

Similarity of Sentence Representations in Multilingual LMs: Resolving Conflicting Literature and a Case Study of Baltic Languages

Maksym DEL, Mark FISHEL

Institute of Computer Science, University of Tartu, Estonia

{maksym,mark}@tartunlp.ai

Abstract. Low-resource languages, such as Baltic languages, benefit from Large Multilingual Models (LMs) that possess remarkable cross-lingual transfer performance capabilities. This work is an interpretation and analysis study into cross-lingual representations of Multilingual LMs. Previous works hypothesized that these LMs internally project representations of different languages into a shared cross-lingual space. However, the literature produced contradictory results. In this paper, we revisit the prior work claiming that "BERT is not an Interlingua" and show that different languages do converge to a shared space in such language models with another choice of pooling strategy or similarity index. Then, we perform cross-lingual representational analysis for the two most popular multilingual LMs employing 378 pairwise language comparisons. We discover that while most languages share joint cross-lingual space, some do not. However, we observe that Baltic languages do belong to that shared space.¹

Keywords: interpretability, similarity, analysis, representations, cross-linguality, mBERT, XLM-R

1 Introduction

Multilingual Language Models such as mBERT (Devlin et al., 2019) or XLM-R (XLM-Roberta; Conneau et al. 2020a) achieve remarkable results on a variety of cross-lingual transfer tasks (Hu et al., 2020; Liang et al., 2020). Many works tried to understand how these models represent multiple languages but achieved incomplete and even conflicting results.

Notably, Muller et al. (2021), Conneau et al. (2020b) and Singh et al. (2019) performed representational similarity analysis comparing encoded sentences in different languages. However, they come up with two opposite conclusions.

¹ The code is available at <https://github.com/TartuNLP/xsim>

In particular, Singh et al. (2019) concluded that “mBERT is not an Interlingua”. They observed that cross-lingual similarity of sentence representations decreased in similarity as the model layer increased. By the term “interlingua,” authors most certainly meant the opposite of this pattern. Muller et al. (2021), on the contrary, found that early layers representations are less aligned than representations from middle layers. See Figure 1 for our reproduction of these conflicting results (we choose Estonian, Latvian, and Lithuanian, as leading examples, together with French and Polish for comparison). In this Figure, we compare languages with English.

This begs the question of which one we should rely on and why. Muller et al. (2021) backs up their representational analysis using probing task in the layer-wise ablation setting; Conneau et al. (2020b) also directly supports this conclusion. Singh et al. (2019), on the other hand, provided explanation based on possible tokenization bias issue. In any case, we believe that cross-lingual representational analysis should provide unambiguous results, so we investigate this issue in depth in this work.

Moreover, we provide a comprehensive study across 378 pairs of languages to have a course-grained view of the cross-lingual similarity. We use Baltic languages as our case study throughout the paper.

Specifically, we consider the following research questions:

1. Which cross-lingual pattern representational similarity suggests in the final analysis? (Answer: convergence pattern from Muller et al. (2021)).
2. Why do results in literature diverge, and how do we interpret them? (Answer: CLS-pooling is a poor choice for sentence summary, and the result is very sensitive to the choice of the similarity index).
3. What is the recommended way to quantify cross-linguality in multilingual LMs? (Answer: mean-pooling with SVCCA/CKA index aligns with the evidence from the behavior analysis literature better than the other options).
4. Does the resulting cross-lingual pattern generalize across all languages? (Answer: yes, it does, except for a few outlier languages).
5. How does representational similarity analysis look for Estonian, Latvian, and Lithuanian? (Answer: they follow the general convergence pattern as the majority of languages and are most similar to each other than to almost all other languages).




Fig. 1: Similarity of mBERT representations decreases (left) vs increases (right). Reproduced from Singh et al. (2019) and Muller et al. (2021) on similar data.

2 Resolving Conflicting Literature

In this section, we address the problem of conflicting evidence and conclusions between Singh et al. (2019) and Muller et al. (2021) (Figure 1).

Singh et al. (2019) as a part of their work took a multilingual Bert model² and used PWCCA cross-lingual similarity analysis index to measure the discrepancy between languages at each layer. Specifically, they chose a bilingual parallel corpus for a pair of languages (say $lang_A$ and $lang_B$) and passed the pair through the model. Then, they extracted hidden states at each layer, did a CLS pooling, and compared the extracted sentence representations. The resulting trend was that at higher layers, the distance between sentence representations is higher than at the lower layers.

Muller et al. (2021) performed a similar procedure with mBERT as part of their work but observed the opposite conclusion: the similarity of language representations is higher at deeper layers and lower at the first layers. In our work, we notice that apart from different datasets, Muller et al. (2021) also used a similar different similarity index (CKA) and pooling technique (mean-pooling).

As the conclusion of what happens to cross-lingual representations inside multilingual models is central to both works, we analyze these differences in detail in a shared setting and resolve the conflicting evidence.

2.1 Setup and Background

2.1.1 Data and model We do a setup similar to the one of Singh et al. (2019) and use the mBERT-cased model and four parallel datasets (en-et, en-it, en-lv, en-pl, en-fr; 10k examples for each pair). The parallel corpus is composed of Singh et al. (2019)'s extension of XNLI (Cross-lingual Natural Language Inference) dataset (Conneau et al., 2018). We choose Estonian, Latvian, and Lithuanian as our case study and use Polish and French to see how results for Baltic languages compare to high-resource Romance and Slavic languages.

We embed the source and target sentences with mBERT and pool the *CLS* tokens or perform a mean-pooling over tokens from each layer for each language pair. Next, we compare two parallel sets of sentence representations using the PWCCA, CKA, or SVCCA similarity indexes.

2.1.2 Representational Similarity Indexes In this subsection, we give a reader brief intuitive explanations and refer to the Kornblith et al. (2019) for a systematic mathematical description of the correlation similarity indexes.

All indexes compare two parallel sets of vectors by maximizing correlations (score 1 represents perfect similarity).

CCA (Canonical Correlation Analysis; Hardoon et al. 2004) is the correlation-based similarity analysis index. As formulated by Morcos et al., CCA "identifies the 'best' (maximizing correlation) linear relationships (under mutual orthogonality and norm constraints) between two sets of multidimensional variates". CCA score can be a mean of the resulting correlation coefficients.

² <https://github.com/google-research/bert/blob/master/multilingual.md>




Fig. 2: CLS-pooled representations compared using three different similarity measurement algorithms.

SVCCA (Singular Vector Canonical Correlation Analysis, Raghu et al. 2017) reduces the sensitivity of CCA to particular dimensions by performing Support Vector Decomposition on parallel vectors first and then applying CCA on the resulting components. The number of resulting CCA coefficients is the hyperparameter for the SVCCA, and we use 20 in this work.

CKA (Centered Kernel Alignment; Kornblith et al. 2019) is another similarity index that works by computing pairwise dot products between two parallel sets of vectors and correlating the resulting distance matrices.

PWCCA (Projection Weighted Canonical Correlation Analysis; Morcos et al. 2018) is an extension of the original CCA (Hardoon et al., 2004) that weights resulting CCA correlation coefficients based on their importance instead of taking a simple mean.

Also, we highlight that PWCCA is only invariant to the translation and isotropic scaling. CKA is also invariant to the orthogonal transforms, and SVCCA is invariant to any invertible linear transform.

2.2 Identifying the Issue

This section aims to find what caused the discrepancy between results in related works. They use different data and code, but we successfully reproduced the patterns with our datasets. However, they also differ in their choice of pooling strategy for sentence representation (CLS-pooling for Singh et al. (2019) and mean-pooling for Muller et al. (2021)) and similarity index (PWCCA for Singh et al. (2019) and CKA for Muller et al. (2021)).

2.2.1 Is this a similarity index issue? To answer the question of this subsection, we compute all similarities between CLS-pooled representations for all three main indexes. Figure 2 presents the results. It shows that as we change the similarity measure from PWCCA, we get a convergence pattern similar to one in Muller et al. (2021).

This explains the discrepancy, but the works additionally differ in sentence representation type. In the following subsection, we investigate this issue.




Fig. 3: Mean-pooled representations compared using three different similarity measurement algorithms.

2.2.2 Is this a pooling strategy issue? To answer the question of this subsection, we compute all similarity measures for representations, but this time obtained by averaging individual token representations (mean-pooling). Figure 3 presents the results. It shows that as we change the pooling type from CLS, we get rid of the divergence pattern Singh et al. observed with CLS-pooling.

PWCCA convergence pattern is less pronounced than other indexes but does not contradict Muller et al. (2021).

The divergence pattern only occurs when we use CLS-pooling and PWCCA metric simultaneously, so in the following subsections, we measure representational similarity with yet another method and question the usefulness of CLS pooling.

2.3 Debunking Divergence Pattern

So we only get a divergence pattern when using PWCCA measure over CLS-pooled representations. In the following subsection, we try cosine similarity as an alternative to the correlational indexes and check the semantic power of CLS pooling on a simple task of cross-lingual sentence matching.

2.3.1 Cosine similarity While PWCCA, SVCCA, and CKA are correlation-based indexes, we might also get insight from using simpler alternatives that have proven useful, especially in the NLP domain. One such metric is cosine similarity. We compute the pairwise cosine similarity between English sentences and their target (parallel) translations reporting an average score. We also report the scores for English and sentences and permuted target and compare the similarity over three pooling strategies and report results in Figure 4.

For the first five layers in CLS-pooling, cosine similarities between parallel sentences are close to one, just as cosine for random sentences. So they all point about in the same direction, and there is no straightforward distinction between translations and not translations.

However, despite representations for CLS being so similar, the question remains how semantically useful they are for cross-lingual similarity analysis and in general. Indeed,




Fig. 4: Cosine similarities of sentence representations under three different pooling strategies. "Target" box in legend means that (average) cosine similarity was measured between parallel sentences ("parallel") or arbitrary pairs of sentences ("random").

even two copies (maybe only slightly altered) of a random matrix would be perfectly similar by all similarity indexes.

2.3.2 Usefulness of CLS Our goal is to find out how cross-lingual mBERT represents languages across layers. A fundamental desired property of cross-linguality is that representations in multiple languages should have close representations. Moreover, these representations should be far closer to representations of other non-parallel sentences.

So let us see how well CLS satisfies these properties. To keep neutral regarding similarity measures, we step away from direct representational similarity analysis and set up a probing task that measures desired properties. We use the same data, and for each English sentence, we find the closest target sentence in the opposite language (out of all 10k targets) and declare "1" if the closest sentence was an actual translation of the source. We declare "0" otherwise. Then we compute the accuracy of this matching task for our language pairs.

We repeat this experiment where we pool the 1st token from each sentence to get a reference point for comparison. We present (averaged over languages) results in Figure 5.

The figure shows that accuracy for the CLS-pooling is almost zero at layers 0-4, which suggests it is not a helpful representation to rely upon when measuring cross-linguality, including using CCA-like measures. Matching by the first token is better than matching by CLS at these layers. In later layers CLS gets higher than first token pooling, but mean-pooling is about 0.3 accuracy points above.

While we showed that mean-pooling empirically performs better than CLS-pooling, there remains the question of why it is the case. We explain that CLS embedding does not have a robust explicit signal (representing a sentence meaning) at the layers other than last.

CLS-pooling is used as a sentence representation because, at the last layer, mBERT uses it to predict the next sentence in the corpus during pretraining. The next sentence prediction task uses the CLS token at the last layer. However, CLS at other hidden layers does not have this strong signal to represent the sentence due to the token mixing procedure in the Transformer layers. Each token position carries pieces of information




Fig. 5: Accuracy of closest sentence vector to each source sentence being an actual translation of this sentence. Measured over three pooling strategies and averaged over four languages (as before).

about itself as well pieces about other tokens, and CLS is only slightly more than a regular token at these layers (5). Mean pooling, on the contrary, gathers distributed information about all tokens at each layer and is thus free of the abovementioned issue.

This section showed that using CLS-pooling and PWCCA measures are suboptimal for cross-linguality representational analysis. They result in a pattern opposite to the expedient one when used together.

3 Analyzing Language Representations

In the previous section, we identified that the combination of PWCCA and CLS-pooling is not very suited for cross-lingual analysis. So in this so we use the combination of mean-pooling and CKA. We perform cross-lingual measurements over mean-pooled representations across two pretrained models and 378 pairwise language comparisons.

3.1 Quantifying cross-linguality Across Languages

3.1.1 Setup This section takes a more in-depth look at cross-linguality in multilingual LMs and explores its structure.

We use all 378 pairwise language comparisons from the extended XNLI dataset (Singh et al., 2019). By computing CKA across all language pairs at all layers, we determine that the seventh layer of XLM-R is the most “cross-lingual”. We choose to focus on XLM-R in addition to the mBERT as this model is more modern and has shown to be superior to mBERT.

Figure 6 also highlights Thai as the most distinct language. *mBERT uncased* pre-training did not include this language. mBERT has more “Excluded” languages, which is logical since it is a weaker cross-lingual model than the XLM-R.

3.1.2 Results and Discussion Figure 6 shows the boxplot covering CKA distances between all pairs of languages for XLM-R and mBERT.




Fig. 6: Per-layer CKA similarity for mBERT (a) and XLM-R (b) at each layer for combinations of 378 pairwise language comparisons. Dots at the bottom show that some language pairs drastically differ in similarity compared to the vast majority of other language pairs.




Fig. 7: All pairwise language similarities at eighth mBERT layer (a) and 7th XLM-R layer (b). There are few languages (Urdu and Hindi and Swahili and Thai) that models seem to locate far away from all others. Estonian, Latvian and Lithuanian do not resemble this property and belong to the main set.




Fig. 8: Agglomerative clustering for the eighth layer of mBERT (a) and the seventh layer of XLM-R (b) based on CKA distances. We can see that some languages (green for mBERT and orange for XLM-R) branch out from the main branch, and models seem to exclude them from the representational space. Estonian is grouped with Finnish (and then Hungarian), forming the Finno-Ugric branch, while Latvian groups with Lithuanian, forming the Baltic branch. Then these two branches merge into a single "Balto-Finno-Ugric" branch, which suggests these languages can be effectively considered collectively in NLP applications and research.

Dots at the bottom suggest that there are clear outlier language pairs. The most cross-lingual layer for XLM-R seems to be 7th, so let us open the box at this layer with Figure 7.

The figure suggests that the outliers are due to Urdu and Hindi (possibly also Swahili and Thai).

However, even within the shared cross-lingual space, languages also exhibit certain relationships. Thus we perform an agglomerative clustering on CKA distances to investigate this phenomenon and present the result in Figures 8.

The linguistic tree in Figure 8 clearly shows that languages in the majority branch (that we consider to be the shared cross-lingual space) structure in a meaningful way. Slavic languages are together, Swahili and Thai are isolated, while Scandinavian languages are again nearby. Urdu and Hindi expectedly occupy their separate branch as outliers.

The figure also shows that only a small subset of mostly low-resource languages like Swahili and Urdu gets its separate top-level branch (we leave finding the exact criteria for exclusion from the shared space for future work), while other languages, including ones analyzed in Singh et al. (2019), are a part of the shared cross-lingual space.

Let us also employ the t-SNE method to demonstrate how models separate extremely low-resource languages (Swahili and Urdu) from the joint space of European Languages (English, German, and French). See Figure 9 for the resulting graph.

Due to the t-SNE algorithm's nature, the representational subspaces' shapes and locations do not carry useful information. However, the graph supports our main point that not all languages are a part of joint cross-lingual space. Baltic languages, however, do belong to the joint space, as we showed in Figure 8.

In summary, in this chapter, we presented a bird's-eye view of how multilingual LMs represent languages across layers. By performing 378 pairwise comparisons we identified that the vast majority of languages share the common interlingual space at the middle layers of the network.




Fig. 9: Exploration of the cross-linguality pattern via t-SNE. In the first layers (0-2), the representations for all five languages are separated, followed by the middle layers (5-8), where the algorithm could not distinguish between German, French, and English while leaving Swahili and Urdu separated. In the last layers, however, the European selection of languages becomes separated again, finishing with its initial state at the last layers.

4 Conclusion

This paper identified, analyzed, and resolved conflicting literature and derived that mean-pooling with SVCCA/CKA similarity measure is the most suitable choice for the cross-lingual representational similarity analysis. Next, we showed that the pattern is not specific to mBERT and is present in other multilingual language models. Finally, we analyzed 378 pairwise language comparisons and found that not all languages share the cross-lingual space equally. We found, however, that Estonian and Baltic languages are grouped and constitute a part of this shared cross-lingual space.

References

- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., Stoyanov, V. (2020a). Unsupervised cross-lingual representation learning at scale, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, pp. 8440–8451. <https://www.aclweb.org/anthology/2020.acl-main.747>
- Conneau, A., Rinott, R., Lample, G., Williams, A., Bowman, S., Schwenk, H., Stoyanov, V. (2018). XNLI: Evaluating cross-lingual sentence representations, *Proceedings of EMNLP 2018*, pp. 2475–2485.
- Conneau, A., Wu, S., Li, H., Zettlemoyer, L., Stoyanov, V. (2020b). Emerging cross-lingual structure in pretrained language models, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, pp. 6022–6034. <https://www.aclweb.org/anthology/2020.acl-main.536>
- Devlin, J., Chang, M.-W., Lee, K., Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding, *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)*, Association for Computational Linguistics, Minneapolis, Minnesota, pp. 4171–4186. <https://www.aclweb.org/anthology/N19-1423>
- Hardoon, D., Szedmák, S., Shawe-Taylor, J. (2004). Canonical correlation analysis: An overview with application to learning methods, *Neural Computation* **16**, 2639–2664.
- Hu, J., Ruder, S., Siddhant, A., Neubig, G., Firat, O., Johnson, M. (2020). XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization, *CoRR abs/2003.11080*. <https://arxiv.org/abs/2003.11080>
- Kornblith, S., Norouzi, M., Lee, H., Hinton, G. E. (2019). Similarity of neural network representations revisited, *ArXiv abs/1905.00414*.
- Liang, Y., Duan, N., Gong, Y., Wu, N., Guo, F., Qi, W., Gong, M., Shou, L., Jiang, D., Cao, G., Fan, X., Zhang, B., Agrawal, R., Cui, E., Wei, S., Bharti, T., Qiao, Y., Chen, J., Wu, W., Liu, S., Yang, F., Majumder, R., Zhou, M. (2020). XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation, *CoRR abs/2004.01401*. <https://arxiv.org/abs/2004.01401>
- Morcos, A. S., Raghu, M., Bengio, S. (2018). Insights on representational similarity in neural networks with canonical correlation, *NeurIPS*.
- Muller, B., Elazar, Y., Sagot, B., Seddah, D. (2021). First align, then predict: Understanding the cross-lingual ability of multilingual BERT, *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Association for

Computational Linguistics, Online, pp. 2214–2231.

<https://www.aclweb.org/anthology/2021.eacl-main.189>

Raghu, M., Gilmer, J., Yosinski, J., Sohl-Dickstein, J. (2017). Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability, *NIPS*.

Singh, J., McCann, B., Socher, R., Xiong, C. (2019). BERT is not an interlingua and the bias of tokenization, *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, Association for Computational Linguistics, Hong Kong, China, pp. 47–55.

<https://www.aclweb.org/anthology/D19-6106>

Received August 14, 2022 , accepted August 15, 2022