

Discovering Representation Sprachbund For Multilingual Pre-Training

Yimin Fan^{1*}, Yaobo Liang², Alexandre Muzio³, Hany Hassan³,
Houqiang Li¹, Ming Zhou⁴ and Nan Duan²

¹University of Science and Technology of China, ²Microsoft Research Asia,

³Microsoft, ⁴Sinovation Ventures

fym0503@mail.ustc.edu.cn, lihq@ustc.edu.cn

{yalia, nanduan, alex.muzio, hanyh}@microsoft.com,

mingzhou926@hotmail.com

Abstract

Multilingual pre-trained models have demonstrated their effectiveness in many multilingual NLP tasks and enabled zero-shot or few-shot transfer from high-resource languages to low-resource ones. However, due to significant typological differences and contradictions between some languages, such models usually perform poorly on many languages and cross-lingual settings, which shows the difficulty of learning a single model to handle massive diverse languages well at the same time. To alleviate this issue, we present a new multilingual pre-training pipeline. We propose to generate language representation from multilingual pre-trained models and conduct linguistic analysis to show that language representation similarity reflect linguistic similarity from multiple perspectives, including language family, geographical sprachbund, lexicostatistics and syntax. Then we cluster all the target languages into multiple groups and name each group as a representation sprachbund. Thus, languages in the same representation sprachbund are supposed to boost each other in both pre-training and fine-tuning as they share rich linguistic similarity. We pre-train one multilingual model for each representation sprachbund. Experiments are conducted on cross-lingual benchmarks and significant improvements are achieved compared to strong baselines.

1 Introduction

The use of pre-trained models is considered a milestone in the development of NLP research. Though early works (Devlin et al., 2019; Radford et al., 2019) on monolingual pre-training (pre-training one model for one language) significantly boosts the performance on the target language, monolingual pre-training can hardly be generalized to multilingual settings because of high training cost and

insufficient corpora resources for many languages.

Multilingual pre-training was proposed to resolve this issue. By using shared vocabulary across languages and pre-training with corpora from multiple languages, multilingual pre-trained models handle cross-lingual tasks in one model. Large scale multilingual pre-trained models provide powerful representation for languages worldwide, enabling significant advances in various multilingual tasks. However, existing widely used multilingual pre-trained models (Lample and Conneau, 2019; Conneau et al., 2020; Huang et al., 2019) perform poorly on many languages and some cross-lingual tasks like zero/few-shot cross-lingual transfer. *e.g.* the performance of zero shot transfer on XNLI task from English data to Urdu language is 15%+ lower than to English (Conneau et al., 2020). Such a huge performance gap is the result of cross-lingual contradictions and differences. This phenomenon is also recognized as negative transfer in transfer learning (Wang et al., 2019).

Many existing works in cross-lingual transfer (K et al., 2020; Pires et al., 2019; Lin et al., 2019) and machine translation (Dabre et al., 2017; Tan et al., 2019) have shown that cross-lingual transfer works best between typologically similar languages. We believe that utilizing similarity between languages is potentially beneficial for large-scale multilingual pre-training. Motivated by this, we propose a new multilingual pre-training pipeline. First, we design a fully data-driven end-to-end way to generate language representation for all languages (108 languages) based on massive multilingual corpora. We represent each language as a 768-dimension vector and use cosine similarity as their similarity measure. With the similarities between languages quantified by their language representation, we automatically divide all languages into a small number of representation sprachbunds. Sprachbund is a linguistic terminology in German that refers to a group of close languages (Sprach means language

* Work is done during internship at Microsoft Research Asia.

and bund means federation in German, so literally it is “language federation” in English). A representation sprachbund is defined as a group of languages with similar language representation. We conduct extensive linguistic analysis and show that languages with similar representation are similar and related from many linguistic perspectives, including language family, geographical sprachbund, lexicostatistics and syntax typology. We believe that training with similar languages in pre-training and fine-tuning are beneficial as they share similar linguistic properties. Second, we train multiple multilingual pre-trained models. Each model is trained with corpora from one representation sprachbund. When handling downstream tasks in one specific language, we fine-tune the model pre-trained with the corresponding representation sprachbund corpora. We conduct experiments on 8 representative cross-lingual tasks from XGLUE (Liang et al., 2020) and XTREME (Hu et al., 2020) including sentence classification, structure prediction, question answering and sentence retrieval. Experiment results show that our model significantly outperforms strong baselines.

Our contributions can be summarized as follows: i) We propose a way to automatically generate language representation from multilingual pre-trained models and massive multilingual corpora. ii) We conduct extensive analysis to show language representation and representation sprachbunds can reflect linguistic language similarity and relatedness from multiple perspectives, therefore they can be considered as new paradigm for clustering similar languages in linguistics. iii) We use representation sprachbunds in multilingual pre-training to alleviate the cross-lingual contradiction and differences, and obtain significant improvements compared with strong baselines.

2 Related Work

Our approach presents a new pipeline of multilingual pre-training. Our representation sprachbund is inspired by linguistic language clustering. Our analysis is closely related to methodology in linguistics.

Multilingual Pre-Training Multilingual pre-training was proposed to pre-train a single model with hundreds of languages. Many works use a large amount of multilingual data (*e.g.*, mC4 (Xue et al., 2020), CCNet (Wenzek et al., 2020)) to pre-train large multilingual models like XLM (Lam-

ple and Conneau, 2019), XLM-R (Conneau et al., 2020), Unicoder (Huang et al., 2019) and mT5 (Xue et al., 2020). Several benchmarks are proposed to evaluate the cross-lingual ability of multilingual pre-trained models, including XGLUE (Liang et al., 2020), XTREME (Hu et al., 2020) and XTREME-R (Ruder et al., 2021)

Language Clustering in Linguistics The linguists propose to classify languages in several ways from different perspectives. There are two main kinds of language clustering: genealogical clustering and typological clustering. In genealogical clustering, languages are clustered into language families (Durbin, 1985; Marcantonio, 2002) by their genetic relatedness. Languages in the same language family have the same ancestral language. In typological clustering, languages are clustered by their typological features, like word order, morphology (Dressler, 1986) and lexicostatistics (Hymes, 1960). Geographical sprachbund (Emeneau, 1980) is also a typological clustering method as it groups languages according to their similar areal features coming from geographical proximity.

Language Clustering in Multilingual NLP Several recent works utilize linguistic knowledge about language clustering in multilingual pre-training and machine translation. Tan et al. (2019) uses language family and language embedding to cluster languages and train machine translation model for each cluster. The language embedding is the language-specific tag added to the input of encoder. Their approach focuses on 23 relatively high resource languages. Fan et al. (2020) clusters languages into several groups according to language family, cultural connection and geographical proximity. They do not obtain any language representation and their language groups are human annotated. Kudugunta et al. (2019) reveals the connection between language SVCCA similarity from NMT models and language family. Their evaluation relies on parallel data. Chung et al. (2020) classifies languages into groups based on their token overlap. Only lexical information of languages is used in their approach. Yu et al. (2021) uses multilingual denoising autoencoder to generate language embeddings and analyze the clusters derived from the embeddings. There are also a few earlier works on generating and analyzing language representation (Tiedemann, 2018; Östling and Tiedemann, 2017).

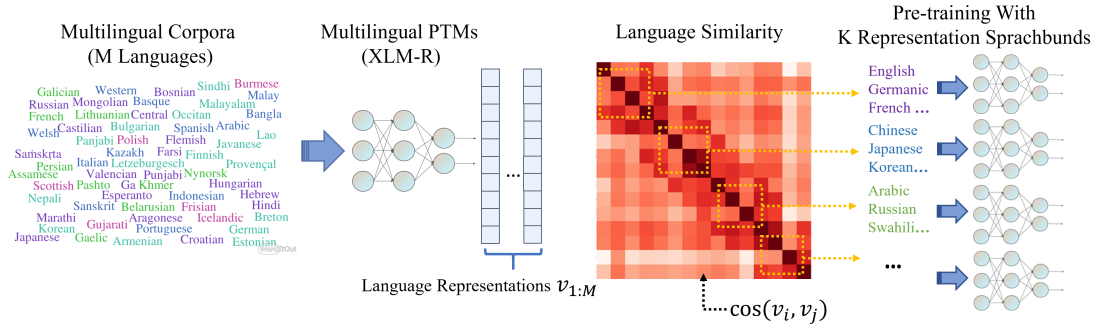


Figure 1: The pipeline of our approach. We first generate language representation for each language with multilingual pre-trained models (XLM-R) and multilingual corpora. We cluster languages into several representation sprachbunds composed of languages with similar representation. We pre-train one model for each representation sprachbund with corpora from that representation sprachbund. Best viewed in color.

Compared to existing works, our approach enjoys the following advantages. First, our approach requires neither parallel data nor linguistic labeling, while most existing works on clustering languages relies on parallel data and linguistic knowledge. Second, our representation sprachbunds contain linguistic features from various aspects, while most existing works focus on language family.

3 Approach

Our proposed pipeline on representation sprachbund for multilingual pre-training is illustrated in Figure 1. Our approach can be divided into two stages: First, we quantify the similarity between languages with generated language representation, and cluster all languages in our corpora into K clusters based on their similarity. Each cluster is called one representation sprachbund. K is a hyper-parameter in our clustering algorithm. Second, we separate our corpora into K parts based on their corresponding representation sprachbund, and pre-train one model with corpora from each representation sprachbund.

3.1 Discovering Representation Sprachbund

Suppose we have M languages in our multilingual corpora C denoted as $L = \{l_1, l_2, \dots, l_M\}$. The corpora of the i th language is denoted as C_i . C_i contains n_i sentences, denoted as $C_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$. Note that there is no need for the sentences $\{s_{ik}\}_{k=1}^{n_i}$ of language l_i and $\{s_{jk}\}_{k=1}^{n_j}$ of l_j to be aligned.

Choenni and Shutova (2020) reveals that the sentence representation from the same language generated by the last layer of multilingual pre-trained models will be very close and form a relatively in-

dependent cluster in the representation space. Motivated by this, a centroid of all sentence representation from the same language can be a reasonable language representation. We employ a transformer-based multilingual pre-trained model, denoted as F . Each language l_i is represented by language representation v_i . We denote the representation of the token k of sentence s from the last layer of F as $F_k(s)$. We then define

$$v_i = \frac{1}{n_i} \sum_{j=1}^{n_i} F_{[CLS]}(s_{ij})$$

We separate all M languages into K clusters via clustering algorithm with input features $\{v_i\}_{i=1}^M$. The cosine similarity between language representation v_i and v_j is used as a similarity metric for the clustering algorithm. The output K clusters are K representation sprachbunds.

3.2 Representation Sprachbund for Multilingual Pre-Training

We denote our K representation sprachbunds as $\{L_1, L_2, \dots, L_K\}$, where $L_k = \{l_{k1}, \dots, l_{kk_i}\}$. We have $\sum_{j=1}^K j_i = M$ as our representation sprachbund is non-overlapping. The corresponding corpora is denoted as $C'_k = \{C'_{k1}, C'_{k2}, \dots, C'_{kk_i}\}$. Note that the corpora of the i th language for discovering representation sprachbund (C_i) and for training models (C'_i) may be different. We train K separate models for K representation sprachbunds. When fine-tuning, we can use the data in language $\{l_{kj}\}_{j=1}^{k_i}$ to fine-tune the k th model if the data is available.

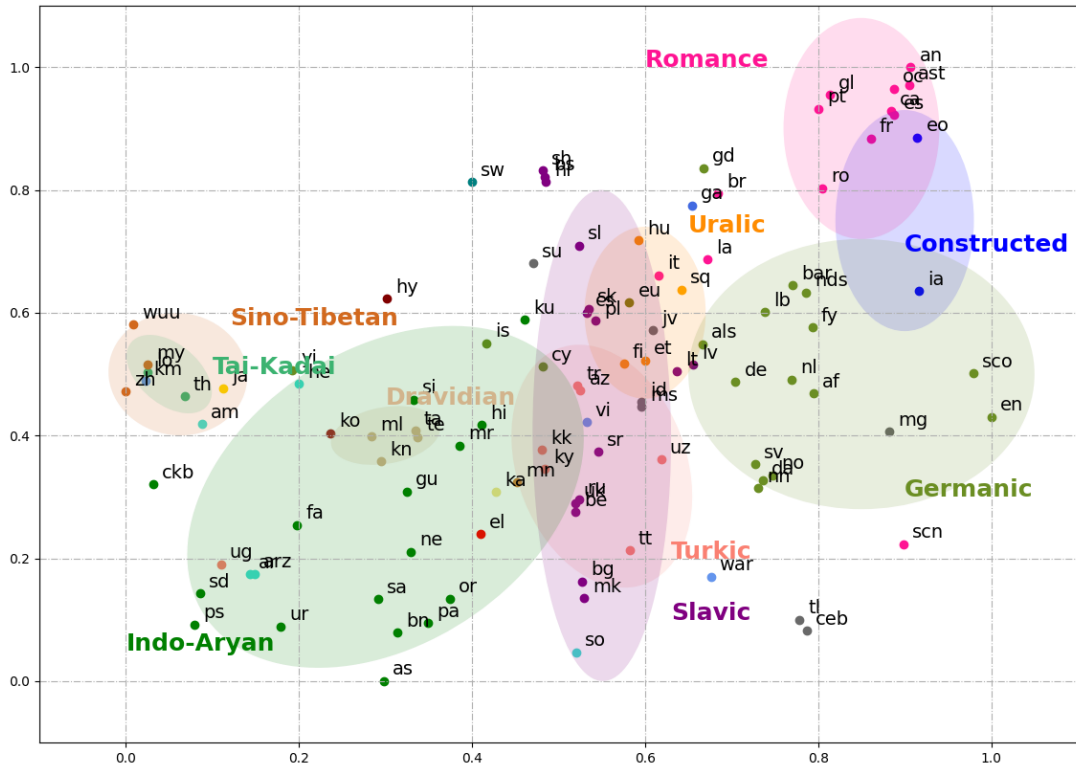


Figure 2: Visualization of language representation (reduced to 2-dimension for visualization). All languages are labeled with ISO 639-1 code. Languages from the same language family are colored the same. We draw ellipse for 10 main language family covering most languages they include. The distribution of language representation has great overlap with the language family and several geographical sprachbunds. Best viewed in color.

4 Representation Sprachbund Discovery and Analysis

4.1 Settings

We collect massive multilingual corpora for discovering representation sprachbund (also for the following multilingual pre-training). We use Wikipedia¹ corpora (100 languages are included, total size 101GB) and a clean version of Common Crawl (CC)² (89 languages are included, total size 2500GB) following Liang et al. (2020). 108 languages are included in our multilingual corpora. Note that we do not use any parallel corpora. The pre-trained model F we use is XLM-R base model implemented by HuggingFace³. As the size of the whole corpora is very large, we use a random sampling strategy to get part of the data for extracting representation. For those languages with less than 10GB of data, we use all the data for extracting representation; for those languages with more than

10GB of data, we sample 10GB out of all the data. We use the method mentioned in Section 3.1 to get the language representation. The dimension of language representation is 768 as in XLM-R base model each token is represented by a 768-dimension vector. We reduce the 768-dimension vectors $v_{1:108}$ to 2-dimension (denoted as $\tilde{v}_{1:108}$) for visualization with the t-SNE algorithm implemented in Scikit-learn Python package⁴ with default parameters. We use min-max normalization to normalize $\tilde{v}_{1:108}$ to range $[0, 1]$. The 2-dimension language representation are visualized in Figure 2.

4.2 Linguistics Analysis

We find that our representation can reflect linguistic similarity and relatedness between languages from different perspectives. We link language representation with several linguistic language similarity measures, also with some linguistic clustering methods. We believe that representation sprachbund is a new paradigm for clustering similar and related languages in linguistics. In Figure 2, each language corresponds to one point (2-dimension

¹https://en.wikipedia.org/wiki/Main_Page

²<https://commoncrawl.org/>

³<https://huggingface.co/xlm-roberta-base>

⁴<https://scikit-learn.org/stable/>

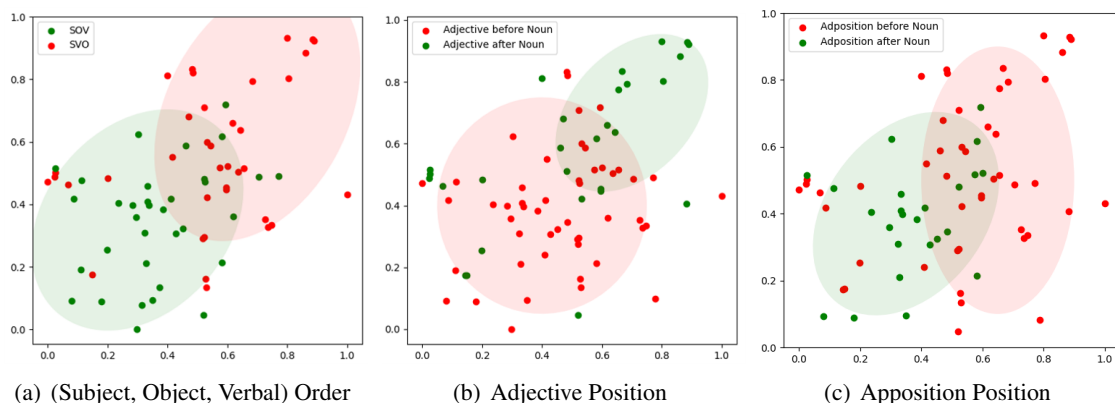


Figure 3: Visualization of the relationship of our language representation and 3 language syntactic features. Languages with the same syntactic features approximately fall in the same region. i.e., have similar language representation. Note that syntactic feature data of several languages is not available. Best viewed in color.

vector) on that figure. We label all points with the ISO 639-1 code⁵ of their corresponding languages.

Relationship with Language Family We find that the distribution of language representation has great overlap and similarity with language family. In Figure 2, the color of each point indicates the language family of its corresponding language. 108 languages are categorized into 22 language families according to Ethnologue⁶. All 108 languages and their corresponding language family can be found in Appendix A. We only label 10 language families with ellipses for clarity. Though there are some special cases, languages within the same language family approximately fall in the same region.

Relationship with Geographical Sprachbund Geographical sprachbund is a group of similar languages from geographical proximity and language contact, while our representation sprachbund is a group of similar languages from representation proximity. We find that our language representation distribution is consistent with many geographical sprachbunds. In Figure 2, on the top-right, Romance and Germanic language representation closeness can be linked with the Western Europe sprachbund from WHORF (1944); on the bottom-left, the closeness between Indo-Aryan, Dravidian and some Sino-Tibetan languages aligns with Indian subcontinent sprachbund proposed in Eme-neau (1956); on the middle, the similar representation from Turkic, Uralic and Mongolic (mn) also

match Altaic sprachbund by BOSWORTH (1962).

Relationship with Lexicostatistics We find that similarity in language representation also reflects lexical similarity between languages. We collect lexical similarity data from Ethnologue on several languages (en, fr, de, pt, ro, ru, es, ca), denoted as Sim_{lex} . We denote the similarity quantified by our language representation as Sim_{data} . We find that Sim_{lex} and Sim_{data} are strongly linear correlated, with Pearson correlation coefficient of 0.83. Lexicostatistics is a method to measure lexical similarity by comparing the percentage of lexical cognates between languages (Hymes, 1960), which is very time-consuming. Our language representation can even further help linguists infer lexical similarity more easily (e.g. linear regression between representation similarity and lexical similarity). The similarity data is shown in Appendix B.

Relationship with Language Syntax Languages have diverse syntactic features defined by linguists and can be classified through these features. We show that the distribution of our language representation implies the syntactic features of corresponding languages. We use the lang2vec Python package (Littell et al., 2017) to query the URIEL database⁷. We choose three representative syntactic features: (subject, object, verbal) word order, adjective position and adposition position. As shown in Figure 3, we find that languages with the same syntactic features approximately have similar language representation.

⁵http://www.infoterm.info/standardization/ISO_639.php

⁶<https://www.ethnologue.com/browse/families>

⁷<http://www.cs.cmu.edu/~dmortens/uriel.html>

Surprise: Help for Exploring Linguistic Mystery Coincidentally, we find that our language representation connect with an existing under-explored linguistic mystery. In Figure 2, Uralic and Austronesian languages (jv, id, ms) have similar language representation. To the best of our knowledge, only a few linguistic works (Ohnishi, 2006, 2009) discussed their similarity and relatedness. The reason for their similarity cannot be explained by language family (genetic relationship) or geographical sprachbund (geographical relationship). Their language representation similarity may be a clue that motivates linguists to find more similarity between them and further explain how their similarity formed.

With the above linguistic analysis, we show that our language representation contain rich linguistic genealogical, geographical, typological, and lexical features of languages, therefore the similarity between language representation can be a good metric for clustering languages. With a 768-dimension vector as numerical feature for each language, we can implement clustering algorithms to cluster similar languages into a representation sprachbund. Our language representation will be released later.

5 Representation Sprachbund For Multilingual Pre-training

5.1 Datasets

We collect massive multilingual corpora for pre-training and use four datasets for downstream task evaluation. The multilingual corpora has been described in Section 4.1. We use XNLI (Conneau et al., 2018), PAWS-X (Yang et al., 2019), NER (Pan et al., 2017), Part of Speech Tagging (POS)(Zeman et al., 2019), MLQA (Lewis et al., 2019), TydiQA (Clark et al., 2020), XQuAD(Artetxe et al., 2020) and cross-lingual sentence retrieval (Artetxe and Schwenk, 2019) as downstream tasks. For cross-lingual sentence retrieval, we collect 21 language pairs and extract 1000 sentence-pairs for each language-pair from tatoeba⁸. This task aims to find the nearest neighbor for each sentence in the other language.

⁸<https://tatoeba.org/eng/downloads>

#i	is the <i>i</i> th representation sprachbund
#1	af als an ast bar br ca ceb da de en eo es el fr fy ga gd gl ia it ku lb nds nl nn no oc pt ro scn sco sq sv tl ur war
#2	ar arz bg bs cy fa hi hr id is mg mk ms ps ru sh sl so sr su sw yi
#3	am as be ckb cs et eu fi he hu ja jv km la lo lt lv mr my ne or pa pl sa sd sk th uk wuu zh
#4	az bn gu hy ka kk kn ko ky ml mn si ta te tt ug uz vi tr

Table 1: Components of 4 representation sprachbunds

5.2 Settings

We use the XLM-R base model as our base model. Fairseq⁹ is used as our pre-training code base. The Huggingface Transformers¹⁰ is used as our fine-tuning code base. We use the hierarchical clustering algorithm implemented by Scikit-learn Python package¹¹ for clustering language representation. We cluster languages into 4 representation sprachbunds. The reason for clustering 4 representation sprachbunds is shown in Section 5.4. Components of representation sprachbunds are shown in Table 1. We use the shared vocabulary of XLM-R base model for reusing the pre-trained parameters and keeping the comparability with the baseline model.

Pre-training Setting We initialize our model with the XLM-R base model parameters and run continual pre-training for 40000 updates on 8 Nvidia V100 GPUs with total batch size 8192. The experiment takes about 8 days. We use Adam optimizer with a linear warm-up and set the learning rate to 3e-5. We pre-train 4 models according to our 4 representation sprachbund corpora. We also randomly create 4 language clusters (each with the same language number as the representation sprachbund). We pre-train 4 models with random language clusters as baseline. To evaluate the impact of continual pre-train corpora, we also continue to pre-train XLM-R base model with our corpora.

Downstream Task Setting For XNLI, we set the learning rate to 5e-6 and train 10 epochs with batch size 32. For POS tagging, we set the learning rate to 2e-5 and train 20 epochs with batch size 32. For MLQA, we set the learning rate to 3e-5, train 2 epochs following BERT for SQuAD with batch size 12. For PAWS-X, NER, TydiQA and XQuAD, we follow the default settings in XTREME (Hu et al., 2020).The downstream task performance on

⁹<https://github.com/pytorch/fairseq>

¹⁰<https://huggingface.co/transformers/>

¹¹<https://scikit-learn.org/stable/>

Tasks	XNLI(Acc)	PAWS-X(Acc)	POS(F1)	NER(F1)	TydiQA(F1)	XQuAD(F1)	MLQA(F1)
XLM-R	75.0	84.2	79.7	60.8	45.9	70.3	65.0
XLM-R CT	74.8	84.8	79.9	60.9	46.3	70.5	65.4
Random	74.8	84.8	79.4	60.9	47.2	70.2	65.5
Ours	75.7	85.4	80.1	63.5	47.8	70.9	66.3

Table 2: Performance of our model fine-tune with English on 7 cross-lingual understanding tasks. XLM-R: directly fine-tune on XLM-R model. XLM-R CT: continue to pre-train XLM-R base model with our corpora. Random: pre-train one model for each random language cluster. Ours: pre-train one model for each representation sprachbund respectively.

Languages	de-en	pt-en	es-en	nl-en	tl-en	ur-en	el-en	af-en	fr-en	it-en	nl-de	de-it	es-pt	de-el	sv-da	da-no	fr-de	it-ro	ar-ru	zh-ja	pl-cs	Avg
XLM-R	89.0	78.7	72.0	77.4	29.9	33.4	54.1	53.1	73.2	65.9	67.0	55.1	75.9	54.5	81.1	89.6	76.0	51.2	43.4	52.6	73.8	64.1
XLM-R CT	85.9	75.5	70.2	77.3	31.4	36.5	52.4	57.9	73.6	63.5	65.4	54.2	73.2	49.9	79.6	89.6	72.3	50.4	46.7	59.6	76.8	63.9
Ours	90.3	80.6	75.3	79.0	32.2	39.8	54.5	57.4	74.6	67.9	68.2	57.8	76.7	55.3	81.0	89.9	77.5	53.3	58.7	61.7	82.1	67.3

Table 3: Performance (Accuracy) of our model on cross-lingual sentence retrieval task without fine-tuning. Languages in different representation sprachbunds are separated with vertical lines.

one specific language is measured by fine-tuning the model pre-trained with the corresponding representation sprachbund corpora. We select the checkpoint with the best performance on the dev set. The results are averaged over three runs. For cross-lingual sentence retrieval, we use the cosine similarity of the average middle layer (the 7th layer of our 12-layer model) embedding for retrieval without fine-tuning.

Evaluation Setting There are three main settings in the fine-tuning stage. (i) **Fine-tune with English.** We fine-tune the model with English labeled data of downstream task. (ii) **Fine-tune with every language.** *e.g.* We fine-tune the model with the French labeled data and test its performance on the French test set. (iii) **Fine-tune with pivot language.** We choose a pivot language l_i in its representation sprachbund m based on similarity with other languages (if $i = \arg \max_i \sum_{l_j \in L_m} \cos(v_i, v_j)$), and fine-tune the model with the pivot language labeled data.

5.3 Main Results

We conduct experiments in several settings on different types of tasks. We find that our approach obtains significant improvements over XLM-R base model and randomly clustered model when applied in pre-training and downstream tasks.

Improving Pre-Training In Table 2, We find that our approach significantly outperforms all baseline models on 7 cross-lingual tasks. The 7 tasks are representative of almost all kinds of cross-lingual understanding tasks, which shows the universal effectiveness of our model. The detailed

results of each task and each language are shown in Appendix C.

Improving Fine-Tuning In Table 4, we also show that without additional costly continual pre-training, directly fine-tuning multiple models with representation sprachbund also improves the performance. We fine-tune each model (XLM-R base) with the labeled pivot language data on XNLI of each representation sprachbund. In Table 4, we find that the improvement is significant (from 75.0 to 76.5). We also show that fine-tuning the model with the labeled data from each language yields significantly better results (77.4). We conclude that fine-tuning with the language similar to target language is likely to boost the performance.

Improving Multilingual Embeddings In Table 3, we show that the performance on cross-lingual sentence retrieval greatly improves with our approach. Though our method is not designed for improving multilingual embeddings, better multilingual embeddings are generated without additional fine-tuning.

Task	XNLI	PAWS-X	POS	NER	TydiQA	XQuAD	MLQA
High	0.2	0.6	0	1.3	1.9	0.5	1.0
Low	1.2	0.6	1.2	3.1	1.4	0.8	1.5

Table 6: Low resource languages has more significant gains with our approach. Low indicates the languages are low resource and isolated.

5.4 Analysis

Gaining More on Low Resource Languages In Table 6, we show that our approach brings more gains to those low resource and isolated languages

Languages	en	fr	es	de	el	ur	bg	ru	ar	hi	sw	tr	vi	th	zh	Avg
XNLI fine-tune with English																
XLM-R	84.8	78.9	79.2	77.7	76.6	67.0	78.6	76.6	72.3	71.1	66.4	73.3	75.6	72.4	74.7	75.0
XLM-R CT	84.5	78.4	79.5	77.2	76.4	66.8	78.3	76.2	73.0	70.6	65.4	73.3	75.3	72.6	74.2	74.8
Random	84.3	78.0	78.5	77.1	76.5	67.3	78.2	76.2	72.2	70.5	66.4	73.5	75.4	72.8	74.7	74.8
Ours	84.1	78.4	80.0	77.8	77.3	68.2	78.4	76.7	74.5	72.1	68.6	74.9	75.3	74.1	74.6	75.7
XNLI fine-tune with pivot language																
XLM-R	83.6	79.5	80.0	79.8	77.4	69.3	80.1	78.4	74.2	73.3	66.9	74.2	78.8	76.1	76.3	76.5
Ours	84.2	80.0	80.4	80.4	77.8	69.9	80.2	78.0	75.5	73.3	70.9	75.0	78.3	76.4	75.7	77.1
XNLI fine-tune with every language																
XLM-R	84.6	79.5	80.9	79.8	79.1	67.1	80.2	78.4	75.5	73.9	71.1	77.2	78.8	76.1	78.4	77.4
Ours	84.1	80.4	81.1	80.4	78.9	66.5	80.3	78.0	77.2	74.3	72.6	76.9	78.3	76.4	78.0	77.6

Table 4: Performance (Accuracy) of our model on XNLI dataset on three settings. Note that we do not use the result in (Conneau et al., 2020), instead we fine-tune the model in our settings. Languages in different representation sprachbunds are separated with vertical lines. Pivot languages are **bold** in the first row.

Languages	en	fr	es	de	el	ur	bg	ru	ar	hi	sw	tr	vi	th	zh	Avg
XLM-R+All	84.8	81.3	82.0	80.5	80.0	71.7	81.6	79.1	78.2	75.6	73.1	78.1	79.4	77.2	79.8	78.8
XLM-R+RSB	85.3	80.9	81.8	80.6	80.0	71.2	81.3	79.4	77.7	75.9	72.6	77.5	79.0	77.4	78.3	78.6
RSB+	85.6	81.4	81.7	80.5	80.9	70.9	82.0	79.5	78.0	76.8	73.8	77.6	78.7	77.3	79.7	79.0

Table 5: Downstream task (XNLI) data efficiency of our model. XLM-R+All: use XNLI data in all languages to fine-tune XLM-R base model. XLM-R+RSB: use XNLI data in the same representation sprachbund to fine-tune XLM-R base model. RSB+: use XNLI data in the same representation sprachbund to fine-tune the pre-trained model. Different representation sprachbunds are separated with vertical lines.

(including ur, ar, sw, tr, vi, th, zh, pl, ja, ko, id, fi, bn, te, tl, af, ms, fa, mr, et, he, jv, eu, yo, my, hu, ta, ml, kk, kn, ka) compared with those high resource languages (including en, de, es, ru, bg, hi, it, fr, nl, pt) which will be beneficial for bridging the large performance gap between high resource and low resource languages. An intuitive explanation is that those low resource and isolated languages suffer from more serious cross-lingual contradictions when trained with those dissimilar high resource languages. When clustered with similar languages, those low resource and isolated languages are likely to benefit a lot.

Achieving Data Efficiency in Downstream Tasks

We show that with the continual pre-training step with representation sprachbund corpora, less data for downstream tasks is needed to achieve high accuracy. In Table 5, on XNLI task, we find that fine-tuning our pre-trained model with all the downstream task data from each representation sprachbund (less than 30% of all the data) achieves better results than fine-tuning XLM-R with all the data (from 78.8 to 79.0). We also find that fine-tuning XLM-R with the representation sprachbund data achieves results comparable with fine-tuning using all the data (78.6 and 78.8), which means that using data from similar languages (though less) works well.

Choosing the Number of Representation Sprachbunds

We conduct experiments to choose the number of representation sprachbunds. We cluster languages into 1,2,4,8 representation sprachbunds, and pre-train 1,2,4,8 models, respectively. We evaluate through fine-tuning with English on the XNLI dataset. As shown in Figure 4, with the increase of the number of representation sprachbunds, the performance also increases. We find that clustering languages into 4 representation sprachbunds is a desirable choice, as from 4 to 8 little gain is obtained but the cost doubles.




Figure 4: Impact of the number of representation sprachbunds on the performance. Performance (Accuracy) is measured by fine-tuning with English labeled data on XNLI task.

6 Conclusion

To reduce the cross-lingual contradictions in pre-training one model for all languages, we propose to merge similar languages into a representation sprachbund and pre-train one model for each representation sprachbund. Results show that our approach outperforms strong baselines in various settings and tasks. We also identify the relationship between our representation sprachbund with linguistic theories. Applications of our representation sprachbund as a paradigm for clustering languages in linguistics will be explored in subsequent work.

Acknowledgements

We thank Ning Wu, Xiaoze Jiang, Yuan Chai, Junhe Zhao, Shunyu Zhang for their useful suggestions on writing this paper.

References

- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. On the Cross-lingual Transferability of Monolingual Representations. In *Proceedings of ACL 2020*.
- Mikel Artetxe and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610.
- C. E. BOSWORTH. 1962. A turco-mongol practice amongst the early ghaznavids? *Central Asiatic Journal*, 7(4):237–240.
- Rochelle Choenni and Ekaterina Shutova. 2020. What does it mean to be language-agnostic? probing multilingual sentence encoders for typological properties.
- Hyung Won Chung, Dan Garrette, Kiat Chuan Tan, and Jason Riesa. 2020. Improving multilingual models with language-clustered vocabularies.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages. In *Transactions of the Association of 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, Guillaume Lample, Ruty Rinott, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. Xnli: Evaluating cross-lingual sentence representations. In *EMNLP*.
- Raj Dabre, Tetsuji Nakagawa, and Hideto Kazawa. 2017. An empirical study of language relatedness for transfer learning in neural machine translation. In *Proceedings of the 31st Pacific Asia Conference on Language, Information and Computation*, pages 282–286. The National University (Phillippines).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- W.U. Dressler. 1986. Morphology: a study of the relation between meaning and form: Joan I. Bybee, (typological studies in language 9.) benjamins, amsterdam; 1985. vii+234pp. *Lingua*, 70(2):197 – 200.
- Marshall Durbin. 1985. A survey of the carib language family. *South American Indian Languages. Retrospect and Prospect*, 325:370.
- M. B. Emeneau. 1956. India as a linguistic area. *Language*, 32(1):3–16.
- M. B. Emeneau. 1980. *Language and linguistic area: essays / by Murray B. Emeneau; selected and introduced by Anwar S. Dil*. Stanford University Press Stanford, Calif.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Edouard Grave, Michael Auli, and Armand Joulin. 2020. Beyond english-centric multilingual machine translation.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization.
- Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2485–2494, Hong Kong, China. Association for Computational Linguistics.

- D. H. Hymes. 1960. [Lexicostatistics so far](#). *Current Anthropology*, 1(1):3–44.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. [Cross-lingual ability of multilingual bert: An empirical study](#). In *International Conference on Learning Representations*.
- Sneha Kudugunta, Ankur Bapna, Isaac Caswell, and Orhan Firat. 2019. [Investigating multilingual NMT representations at scale](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1565–1575, Hong Kong, China. Association for Computational Linguistics.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#). *CoRR*, abs/1901.07291.
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. [Mlqa: Evaluating cross-lingual extractive question answering](#). *ArXiv*, abs/1910.07475.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenge Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastasopoulos, Patrick Littell, and Graham Neubig. 2019. [Choosing transfer languages for cross-lingual learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3125–3135, Florence, Italy. 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.
- Angela Marcantonio. 2002. The uralic language family: facts, myths and statistics.
- Koji Ohnishi. 2006. Word-initial consonant correspondence laws between uralic and austronesian language families: The origin of the uralic language family from a language closely kin to western oceanic and sulawesi language group.
- Koji Ohnishi. 2009. [Sound-correspondence laws of word-initial consonants between proto-indo-european and austronesian languages](#). *Artificial Life and Robotics*, 14(4):567–573.
- Robert Östling and Jörg Tiedemann. 2017. [Continuous multilinguality with language vectors](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 644–649, Valencia, Spain. Association for Computational Linguistics.
- 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 ACL 2017*, pages 1946–1958.
- 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.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Graham Neubig, and Melvin Johnson. 2021. [Xtreme-r: Towards more challenging and nuanced multilingual evaluation](#).
- Xu Tan, Jiale Chen, Di He, Yingce Xia, Tao Qin, and Tie-Yan Liu. 2019. [Multilingual neural machine translation with language clustering](#).
- Jörg Tiedemann. 2018. [Emerging language spaces learned from massively multilingual corpora](#).
- Zirui Wang, Zihang Dai, Barnabas Póczos, and Jaime Carbonell. 2019. [Characterizing and avoiding negative transfer](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Édouard Grave. 2020. [Ccnets: Extracting high quality monolingual datasets from web crawl data](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 4003–4012.
- BENJAMIN LEE WHORF. 1944. [The relation of habitual thought and behavior to language](#). *ETC: A Review of General Semantics*, 1(4):197–215.

- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. [mt5: A massively multilingual pre-trained text-to-text transformer](#).
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. PAWS-X: A cross-lingual adversarial dataset for paraphrase identification. In *Proceedings of EMNLP 2019*, pages 3685–3690.
- Dian Yu, Taiqi He, and Kenji Sagae. 2021. [Language embeddings for typology and cross-lingual transfer learning](#). 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 7210–7225, Online. Association for Computational Linguistics.
- Daniel Zeman, Joakim Nivre, Mitchell Abrams, Noëmi Aeppli, Željko Agić, Lars Ahrenberg, Gabrielè Aleksandravičiūtė, Lene Antonsen, Katya Aplonova, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, Colin Batchelor, John Bauer, Sandra Bellato, Kepa Bengoetxea, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Agnè Bielinskienė, Rogier Blokland, Victoria Bobicev, Loïc Boizou, Emanuel Borges Völker, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Kristina Brokaić, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Tatiana Cavalcanti, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čepľo, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Alessandra T. Cignarella, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Elvis de Souza, Arantza Diaz de Ilarraza, Carly Dickerson, Bamba Dione, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Drohanova, Puneet Dwivedi, Hanne Eckhoff, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Olga Erina, Tomáš Erjavec, Aline Etienne, Wograinne Evelyn, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Kazunori Fujita, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Bernadeta Griciūtė, Matias Groni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Mika Hämäläinen, Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Johannes Heinecke, Felix Hennig, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Takumi Ikeda, Radu Ion, Elena Irimia, Olájidé Ishola, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Markus Juutinen, Hüner Kaşıkara, Andre Kaasen, Nadezhda Kabaeva, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Elena Klementieva, Arne Köhn, Kamil Kopacewicz, Natalia Kotsyba, Jolanta Kovalevskaitė, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Lucia Lam, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phng Lee Hong, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Maria Liovina, Yuan Li, Nikola Ljubei, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Mackentanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Măranduc, David Mareek, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Sarah McGuinness, Gustavo Mendonça, Niko Miekka, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Maria Mitrofan, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Tomohiko Morioka, Shinsuke Mori, Shigeki Moro, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Robert Munro, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horriacek, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lng Nguyễn Thị, Huyền Nguyễn Thị Minh, Yoshihiro Nikaido, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Atul Kr. Ojha, Adédayo Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Angelika Peljak-Lapińska, Siyao Peng, Cenel-Augusto Perez, Guy Perrier, Daria Petrova, Slav Petrov, Jason Phelan, Jussi Piitulainen, Tommi A Pirinen, Emily Pitler, Barbara Plank, Thierry Poibeau, Larisa Ponomareva, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Peng Qi, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Ivan Riabov, Michael Rießler, Erika Rimkutė, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roşca, Olga Rudina, Jack Rueter, Shoval Sadde, Benoît Sagot, Shadi Saleh, Alessio Salomoni, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Dage Särg, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamel Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Hiroyuki Shirasu, Muh Shohibussirri, Dmitry Sichinava, Aline Silveira, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Shingo Suzuki, Zsolt Szántó, Dima Taji, Yuta Takahashi, Fabio Tamburini, Takaaki

Tanaka, Isabelle Tellier, Guillaume Thomas, Lisi Torga, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Andrius Utkas, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Lars Wallin, Abigail Walsh, Jing Xian Wang, Jonathan North Washington, Maximilian Wendt, Seyi Williams, Mats Wirén, Christian Wittern, Tsegay Woldemariam, Tak-sum Wong, Alina Wróblewska, Mary Yako, Naoki Yamazaki, Chunxiao Yan, Koichi Yasuoka, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Manying Zhang, and Hanzhi Zhu. 2019. [Universal dependencies 2.5](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

A Linguistic Language Family of All Languages

Language Family	Languages
Germantic	af,als,bar,cy,da,de,en,fy,gd,is,lb,nds,nl,nn,no,sco,sv,yi
Greek	el
Japonic	ja
Sino-Tibetan	my,wuu,zh
Turkic	az,kk,ky,tr,tt,ug,uz
Uralic	et,fi,hu
Austroasiatic	km,vi,war
Dravidian	kn,ml,ta,te
Slavic	be,bg,bs,cs,hr,lt,lv,mk,pl,ru,sh,sk,sl,sr,uk
Kartvelian	ka
Niger-Congo	sw
Austronesian	ceb,id,jv,mg,ms,su,tl
Armenian	hy
Koreanic	ko
Albanian	sq
Tai-Kadai	lo,th
Romance	an,ast,br,ca,es,fr,gl,it,la,oc,pt,ro,scn,eu
Constructed	eo,ia
Afro-Asiatic	am,ar,arz,he,so
Celtic	ga
Indo-Aryan	as,bn,ckb,fa,gu,hi,ku,mr,ne,or,pa,ps
Mongolic	sa,sd,si,ur

Table 7: Languages and their corresponding language family

B Lexical similarity and Embedding similarity

Languages	ca	en	fr	de	pt	ro	ru	es
ca	1.00	-	0.85	-	0.85	0.73	-	0.85
en	-	1.00	0.27	0.60	-	-	0.24	-
fr	0.85	0.27	1.00	0.28	0.75	0.75	-	0.75
de	-	0.60	0.28	1.00	-	-	-	-
pt	0.85	-	0.75	-	1.00	0.72	-	0.88
ro	0.73	-	0.75	-	0.72	1.00	0.72	0.71
ru	-	0.24	-	-	-	-	1.00	-
es	0.85	-	0.75	-	0.88	0.71	-	1.00

Table 8: Lexical similarity from Ethnologue. Some data is missing.

Languages	ca	en	fr	de	pt	ro	ru	es
ca	1.00	0.08	0.76	0.22	0.63	0.56	0.23	0.81
en	0.08	1.00	0.26	0.38	0.28	0.17	0.31	0.00
fr	0.76	0.26	1.00	0.49	0.63	0.65	0.47	0.68
de	0.22	0.38	0.49	1.00	0.47	0.49	0.59	0.26
pt	0.63	0.28	0.63	0.47	1.00	0.61	0.45	0.64
ro	0.56	0.17	0.65	0.49	0.61	1.00	0.48	0.56
ru	0.23	0.31	0.47	0.59	0.45	0.48	1.00	0.24
es	0.81	0.00	0.68	0.26	0.64	0.56	0.24	1.00

Table 9: Language embedding similarity

C Detailed Results of Cross-lingual Tasks

Languages	ar	hi	de	en	es	zh	vi	Avg
XLM-R	55.3	61.3	62.1	80.0	68.1	61.5	66.9	65.0
XLM-R CT	57.1	61.7	61.9	80.1	68.0	61.5	67.3	65.4
Random	57.2	62.3	62.0	81.0	67.6	61.6	67.2	65.5
Ours	59.2	63.1	62.9	80.4	69.3	61.4	67.6	66.3

Table 10: Performance (F1 score) on MLQA dataset. We examine the results on MLQA on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.

Languages	de	en	es	fr	ja	zh	ko	Avg
XLM-R	87.5	94.4	88.5	88.5	75.9	80.1	74.7	84.2
XLM-R CT	87.7	94.4	88.8	88.5	77.2	81.1	75.6	84.8
Random	87.0	94.4	88.8	88.9	77.4	80.0	76.7	84.8
Ours	88.4	94.9	89.1	89.6	78.2	80.1	77.2	85.4

Table 11: Performance (F1 score) on PAWS-X dataset. We examine the results on PAWS-X on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.

Languages	ar	ru	bg	hi	ur	de	el	en	es	fr	it	nl	pt	th	zh	pl	tr	vi	Avg
XLM-R	69.5	86.6	88.5	72.2	59.8	92.3	87.6	96.4	88.8	89.3	92.1	88.6	90.0	58.6	59.7	84.0	73.6	56.7	79.7
XLM-R CT	69.4	86.1	88.6	68.2	61.5	91.5	88.5	96.4	89.6	89.8	92.2	88.9	90.0	57.8	62.9	84.0	74.1	57.9	79.9
Random	68.5	86.5	88.6	69.2	59.7	92.3	87.9	96.3	88.1	88.9	92.0	88.4	90.0	58.2	60.5	83.9	72.9	57.1	79.4
Ours	69.3	85.4	88.2	70.8	60.1	92.4	88.6	95.6	89.1	89.4	92.3	88.9	90.9	58.5	63.0	84.7	75.5	58.5	80.1

Table 12: Performance (F1 score) on POS tagging dataset. We examine the results on POS tagging on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.

Languages	en	ar	ru	sw	id	fi	bn	ko	te	Avg
XLM-R	60.7	52.5	50.1	50.2	63.8	51.2	31.1	23.0	30.9	45.9
XLM-R CT	60.1	53.7	50.6	54.5	66.1	52.4	31.0	21.6	26.7	46.3
Random	61.2	52.4	50.1	51.1	65.5	50.1	38.3	22.2	33.5	47.2
Ours	62.8	54.6	51.6	49.3	66.1	50.8	39.1	22.6	33.7	47.8

Table 13: Performance (F1 score) on TydiQA dataset. We examine the results on TydiQA on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.

Languages	en	es	de	el	ru	hi	ar	th	zh	tr	vi	Avg
XLM-R	83.0	76.1	73.2	72.4	73.3	68.0	66.0	68.0	51.7	67.4	73.8	70.3
XLM-R CT	83.4	75.9	73.4	72.4	73.7	69.2	65.8	68.1	52.8	66.7	74.0	70.5
Random	82.8	76.0	72.8	71.9	73.2	68.9	65.5	67.4	52.6	66.8	74.1	70.2
Ours	82.8	75.5	74.5	73.0	73.4	69.5	67.3	68.0	53.1	68.5	73.8	70.9

Table 14: Performance (F1 score) on XQuAD dataset. We examine the results on XQuAD on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.

Languages	en	de	el	tl	af	nl	ur	fr	pt	es	it	ar	id	ms	hi	fa	bg	ru	sw	fa
XLM-R	83.0	74.3	72.5	71.4	74.6	80.4	50.1	76.9	77.9	70.8	77.2	45.1	50.1	56.4	66.1	40.7	77.3	63.6	66.9	40.7
XLM-R CT	82.5	74.2	72.8	70.9	74.8	80.0	53.3	77.0	77.2	71.3	77.4	47.5	47.4	62.2	65.6	42.2	76.6	63.7	65.9	42.2
Random	82.7	74.2	72.6	71.2	74.7	80.2	51.7	77.0	77.6	71.0	77.3	46.3	48.7	59.3	65.9	41.4	77.0	63.7	66.4	41.4
Ours	83.1	76.4	74.2	73.1	78.2	81.8	47.2	78.7	78.2	74.5	78.4	51.1	49.6	68.7	65.2	45.4	78.0	64.1	67.0	45.4
Languages	mr	et	ja	zh	he	jv	eu	fi	yo	my	hu	ta	te	vi	ml	tr	ko	kk	bn	ka
XLM-R	59.5	72.3	18.3	24.7	52.1	58.5	60.2	75.3	41.5	51.7	76.1	53.7	47.0	65.6	61.3	73.9	49.7	44.6	66.3	65.5
XLM-R CT	60.9	71.7	18.1	21.9	50.8	58.2	60.1	74.5	38.3	51.2	75.2	52.9	48.2	65.7	60.7	74.5	49.0	49.6	66.2	64.7
Random	60.2	72.0	18.2	23.3	51.5	58.3	60.2	74.9	39.9	51.4	75.7	53.3	47.6	65.7	61.0	74.2	49.4	47.1	66.3	65.1
Ours	62.8	73.8	19.4	21.9	53.4	63.2	65.7	76.0	51.6	56.0	76.9	55.6	55.0	66.8	65.0	77.0	53.0	50.8	70.5	67.1

Table 15: Performance (F1 score) on NER dataset. We examine the results on NER on fine-tuning with English setting. Languages in different embedding sprachbunds are separated with vertical lines.