

On the Language Neutrality of Pre-trained Multilingual Representations

Jindřich Libovický¹ and Rudolf Rosa² and Alexander Fraser¹

¹Center for Information and Language Processing, LMU Munich, Germany

²Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic

{libovicky, fraser}@cis.lmu.de rosa@ufal.mff.cuni.cz

Abstract

Multilingual contextual embeddings, such as multilingual BERT (mBERT) and XLM-RoBERTa, have proved useful for many multi-lingual tasks. Previous work probed the cross-linguality of the representations indirectly using zero-shot transfer learning on morphological and syntactic tasks. We instead focus on the language-neutrality of mBERT with respect to lexical semantics. Our results show that contextual embeddings are more language-neutral and in general more informative than aligned static word-type embeddings which are explicitly trained for language neutrality. Contextual embeddings are still by default only moderately language neutral, however, we show two simple methods for achieving stronger language neutrality: first, by unsupervised centering of the representation for languages, and second by fitting an explicit projection on small parallel data. In addition, we show how to reach state-of-the-art accuracy on language identification and match performance of statistical methods for word alignment in parallel sentences.

1 Introduction

Multilingual BERT (mBERT; Devlin et al. (2019)) gained popularity as a contextual representation for many multilingual tasks, e.g., dependency parsing (Kondratyuk and Straka, 2019a; Wang et al., 2019), cross-lingual natural language inference (XNLI) or named-entity recognition (NER) (Pires et al., 2019; Wu and Dredze, 2019; Kudugunta et al., 2019). Recently, a new pre-trained model XLM-RoBERTa (XLM-R; Conneau et al. (2019)) claimed to outperform mBERT both on XNLI and NER tasks. DistilmBERT (Sanh et al., 2019) promises to deliver comparable results to mBERT at a significantly lower computational cost.

Pires et al. (2019) present an exploratory paper showing that mBERT can be used cross-lingually for zero-shot transfer in morphological and syntactic tasks, at least for typologically similar languages. They also study an interesting semantic task, sentence-retrieval, with promising initial results. Their work leaves many open questions in terms of how good the cross-lingual mBERT representation is for lexical semantics, motivating our work.

In this paper, we directly assess cross-lingual properties of multilingual representation on tasks where lexical semantics plays an important role and present two simple methods of achieving better language neutrality.

Multilingual capabilities of representations are often evaluated by zero-shot transfer (Hu et al., 2020). However, in such a setup, we can never be sure if the probing model did not overfit for the language used for training. The training of models is usually done using a validation set from the same language as the training set (otherwise it would not be zero-shot), even when it would have been better to stop the training earlier. This overfitting on the original language can pose a disadvantage for information-richer representations.

To avoid such methodological issues, we selected tasks that only involve a direct comparison of the representations: cross-lingual sentence retrieval, word alignment (WA), and machine translation quality estimation (MT QE). Additionally, we explore how the language is represented in the embeddings by training language ID classifiers and by assessing how the representation similarity corresponds to phylogenetic language families.

Our results show that contextual representations are more language neutral than static word embeddings which have been explicitly trained to represent matching word similarly. However, contextual representations still strongly carry information about the language.

By appropriately modifying the contextual embeddings to be language neutral, we reach state-of-the-art results on language identification and word alignment using simple straightforward setups: centering the representations for languages or fitting an explicit projection on small parallel data.

We further show that XLM-RoBERTa (XLM-R; Conneau et al. (2019)) outperforms mBERT in sentence retrieval and MT QE, while offering a similar performance for language identification and WA.

2 Related Work

Multilingual representations, mostly mBERT, were already tested in a wide range of tasks. Despite many positive results, the findings in the literature also indicate limited language neutrality.

- Wang et al. (2019) reached impressive results in zero-shot dependency parsing. However, the representation used for the parser was a bilingual projection of the contextual embeddings based on word alignment trained on parallel data.
- Pires et al. (2019) examined the cross-lingual properties of mBERT on zero-shot NER and part-of-speech (POS) tagging but the success of zero-shot transfer strongly depends on how typologically similar the languages are. Similarly, Wu and Dredze (2019) trained good multilingual models for POS tagging, NER, and XNLI, but struggled to achieve good results in the zero-shot setup. Rönqvist et al. (2019) draw similar conclusions for language-generation tasks.
- Pires et al. (2019) also assessed mBERT on cross-lingual sentence retrieval between three language pairs. They observed that if they subtract the average difference between the embeddings from the target language representation, the retrieval accuracy significantly increases. We systematically study this idea in the later sections.
- XTREME (Hu et al., 2020), a recently introduced benchmark for multilingual representation evaluation, assesses the representation on a wider range of zero-shot transfer tasks that include natural language inference (Conneau et al., 2018) and question answering (Artetxe et al., 2019; Lewis et al., 2019), their results show a clearly superior performance of XLM-R compared to mBERT.

The literature clearly shows that downstream task models can extract relevant features from the multilingual representations (Wu and Dredze, 2019; Kudugunta et al., 2019; Kondratyuk and Straka, 2019a). But they do not directly show language-neutrality, i.e., to what extent similar phenomena are represented similarly across languages.

Based on previous work it is impossible to say if this happens in a language-agnostic way or if it is based on some implicit language identification. Our choice of evaluation tasks eliminates this risk by directly comparing the representations.

3 Centering Representations

One way to achieve stronger language neutrality is by suppressing the language identity, only keeping what encodes the sentence meaning. It can be achieved for instance using an explicit projection, however, training such a projection requires parallel data. Instead, we explore a simple unsupervised method: representation centering.

Following Pires et al. (2019), we hypothesize that a sentence representation in mBERT is composed of a language-specific component, which identifies the language of the sentence, and a language-neutral component, which captures the meaning of the sentence in a language-independent way. We assume that the language-specific component is similar across all sentences in the language.

We estimate the language centroid as the mean of the representations for a set of sentences in that language and subtract the language centroid from the contextual embeddings. We try to remove the

language-specific information from the representations by centering the sentence representations in each language so that their average lies at the origin of the vector space.

The intuition behind this is that within one language, certain phenomena (e.g. function words) would be very frequent, thus being quite prominent in the mean of the representations for that language (but not for a different language), while the phenomena that vary among sentences of the language (and thus presumably carry most of the meaning) would be averaged out in the centroid. We thus hypothesize that by subtracting the centroid, we remove the language-specific features (without much loss of the meaning content), making the meaning-bearing features more prominent.

We analyze the semantic properties of the original and the centered representations on a range of probing tasks. For all tasks, we test all layers of the model. For tasks utilizing a single-vector sentence representation, we test both the [cls] token vector and mean-pooled states.

4 Probing Tasks

We employ five probing tasks to evaluate the language neutrality of the representations.

The first two tasks analyze the contextual embeddings. The other three tasks are cross-lingual NLP problems, all of which can be treated as a general task of a cross-lingual estimation of word or sentence similarities. Supposing we have sufficiently language-neutral representations, we can estimate these similarities using the cosine distance of the representations; the performance in these tasks can thus be viewed as a measure of the language-neutrality of the representations.

Moreover, in addition to such an unsupervised approach, we can also utilize actual training data for the tasks to further improve the performance of the probes; this does not tell us much more about the representations themselves, but leads to a nice by-product of reaching state-of-the-art accuracies for two of the tasks.

Language Identification. With a representation that captures all phenomena in a language-neutral way, it should be difficult to determine what language the sentence is written in. Unlike our other tasks, language ID requires fitting a classifier. We train a linear classifier on top of a sentence representation.

Language Similarity. Previous work (Pires et al., 2019; Wang et al., 2019) show that models can be transferred better between more similar languages, which suggests that similar languages tend to get similar representations. We quantify this observation by V-measure between language families and hierarchical clustering of the language centroids (Rosenberg and Hirschberg, 2007). We cluster the centroid by their cosine distance using Nearest Point Algorithm and stop the clustering with a number of clusters equal to the number of language families in the data.

Parallel Sentence Retrieval. For each sentence in a multi-parallel corpus, we compute the cosine distance of its representation with representations of all sentences on the parallel side of the corpus and select the sentence with the smallest distance.

Besides the plain and centered representations, we evaluate explicit projection of the representations into the “English space”. The projection is fitted by minimizing the element-wise mean square error between the representation of an English sentence and a linear projection of the representation of its translation.

Word Alignment. WA is the task of matching words which are translations of each other in parallel sentences. WA is a key component of statistical machine translation systems (Koehn, 2009). While sentence retrieval could be done with keyword spotting, computing bilingual WA requires resolving detailed correspondence on the word level. Unsupervised statistical methods trained on parallel corpora (Och and Ney, 2003; Dyer et al., 2013) still pose a strong baseline for the task. In a work parallel to ours, Sabet et al. (2020), present a more complex alternative way of leveraging contextual representations for word alignment that outperforms the statistical methods.

For a pair of parallel sentences, we find the WA as a minimum weighted edge cover of a bipartite graph. We create an edge for each potential alignment link, weight it by the cosine distance of the token

representations, and find the WA as a minimum weighted edge cover of the resulting bipartite graph. Unlike statistical methods, this does not require parallel data for training.

To make the algorithm prefer monotonic alignment, we add distortion penalty of $1/d$ to each edge where d is the difference in the absolute positions of the respective tokens in the sentence. We add the penalty with a weight that is hyper-parameter of the method estimated on a devset.

We keep the tokenization that is provided in the word alignment dataset. In the matching phase, we represent the tokens that get split into multiple subwords as the average of the embeddings of the subwords.

Note that this algorithm is invariant to representation centering. Centering the representation would shift all vectors by a constant, therefore all weights would change by the same offset, not influencing the edge cover. We evaluate WA using F_1 over sure and possible alignments in manually aligned data.

MT Quality Estimation. MT QE assesses the quality of an MT system output without having access to a reference translation. Semantic adequacy that we can estimate by comparing representations of the source sentence and translation hypothesis can be a strong indicator of the MT quality. The standard evaluation metric is the Pearson correlation with the Translation Error Rate (TER)—the number of edit operations a human translator would need to do to correct the system output. QE is a more challenging task than the previous ones because it requires capturing more subtle differences in meaning.

We evaluate how cosine distance of the representation of the source sentence and of the MT output reflects the translation quality. In addition to plain and centered representations, we also test trained bilingual projection, and a fully supervised regression trained on the shared task training data.

We use the same bilingual projection into English space fitted by linear regression on the small parallel data that we used for sentence retrieval.

For the supervised regression, we use a multilayer perceptron directly predicting the value of the translation error rate provided in the training data.

5 Probed Models

Static word embeddings. As one of the baselines in all our experiments, we use aligned static word embeddings. Unlike the hidden states of the pre-trained Transformers, they do not capture sentence context. However, they were explicitly trained to be language-neutral with respect to lexical semantics. We represent sentences as an average of the embeddings of the words.

Multilingual BERT (Devlin et al., 2019) is a deep Transformer (Vaswani et al., 2017) encoder that is trained in a multi-task learning setup, first, to be able to guess what words were masked-out in the input and, second, to decide whether two sentences follow each other in a coherent text.

We use a pre-trained mBERT model that was made public with the BERT release.¹ The model dimension is 768, hidden layer dimension 3072, self-attention uses 12 heads, the model has 12 layers. It uses a vocabulary of 120k wordpieces that is shared for all languages.

It is trained using a combination of a masked language model (MLM) objective and sentence-adjacency objective. For the MLM objective, 15% of input subwords are masked out and the model predicts the masked subwords. For the sentence-adjacency objective, a special [cls] token is prepended to the input. The embedding corresponding to this token is used as an input to a classifier predicting if the input sentences are adjacent.

Therefore, for models based on mBERT, we experiment both with [cls] vector and the *mean-pooled* vector, i.e., average embeddings for the rest of the tokens.

UDify. The UDify model (Kondratyuk and Straka, 2019a) uses mBERT to train a single model for dependency parsing and morphological analysis of 75 languages. During training, mBERT is fine-tuned, which improves the accuracy. Results on zero-shot parsing suggest that the fine-tuning leads to better language neutrality with respect to morphology and syntax.

¹<https://github.com/google-research/bert>

Ing-free. In this experiment, we try to make the representations more language-neutral by removing the language identity from the model using an adversarial approach. We continue training mBERT in a multi-task learning setup with the MLM objective (Devlin et al., 2019) without the sentence adjacency objective, i.e., the same way as XLM-R. It is trained jointly with adversarial language ID classifiers (Elazar and Goldberg, 2018) using the same dataset as for the language ID tasks. The classifier is separated from the rest of the model by a gradient-reversal layer (Ganin and Lempitsky, 2015), which negates the gradients flowing from the classifier into the model. Intuitively, we can say that the rest of the model is trying to fool the classifier, whereas the classifier tries to improve.

DistillmBERT. This model was inferred from mBERT by knowledge distillation (Sanh et al., 2019). The model only has 6 layers instead of 12, the rest of the hyperparameters remain the same. The model is initialized with a subset of the original mBERT parameters and trained on similar training data. The model was optimized towards cross-entropy of its output distribution with respect to output of the teacher mBERT model while keeping the MLM objective in the multitask learning setup. The model is forced to use a smaller space to obtain the representation and therefore it might leverage the similarities between language and thus reach better language neutrality.

XLM-RoBERTa. Conneau et al. (2019) claim that the original mBERT is under-trained and train a similar model on a larger dataset that consists of two terabytes of plain text extracted from CommonCrawl (Wenzek et al., 2019). Unlike mBERT, XLM-R uses a SentencePiece-based vocabulary (Kudo and Richardson, 2018) of 250k tokens, the rest of the architecture remains the same as in the case of mBERT. The model is trained using the MLM objective, only without the sentence adjacency prediction.

6 Experimental Setup

To train the language ID classifier, for each of the BERT languages we randomly select 110k sentences of at least 20 characters from Wikipedia, and keep 5k for validation and 5k for testing for each language. The training data is also used for estimating the language centroids and training the *Ing-free* version of the model.

For parallel sentence retrieval, we use a multi-parallel corpus of test data from the WMT14 evaluation campaign (Bojar et al., 2014) with 3,000 sentences in Czech, English, French, German, Hindi, and Russian. To compute the linear projection (for the special linear projection experimental condition), we used the WMT14 development data (500–3000 sentences per language pair).

We use manually annotated WA datasets to evaluate word alignment between English on one side and Czech (2.5k sent.; Mareček (2016)), Swedish (192 sent.; Holmqvist and Ahrenberg (2011)), German (508 sent.), French (447 sent.; Och and Ney (2000)) and Romanian (248 sent.; Mihalcea and Pedersen (2003)) on the other side. We compare the results with FastAlign (Dyer et al., 2013) and Efmara (Östling and Tiedemann, 2016) models which were provided with 1M additional parallel sentences from ParaCrawl (Esplà et al., 2019).

For MT QE, we use English-German training and test data provided for the WMT19 QE Shared Task (Fonseca et al., 2019) which consist of source sentences, automatic translations, and manually corrected reference translations. For the supervised estimation, we use a multi-layer perceptron with a hidden layer of size 256, trained to estimate the HTER value using the mean-square-error loss.

For the static word embeddings, we use pre-trained tables provided by Joulin et al. (2018).² The embeddings were trained on Wikipedia and aligned with a projection trained on small bilingual dictionaries. The number of word types captured in the embedding tables span from 350k for Romanian to 2.5M for English.

The experiments with contextualized embeddings are implemented using the Transformers package (Wolf et al., 2019), which we also use for obtaining the pre-trained models; except for UDify, which was obtained from (Kondratyuk and Straka, 2019b). The *Ing-free* mBERT version was finetuned using the same data that was used for language identification.

Our source code is available at <https://github.com/jlibovicky/assess-multilingual-bert>.

²<https://fasttext.cc/docs/en/aligned-vectors.html>

	mBERT	UDify	lng-free	Distil	XLM-R
[cls]	.935	.938	.796	.953	—
[cls], cent.	.867	.851	.337	.826	N/A
mean-pool	.960	.959	.951	.953	.950
mean-pool, cent.	.853	.854	.855	.826	.846

Table 1: Accuracy of language identification, values from the best-scoring layers.




Figure 1: Language ID accuracy for different layers of mBERT.

7 Results

Language Identification. Table 1 shows that for mBERT, centering the sentence representations considerably decreases the accuracy of language ID, especially in the case of mean-pooled embeddings. This indicates that the centering procedure indeed removes the language-specific information to a great extent.

For comparison the state-of-the-art language ID model from FastText (Grave et al., 2018) reaches accuracy of 91.4 % with a pretrained model, and 91.8 % when retrained on our training data, i.e., slightly worse than our best model based on mBERT. Langid.py (Lui and Baldwin, 2012) reaches 90.1 % when trained on the same dataset.

Adversarial fine-tuning prevented the language identification only from the [cls] vector and only marginally for mean-pooling. This supports the hypothesis that language identity is derived from the presence of function words and structures and representation centering suppresses these frequent phenomena.

Centering the representations within languages requires knowing the language in advance. Centering adds language-specific information to the representation which the classifier might take advantage of. However, because the centering decreases the accuracy, we can interpret this as removing information about the language.

For further comparison, we conduct the same experiment with aligned word embeddings for 44 languages (Joulin et al., 2018), the language ID accuracy is 99.5% with a drop to 2.3% being the same as assigning language by chance which supports our intuition about centering being a removal of frequent patterns. Note however, that such an experiment cannot be considered language identification because we need to know the language in advance to use the matching embeddings table, so the accuracy is not comparable with other experiments.

Language Similarity. Figure 2 is a tSNE plot (Maaten and Hinton, 2008) of the language centroids, showing that the similarity of the centroids tends to correspond to the similarity of the languages. Table 2 confirms that the hierarchical clustering of the language centroids mostly corresponds to the language families.

XLM-R not only has a slightly worse performance in language ID, but also has worse performance in capturing language similarity. We hypothesize that this is because of the different approaches used in training the models. The next-sentence prediction used to train mBERT leads to stronger language-specific information because this sort of information is helpful in determining if two sentences are adjacent.




Figure 2: Language centroids of the mean-pooled representations from the 8th layer of cased mBERT on a tSNE plot with highlighted language families.

mBERT	82.42
UDify	80.03
Ing-free	80.59
XLM-R	69.35
Distil	81.30
random	62.14

Table 2: Clustering of language centroids, evaluated with V-Measure against genealogical language families with at least three mBERT languages.

	SWE	mBERT	UDify	Ing-free	Distil	XLM-R
[cls]	—	.639	.462	.549	.420	—
[cls], cent.	—	.684	.660	.686	.505	—
[cls], proj.	—	.915	.933	.697	.830	—
mean-pool	.113	.776	.314	.755	.600	.883
mean-pool, cent.	.496	.838	.564	.828	.770	.923
mean-pool, proj.	.650	.983	.906	.983	.980	.996

Table 3: Average accuracy for sentence retrieval over all 30 language pairs compared to static bilingual word embeddings (SWE).

Parallel Sentence Retrieval. Results in Table 3 reveal that the representation centering dramatically improves the retrieval accuracy, showing that it makes the representations more language-neutral. However, an explicitly learned projection of the representations leads to a much greater improvement, reaching a close-to-perfect accuracy, even though the projection was fitted on relatively small parallel data. The accuracy is usually higher for mean-pooled states than for the [cls] embedding and varies among the languages too (see Table 4).

The accuracy also varies according to the layer of mBERT used (see Figure 3). The best-performing is the 8th layer, both for mBERT and XLM-R. This is consistent both among models and among tasks.

Similar trends hold for all models. XLM-R significantly outperforms all models. The UDify model that was finetuned for syntax seems to significantly lose semantic abilities. Adversarial finetuning did not improve the performance.

	cs	de	en	es	fr	ru
cs	—	.812	.803	.821	.795	.836
de	.806	—	.845	.833	.818	.816
en	.783	.834	—	.863	.860	.809
es	.805	.824	.863	—	.869	.822
fr	.784	.822	.861	.859	—	.811
ru	.828	.820	.810	.826	.817	—

	cs	de	en	es	fr	ru
en	—	.917	.935	.941	.926	.919
cs	.925	—	.907	.913	.896	.923
de	.938	.913	—	.921	.904	.912
es	.936	.907	.916	—	.934	.908
fr	.928	.903	.917	.935	—	.905
ru	.920	.910	.918	.910	.903	—

Table 4: Sentence retrieval scores for the 8th layer of mBERT and XLM-R models.

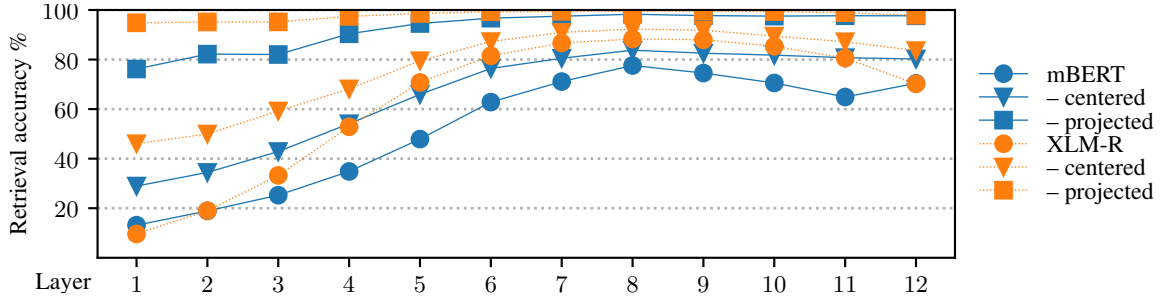


Figure 3: Accuracy of sentence retrieval for mean-pooled contextual embeddings from BERT layers.

en-	FastAlign	Efmaral	SWE	mBERT	UDify	Ing-free	Distil	XLM-R
cs	.692	.729	.501 – .540	.738	.708	.744	.660	.731
sv	.438	.501	.272 – .331	.478	.459	.468	.454	.461
de	.741	.759	.473 – .515	.767	.731	.768	.723	.762
fr	.583	.589	.371 – .435	.612	.581	.607	.582	.591
ro	.690	.742	.448 – .470	.703	.696	.704	.669	.732

Table 5: Maximum F_1 score (usually the 8th layer) for WA across layers compared with FastAlign baseline. For static word embeddings (SWE), we report the difference from introducing the distortion penalty.

Representation centering drastically improves accuracy. An additional 50% error reduction is achievable via learning a projection on relatively small parallel data leading to close-to-perfect accuracy.

Word Alignment. Table 5 shows that WA based on mBERT and XLM-R representations match the state-of-the-art aligners trained on a large parallel corpus. WA techniques based on multilingual contextual representations can thus be used as a replacement of state-of-the-art statistical methods without the use of parallel data.

The results show that word-level semantics is well captured by the contextual embeddings. Furthermore, the distortion penalty does not seem to influence the alignment quality when using the contextual embeddings, whereas for the static word embeddings, it can make a difference of 3–6 F_1 points. This shows that the contextual embeddings encode information about the relative word position in the sentence across languages. However, their main advantage is still the context-awareness, which allows accurate alignment of function words.

Similarly to sentence retrieval, we experimented with explicit projection trained on parallel data. We used an expectation-maximization approach that alternately aligned the words and learned a linear projection between the representations. This algorithm only brings a negligible improvement of .005 F_1 points.

MT Quality Estimation. Table 6 reveals that measuring the distance of non-centered sentence vectors does not correlate with MT quality at all; centering or explicit projection only leads to a mild correlation. Unlike sentence retrieval, QE is more sensitive to subtle differences between sentences, while the projection only seems to capture rough semantic correspondence. Note also that Pearson correlation used as an evaluation metric for QE might not favor the cosine distance because semantic similarity might not linearly correspond to HTER.

Supervised regression using either only the source or only MT output also shows a respectable correlation. The source sentence embedding alone can be used for a reasonable QE. This means that the source sentence complexity is already a strong indicator of the translation quality. The fact that using the target sentence embedding alone leads to almost as good results as using both the source and the hypothesis suggests that the structure of the translation hypothesis is what plays the important role. We must interpret the modest gain from concatenating the sentence representations as QE not being a suitable task for probing semantic properties of multilingual representations, because semantic adequacy is only a

	SWE	mBERT	UDify	lng-free	Distil	XLM-R
centered	.020	.005	.039	.026	.001	.001
projection	.038	.163	.167	.136	.241	.190
supervised: SRC only	.349	.362	.368	.349	.342	.388
supervised: TGT only	.339	.352	.375	.343	.344	.408
supervised full	.332	.419	.413	.411	.389	.431

Table 6: Correlation of estimated MT quality with HTER for WMT19 English-to-German translation.

marginally important aspect of MT QE.

The experiments with QE show that all tested contextual sentence representations carry information about sentence difficulty for MT and structural plausibility, however, unlike lexical-semantic features, this information is not accessible via simple embedding comparison.

8 Conclusions

Using a set of semantically oriented tasks, we showed that unsupervised BERT-based multilingual contextual embeddings capture similar semantic phenomena quite similarly across different languages. Surprisingly, in cross-lingual semantic similarity tasks, employing cosine similarity of the contextual embeddings without any tuning or adaptation clearly and consistently outperforms cosine similarity of static multilingually aligned word embeddings, even though these were explicitly trained to be language-neutral using bilingual dictionaries.

Nevertheless, we found that vanilla contextual embeddings contain a strong language identity signal, as demonstrated by their state-of-the-art performance for the language identification task. We hypothesize this is due to the sentence-adjacency objective used during training, because language identity is a strong feature for adjacency. We thus explored two ways of removing the language ID from the representations, in an attempt to make them even more cross-lingual. While adversarial fine-tuning of mBERT did not meet the expectations, a simpler unsupervised approach of language-specific centering of the representations managed to reach the goal to some extent, leading to higher performance of the centered representations in the probing tasks; the adequacy of the approach is also confirmed by a strong performance of the computed language centroids in estimating language similarity. Still, an even stronger language-neutrality of the representations can be achieved by fitting a supervised linear projection on a small set of parallel sentences.

Although representation centering leads to satisfactory language neutrality, it still requires knowing in advance what the language is. The future work thus should focus on representations that are more language-neutral by default, not requiring subsequent language-dependent modifications.

Acknowledgments

We would like to thank Philipp Dufter and Masoud Jalili Sabet for fruitful extensive discussions of the work.

References

- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. On the cross-lingual transferability of monolingual representations. *CoRR*, abs/1910.11856.
- Ondrej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA, June. 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, October–November. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, F. Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. *ArXiv*, abs/1911.02116.
- 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, June. Association for Computational Linguistics.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia, June. Association for Computational Linguistics.
- Yanai Elazar and Yoav Goldberg. 2018. Adversarial removal of demographic attributes from text data. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 11–21, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Miquel Esplà, Mikel Forcada, Gema Ramírez-Sánchez, and Hieu Hoang. 2019. ParaCrawl: Web-scale parallel corpora for the languages of the EU. In *Proceedings of Machine Translation Summit XVII Volume 2: Translator, Project and User Tracks*, pages 118–119, Dublin, Ireland, 19–23 August. European Association for Machine Translation.
- Erick Fonseca, Lisa Yankovskaya, André F. T. Martins, Mark Fishel, and Christian Federmann. 2019. Findings of the WMT 2019 shared tasks on quality estimation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 1–10, Florence, Italy, August. Association for Computational Linguistics.
- Yaroslav Ganin and Victor Lempitsky. 2015. Unsupervised domain adaptation by backpropagation. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1180–1189, Lille, France, 07–09 Jul. PMLR.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning word vectors for 157 languages. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Maria Holmqvist and Lars Ahrenberg. 2011. A gold standard for English-Swedish word alignment. In *Proceedings of the 18th Nordic Conference of Computational Linguistics (NODALIDA 2011)*, pages 106–113, Riga, Latvia, May. Northern European Association for Language Technology (NEALT).
- 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. *CoRR*, abs/2003.11080.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2984, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Philipp Koehn. 2009. *Statistical Machine Translation*. Cambridge University Press.
- Dan Kondratyuk and Milan Straka. 2019a. 75 languages, 1 model: Parsing universal dependencies universally. 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 2779–2795, Hong Kong, China, November. Association for Computational Linguistics.
- Dan Kondratyuk and Milan Straka. 2019b. UDify pretrained model. LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University, <http://hdl.handle.net/11234/1-3042>.

- Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium, November. Association for Computational Linguistics.
- 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, November. Association for Computational Linguistics.
- Patrick S. H. Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2019. MLQA: evaluating cross-lingual extractive question answering. *CoRR*, abs/1910.07475.
- Marco Lui and Timothy Baldwin. 2012. langid.py: An off-the-shelf language identification tool. In *Proceedings of the ACL 2012 System Demonstrations*, pages 25–30, Jeju Island, Korea, July. Association for Computational Linguistics.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605.
- David Mareček. 2016. Czech-english manual word alignment. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Rada Mihalcea and Ted Pedersen. 2003. An evaluation exercise for word alignment. In *Proceedings of the HLT-NAACL 2003 Workshop on Building and Using Parallel Texts: Data Driven Machine Translation and Beyond*, pages 1–10.
- Franz Josef Och and Hermann Ney. 2000. Improved statistical alignment models. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 440–447, Hong Kong, October. Association for Computational Linguistics.
- Franz Josef Och and Hermann Ney. 2003. A systematic comparison of various statistical alignment models. *Computational Linguistics*, 29(1):19–51.
- Robert Östling and Jörg Tiedemann. 2016. Efficient word alignment with Markov Chain Monte Carlo. *Prague Bulletin of Mathematical Linguistics*, 106:125–146, October.
- 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, July. Association for Computational Linguistics.
- Samuel Rönnqvist, Jenna Kanerva, Tapio Salakoski, and Filip Ginter. 2019. Is multilingual BERT fluent in language generation? In *Proceedings of the First NLPL Workshop on Deep Learning for Natural Language Processing*, pages 29–36, Turku, Finland, 30 September. Linköping University Electronic Press.
- Andrew Rosenberg and Julia Hirschberg. 2007. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 410–420, Prague, Czech Republic, June. Association for Computational Linguistics.
- Masoud Jalili Sabet, Philipp Dufter, and Hinrich Schütze. 2020. Simalign: High quality word alignments without parallel training data using static and contextualized embeddings. *CoRR*, abs/2004.08728.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. In *NeurIPS EMC2 Workshop*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30*, pages 6000–6010, Long Beach, CA, USA, December. Curran Associates, Inc.
- Yuxuan Wang, Wanxiang Che, Jiang Guo, Yijia Liu, and Ting Liu. 2019. Cross-lingual BERT transformation for zero-shot dependency parsing. 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 5725–5731, Hong Kong, China, November. Association for Computational Linguistics.
- Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2019. Ccnet: Extracting high quality monolingual datasets from web crawl data. *CoRR*, abs/1911.00359.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R'emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT. 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 833–844, Hong Kong, China, November. Association for Computational Linguistics.