingual transfer experiments in both zero-shot and few-shot settings. The experimental results showed that zero-shot, cross-lingual models are not enough to resolve the lack of large-scale labeled data in their current state but that the few-shot models' performance was comparable to or slightly better than that of the monolingual models.
- We discovered that DistilBERT, a pre-trained multilingual language model, can quickly leverage multilingual representation space, adapting its knowledge from high-resource languages to low-resource ones.

2. Related Work

2.1. Phrase Break Prediction

A phrase break prediction refers to a pause or boundary between phrases, mainly for intonational or breath reasons, and plays an important role in comprehending the structure of sentences. In TTS systems, phrase break prediction can be one of the most important front-end modules [23] along with text normalization [4] and G2P [5]. Although previous studies have indicated that the use of a deep learning-based model architecture [6–9] in phrase break prediction results in better performance than traditional methods such as a hidden Markov model (HMM) based model [2], most models are trained on a large-scale monolingual dataset. Futamata et al. [7] employed contextual text representations obtained from the BERT model pre-trained on Japanese Wikipedia, along with various linguistic features such as part-

Table 1: *Distribution of the dataset. The number of utterances and their ratio to English per language, the average number of subwords per utterance, the average subword length, and the average phrasing information ratio in an utterance. The standard deviation is shown in parentheses.*

| Language | Utterances | Ratio | Subwords | Subword length | AP | IP | SB |
|------------|------------|-------|---------------|----------------|------|------|-----|
| English | 69,169 | 1.00 | 17.89 (11.42) | 3.47 (2.17) | 84.2 | 7.7 | 8.1 |
| French | 8,322 | 0.12 | 18.17 (4.99) | 3.21 (2.20) | 73.6 | 17.6 | 8.7 |
| Spanish | 11,101 | 0.16 | 21.47 (11.66) | 3.35 (2.19) | 86.4 | 6.4 | 7.2 |
| Portuguese | 10,000 | 0.15 | 17.46 (4.22) | 2.94 (1.78) | 88.2 | 3.4 | 8.4 |

of-speech (POS) and dependency tree (DEP), and showed significant improvement compared to the conventional methods. However, these approaches for building phrase break prediction models typically demand a substantial amount of labeled data, which can be a significant challenge for low-resource languages.

2.2. Cross-lingual Transfer

Most supervised learning-based models heavily rely on a large amount of labeled data to achieve optimal performance, which can be challenging due to the data scarcity issue. To resolve this issue, cross-lingual transfer from high-resource to low-resource languages can be employed [21, 24]. A cross-lingual transfer is a strategy by which a model can solve diverse downstream tasks from a resource-poor target language despite receiving only a few examples or none at all in that language during the training process. Typically, a model is pre-trained using a large corpus of monolingual data in different languages, including more than 100 diverse languages, and can also make use of multilingual corpora. Then, the model is fine-tuned with a specific target language [25]. In recent studies, cross-lingual transfer in NLP comes with techniques that rely on continuous cross-lingual representation spaces for different languages [26] or multilingual transformer models that are pre-trained on large, multilingual corpora through the language modeling objectives [12, 16]. Due to the advantages of using a shared representation, multilingual language model-based zero-shot [19, 20] or few-shot [15] cross-lingual transfer studies are making progress.

3. Dataset Description

To assess the cross-lingual transfer learning framework for phrase break prediction, we collected utterances from various domains in four languages: English (EN), French (FR), Spanish (ES), and Portuguese (PT). These languages share typologically similar features: they belong to the Indo-European language family [27], use the Latin alphabet [28], and have the same subject-verb-object (SVO)¹ word order [27, 29]. To prevent bias towards specific domains in our datasets, we collected utterances from diverse sources, such as news articles, blog posts, community websites, and the spoken language corpus. We then manually curated utterances containing sentences that satisfy the following criteria: 1) complete sentence structure, and 2) a minimum of 4 words and a maximum of 25 words per sentence.

Each utterance consists of several phrasing information, which mainly comprises an intonation phrase (IP), accent

¹There are seven categories of word order found among the various human languages: subject-verb-object (SVO), subject-object-verb (SOV), object-verb-subject (OVS), object-subject-verb (OSV), verb-object-subject (VOS), and verb-subject-object (VSO) [27].

phrase (AP), and end of the sentence (SB). IP refers phonetic pause inserted between phrases, and AP refers to a delimitation. Following the annotation method of Futamata et al. [7], we constructed a phrase break dataset in which phrasing information was manually annotated by a linguistic expert. Each silence between words and word transitions in the recorded speech of collected utterances was labeled as a phrase break.

Since some text preprocessing techniques for a particular language may degrade performance in other languages [30, 31], we applied only minimal preprocessing. For tokenization, we used a shared WordPiece [32] vocabulary instead of language-specific tokenization. Most of the utterances in our dataset around 69,000 are in English, whereas the other three languages have much fewer scripts, roughly 12%-16% of the English dataset, primarily due to the higher cost of the annotation required. Each utterance consisted of an average of 18.27 ($SD = 10.58$) subwords, and the average length of a subword was 3.38 ($SD = 2.14$). In total, our dataset contained 98,592 utterances with 1.8 million subwords in four languages. Table 1 shows the distribution of our dataset by language.

4. Zero-Shot Transfer

First, we focus on zero-shot, cross-lingual transfer for phrase break prediction in four languages: training on one language and testing on other unseen languages.

4.1. Experimental Setup

From our manually-curated dataset, we randomly select 6,000, 200, and 200 utterances as training, validation, and test sets for each language, respectively. We use the shared multilingual WordPiece vocabulary for utterances across all languages and take sequences of subwords as model input.

We follow the same architecture as DistilmBERT [11], which is a distilled version of multilingual BERT (mBERT) [12] trained on monolingual corpora in 104 languages, including English, French, Spanish, and Portuguese. In detail, the model consisted of 6 transformer layers with 12 heads in each layer, 768 hidden dimensions, making up 134M parameters². DistilmBERT is fine-tuned with a maximum sequence length of 128 subword tokens to accommodate relatively long utterances for 5 epochs³. To select the best hyperparameters, we use the validation set performance. We grid-search for the following hyperparameter values and select those in bold as our best: a batch

²Note that DistilmBERT has relatively smaller parameters compared to the 177M parameters of mBERT-base and is about twice as fast as mBERT-base [11]. In our preliminary experiments, we confirmed that the distilled model was enough to effectively train for phrase break prediction, our downstream task, in various languages.

³DistilmBERT was trained on a single NVIDIA Tesla V100 GPU for every cross-lingual transfer experiment including subsequent few-shot settings.

Table 2: Macro-F1 results of zero-shot cross-lingual transfer experiments. Monolingual results are underlined, and the best results are **bolded**.

| Train \ Test | EN | FR | ES | PT |
|--------------|---------------------|---------------------|---------------------|---------------------|
| EN | <u>82.12</u> | 76.98 | 72.16 | 86.39 |
| FR | 73.71 | <u>92.37</u> | 67.37 | 76.84 |
| ES | 71.63 | 66.18 | <u>91.40</u> | 70.38 |
| PT | 78.78 | 76.72 | 68.37 | <u>92.62</u> |

size from {16, 32}, a learning rate of Adam optimizer [33] from {2e-5, **5e-5**, 5e-6}, and a dropout rate of the last layer from {0.1, 0.2}. The loss function is cross-entropy loss, and we measure the model performance using a macro-averaged F1 score.

4.2. Results and Analysis

Table 2 presents the macro-averaged F1 scores for different training and test languages, with rows and columns representing each language, respectively. The numbers underlined refer to the results in the monolingual setting. As expected, the performance drops in zero-shot results for all target languages compared to the monolingual results, with the degree of the decline varying significantly across languages. For example, when training on Spanish and testing on French, we observe the largest performance drop of 26.19%. On the other hand, when training on Portuguese and testing on English, it is only 3.34%, which is the smallest drop, and very close to the result of monolingual models. Notably, given the monolingual results (underlined numbers in Table 2), training and testing in English produces the poorest performance compared to other languages, even though English has the highest proportion of pre-training corpora used in multilingual model training. These results imply that zero-shot cross-lingual transfer, even without using a single target language instance, can mitigate the data scarcity issue to some extent. However, it is also clear that it cannot achieve the same level of performance as monolingual models, which could be a feasible alternative for real-world TTS front-end applications. Therefore to further explore beyond the limitation of a zero-shot setting, we proceed to investigate cross-lingual transfer in few-shot settings in the following section.

5. Few-Shot Transfer

To improve the inadequate results in zero-shot transfer, we now investigate the effectiveness of few-shot, cross-lingual transfer with a modest number of target language instances in conjunction with the source language data. Our objective is to narrow the performance gap with the monolingual model.

5.1. Experimental Setup

Unlike in zero-shot transfer, we employ a two-step fine-tuning approach in this setting. For the first fine-tuning step, we consider two different scenarios – *English-Only* and *Augmented* – to analyze each few-shot, cross-lingual transfer performance. In both scenarios, we expand a source language dataset in English to 60,000 examples to take advantage of the abundance of available data. For the first scenario (*English-Only*), we use one of the general cross-lingual transfer methods, which fine-tunes the model solely on a large amount of English data. Additionally, we consider another scenario (*Augmented*) based on the generally accepted correlation between larger dataset size and better




Figure 1: Results of few-shot, cross-lingual transfer experiments with varying numbers of target language examples k . The solid lines indicate the *English-Only* scenario, and the dashed lines indicate the *Augmented* scenario.

performance [34]. In this scenario, English and all other languages except for the target language are included during the training. Note that the datasets used for training, validation, and testing are identical to those used in the zero-shot experiment. This approach reflects a realistic scenario that utilizes all available datasets.

After the first fine-tuning step, the second fine-tuning step is conducted in the same manner for both scenarios. During this step, we progressively increase the size of the target language examples k in powers of 2 as follows: $k \in \{4, 8, 16, \dots, 4096\}$. Note that the model architecture and hyperparameters used in this experiment are the same as those in the zero-shot setting.

5.2. Results and Analysis

The overall results of few-shot, cross-lingual transfer in both scenarios across all languages are presented in Figure 1. The results depicted in Figure 1 indicate consistent and significant improvements in performance across all languages and scenarios, as the number of target language instances represented by k increases. Even with very few instances, less than 100 in total, substantial performance gains are observed in comparison to zero-shot setting in all cases. We can also observe that when trained on only a few target examples (i.e., from 4 to 256), the score rapidly increased during the initial stages, then gradually slowed down as more examples were added, and eventually converged. Interestingly, in some cases, the performance even declined with a larger number of examples, particularly when k was greater than 2,000. Considering the results in Figure 1, we conjecture that DistilMBERT, a pre-trained multilingual language model, can leverage multilingual representation space effectively and transfer its knowledge of the phrase break prediction task to a new unseen language.

5.2.1. Results in English-Only Scenario

Table 3 shows a detailed breakdown of the few-shot transfer results for the *English-Only* scenario. To highlight the performance changes with respect to k , we included the zero-shot performance from a model exclusively fine-tuned on an English dataset, as well as the few-shot performance, and the differences

Table 3: Results of few-shot, cross-lingual transfer in the English-Only scenario. k denotes the number of target language instances, score denotes a macro-averaged F1 score, and Δ denotes difference with respect to the zero-shot setting.

| Language | zero-shot | $k=8$ | | $k=32$ | | $k=128$ | | $k=256$ | | $k=512$ | | $k=2048$ | |
|----------|-----------|-------|----------|--------|----------|---------|----------|--------------|----------|---------|----------|--------------|----------|
| | score | score | Δ | score | Δ | score | Δ | score | Δ | score | Δ | score | Δ |
| FR | 79.69 | 86.42 | 6.73 | 86.75 | 7.06 | 89.79 | 10.1 | 89.41 | 9.72 | 90.59 | 10.9 | 91.89 | 12.2 |
| ES | 72.01 | 76.37 | 4.36 | 76.96 | 4.95 | 77.45 | 5.44 | 84.93 | 12.92 | 87.24 | 15.23 | 91.79 | 19.78 |
| PT | 85.89 | 88.22 | 2.33 | 87.98 | 2.09 | 90.82 | 4.93 | 92.85 | 6.96 | 92.25 | 6.36 | 90.96 | 5.07 |

Table 4: Results of few-shot, cross-lingual transfer in the Augmented scenario. k denotes the number of target language instances, score denotes a macro-averaged F1 score, and Δ denotes difference with respect to the zero-shot setting.

| Language | zero-shot | $k=8$ | | $k=32$ | | $k=128$ | | $k=256$ | | $k=512$ | | $k=2048$ | |
|----------|-----------|-------|----------|--------|----------|---------|----------|---------|----------|--------------|----------|--------------|----------|
| | score | score | Δ | score | Δ | score | Δ | score | Δ | score | Δ | score | Δ |
| FR | 79.21 | 87.38 | 8.4 | 88.06 | 8.85 | 90.13 | 10.92 | 89.77 | 10.56 | 91.19 | 11.98 | 91.9 | 12.69 |
| ES | 69.57 | 72.5 | 2.93 | 72.73 | 3.16 | 77.73 | 8.16 | 84.79 | 15.22 | 87.41 | 17.84 | 90.89 | 21.32 |
| PT | 81.43 | 87.52 | 6.09 | 88.68 | 7.25 | 88.53 | 7.1 | 91.15 | 9.72 | 91.99 | 10.56 | 90.91 | 9.48 |

between them. As shown in Table 3, the zero-shot transfer performance is enhanced in French and Portuguese and remains nearly the same in Spanish since the size of the English training set increases in the first fine-tuning step. This implies that training on a larger dataset can lead to learning more knowledge about the phrase break prediction task. We notice a consistent pattern in which performance gradually improves and achieves comparable results to a monolingual model trained with supervised learning. Despite a marginal difference of 0.48, the French few-shot, cross-lingual model (91.89, $k=2,048$) scored slightly lower than the monolingual model (92.37), whereas the Spanish and Portuguese models outperformed when $k=2,048$ and $k=256$, respectively. Surprisingly, in the case of Portuguese, the model obtained the best F1 score when only 256 training instances were used. Furthermore, the performance in Spanish improved to 91.99 even when k was set to 4,096.

5.2.2. Results in Augmented Scenario

Table 4 presents the results of few-shot transfer for the *Augmented* scenario. As in the previously presented findings, we can see a steady improvement as k increases, following a similar pattern to the *English-Only* scenario. On the other hand, zero-shot performance generally decreased by less than 4%. Interestingly, even with a relatively small amount of high-resource language data (as shown in Table 2), the performance in Spanish and Portuguese was lower than that of zero-shot transfer. This demonstrates that the model is not able to generalize well across different languages during the initial fine-tuning step, as the datasets of different languages and sizes are trained simultaneously. Consequently, this affects the next fine-tuning step, where the target language examples are added to the training set. As you can see in Figure 1, languages with the lower zero-shot performance show poorer performance than the *English-Only* scenario. A notable observation is that, as k reaches a large size (greater than 1,000), comparable performance is reached, and even beyond that point, performance gradually improves.

5.2.3. Analysis

As we have conducted few-shot transfer experiments in two scenarios, we can verify that cross-lingual transfer is possible in a phrase break prediction task and report considerable gains with

only a small amount of labeled data. Since PLMs were introduced, training a large amount of labeled data has become more essential, as it always leads to better performance on various downstream tasks, including TTS front-end modules. However, our findings demonstrate that few-shot transfer can fully address the lack of labeled data and narrow the gap with that, while zero-shot transfer, the most cost-effective approach that does not use any target example, cannot be used as it is. Furthermore, the fact that a distilled version of the multilingual language model was enough to leverage its representation ability will be more beneficial for real-time TTS applications.

6. Conclusion

In this study, we present the first comprehensive research to verify the effectiveness of cross-lingual transfer for phrase break prediction in less-resourced languages using a pre-trained multilingual language model. We investigate zero-shot and few-shot transfer learning settings to find an inexpensive alternative for low-resource language adaptation. While our experimental results in the zero-shot setting did not rise to the performance of the monolingual model trained on a large-scale dataset with consistent labeling, they demonstrated the capability of the multilingual model to perform the transfers required for this task. Moreover, we found that few-shot transfer using only a small number of annotated examples led to performance comparable to that of the monolingual model. Our findings indicate that a few-shot, cross-lingual model can be an effective solution to address the challenge of limited annotated data in TTS front-end applications. In future work, we plan to investigate the efficacy of cross-lingual transfer for other TTS front-end tasks, such as grapheme-to-phoneme (G2P) conversion, in languages with limited resources. Furthermore, in future studies, we will explore whether the transfer learning approach works best even between typologically distinct languages, such as English and Korean.

7. Acknowledgements

This work was supported by Voice&Avatar, NAVER Cloud, Seongnam, Korea.

8. References

- [1] X. Sun and T. H. Applebaum, “Intonational phrase break prediction using decision tree and n-gram model.” in *Proc. INTERSPEECH*, 2001, pp. 537–540.
- [2] H. Schmid and M. Atterer, “New statistical methods for phrase break prediction,” in *COLING 2004: Proceedings of the 20th International Conference on Computational Linguistics*, 2004, pp. 659–665.
- [3] Y. Qian, Z. Wu, X. Ma, and F. Soong, “Automatic prosody prediction and detection with conditional random field (CRF) models,” in *2010 7th International Symposium on Chinese Spoken Language Processing*, 2010, pp. 135–138.
- [4] H. Zhang, R. Sproat, A. H. Ng, F. Stahlberg, X. Peng, K. Gorman, and B. Roark, “Neural models of text normalization for speech applications,” *Computational Linguistics*, vol. 45, no. 2, pp. 293–337, 2019.
- [5] Y. Zhang, L. Deng, and Y. Wang, “Unified mandarin TTS front-end based on distilled BERT model,” in *arXiv preprint arXiv:2012.15404*, 2020.
- [6] R. Liu, B. Sisman, F. Bao, J. Yang, G. Gao, and H. Li, “Exploiting morphological and phonological features to improve prosodic phrasing for mongolian speech synthesis,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 274–285, 2020.
- [7] K. Futamata, B. Park, R. Yamamoto, and K. Tachibana, “Phrase break prediction with bidirectional encoder representations in japanese text-to-speech synthesis,” in *Proc. INTERSPEECH*, 2021, pp. 3126–3130.
- [8] M. Kunešová and M. Řezáčková, “Detection of prosodic boundaries in speech using Wav2Vec 2.0,” in *International Conference on Text, Speech, and Dialogue*. Springer, 2022, pp. 377–388.
- [9] A. Meshnikova and D. Kocharov, “Prosodic boundaries prediction in russian using morphological and syntactic features,” in *Conference on Artificial Intelligence and Natural Language*. Springer, 2019, pp. 126–135.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.
- [11] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter,” in *NeurIPS EMC² Workshop*, 2019.
- [12] T. Pires, E. Schlinger, and D. Garrette, “How multilingual is multilingual BERT?” in *Proc. ACL*, 2019, pp. 4996–5001.
- [13] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu, “TinyBERT: Distilling BERT for natural language understanding,” in *Findings of EMNLP*, 2020, pp. 4163–4174.
- [14] K. Fort, *Collaborative annotation for reliable natural language processing: Technical and sociological aspects*. John Wiley & Sons, 2016.
- [15] A. Lauscher, V. Ravishankar, I. Vulić, and G. Glavaš, “From zero to hero: On the limitations of zero-shot cross-lingual transfer with multilingual transformers,” in *Proc. EMNLP*, 2020, pp. 4483–4499.
- [16] A. Conneau and G. Lample, “Cross-lingual language model pre-training,” *Advances in Neural Information Processing Systems*, vol. 32, pp. 7059–7069, 2019.
- [17] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, “Unsupervised cross-lingual representation learning at scale,” in *Proc. ACL*, 2020, pp. 8440–8451.
- [18] L. Xue, N. Constant, A. Roberts, M. Kale, R. Al-Rfou, A. Sidhant, A. Barua, and C. Raffel, “mT5: A massively multilingual pre-trained text-to-text transformer,” in *Proc. NAACL-HLT*, 2021, pp. 483–498.
- [19] S. Wu and M. Dredze, “Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT,” in *Proc. EMNLP*, 2019.
- [20] M. Artetxe and H. Schwenk, “Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond,” *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 597–610, 2019.
- [21] J. Pfeiffer, I. Vulić, I. Gurevych, and S. Ruder, “MAD-X: An adapter-based framework for multi-task cross-lingual transfer,” in *Proc. EMNLP*, 2020.
- [22] F. Nooralahzadeh, G. Bekoulis, J. Bjerva, and I. Augenstein, “Zero-shot cross-lingual transfer with meta learning,” in *Proc. EMNLP*, 2020, pp. 7654–7673.
- [23] V. Klimkov, A. Nadolski, A. Moinet, B. Putrycz, R. Barra-Chicote, T. Merritt, and T. Drugman, “Phrase break prediction for long-form reading TTS: Exploiting text structure information,” in *Proc. INTERSPEECH*, 2017, pp. 1064–1068.
- [24] S. Schuster, S. Gupta, R. Shah, and M. Lewis, “Cross-lingual transfer learning for multilingual task oriented dialog,” *arXiv preprint arXiv:1810.13327*, 2018.
- [25] G. Lample and A. Conneau, “Cross-lingual language model pre-training,” *arXiv preprint arXiv:1901.07291*, 2019.
- [26] S. Ruder, I. Vulić, and A. Søgaard, “A survey of cross-lingual word embedding models,” *Journal of Artificial Intelligence Research*, vol. 65, pp. 569–631, 2019.
- [27] M. S. Dryer and M. Haspelmath, *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig., 2013.
- [28] G. Grefenstette and J. Nioche, “Estimation of english and non-english language use on the WWW,” in *Proc. RIAO (Recherche d’Informations Assistée par Ordinateur)*, 2000, pp. 237–246.
- [29] B. Comrie, *The world’s major languages*. New York, NY: Oxford University Press, 1987.
- [30] T. Kudo and J. Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” in *Proc. EMNLP*, 2018, pp. 66–71.
- [31] I. Adebara and M. Abdul-Mageed, “Towards afrocetric NLP for african languages: Where we are and where we can go,” in *Proc. 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 3814–3841.
- [32] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey *et al.*, “Google’s neural machine translation system: Bridging the gap between human and machine translation,” *arXiv preprint arXiv:1609.08144*, 2016.
- [33] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [34] P. Delobelle, T. Winters, and B. Berendt, “Robbert: a dutch roberta-based language model,” in *Findings of EMNLP*, 2020, pp. 3255–3265.