

The Impact of Cross-Lingual Adjustment of Contextual Word Representations on Zero-Shot Transfer

Pavel Efimov^{1*}, Leonid Boytsov^{2†}, Elena Arslanova³, Pavel Braslavski^{3,4}

¹ITMO University, Saint Petersburg, Russia

²Bosch Center for Artificial Intelligence, Pittsburgh, USA

³Ural Federal University, Yekaterinburg, Russia

⁴HSE University, Moscow, Russia

pavel.vl.efimov@gmail.com, leo@boytsov.info

contilen@gmail.com, pbras@yandex.ru

Abstract

Large multilingual language models such as mBERT or XLM-R enable zero-shot *cross-lingual* transfer in various IR and NLP tasks. Cao et al. (2020) proposed a *data-* and *compute-efficient* method for cross-lingual adjustment of mBERT that uses a *small* parallel corpus to make embeddings of related words across languages similar to each other. They showed it to be effective in NLI for five European languages. In contrast we experiment with a typologically diverse set of languages (Spanish, Russian, Vietnamese, and Hindi) and extend their original implementations to new tasks (XSR, NER, and QA) and an additional training regime (continual learning). Our study reproduced gains in NLI for four languages, showed improved NER, XSR, and *cross-lingual* QA results in three languages (though some cross-lingual QA gains were not statistically significant), while *mono-lingual* QA performance never improved and sometimes degraded. Analysis of distances between contextualized embeddings of related and unrelated words (across languages) showed that fine-tuning leads to “forgetting” some of the cross-lingual alignment information. Based on this observation, we further improved NLI performance using continual learning. Our software is publicly available <https://github.com/pefimov/cross-lingual-adjustment>.

1 Introduction

Large *multi-lingual* language models such as mBERT or XLM-R pre-trained on large *unpaired* multilingual corpora enable zero-shot *cross-lingual* transfer (Libovický et al., 2019; Pires et al., 2019), which is sometimes effective even for languages not seen at the pre-training stage (Ebrahimi et al., 2021; Muller et al., 2021). Contextualized word representations produced by the models can be further aligned (across languages) with a modest amount of parallel data, e.g., by finding a rotation matrix using a bilingual dictionary (Mikolov et al., 2013). Such a post hoc alignment can improve zero-shot transfer for text parsing and understanding (Kulshreshtha et al., 2020; Wang et al., 2019b; Wang et al., 2019a). This approach is a more data- and computationally-efficient alternative to training a machine translation system on a large parallel corpus.

Cao et al. (2020) used a small parallel corpus for direct cross-lingual adjustment of the mBERT model and found this adjustment procedure to be more effective than the post hoc rotation. However, they experimented only with a single task: cross-lingual NLI, and restricted their study to European languages. In that, we are not aware of any systematic study that tested this procedure on a set of diverse languages and tasks.

To fill this gap, we rigorously evaluated the method using four typologically diverse languages (Spanish, Russian, Vietnamese, and Hindi) and four different tasks all of which are crucial to information retrieval and text mining: natural-language inference (NLI), question-answering (QA), named-entity recognition (NER), and cross-lingual sentence retrieval (XSR). QA and XSR are core information mining tasks directly targeting a user’s information need (Baeza-Yates et al., 1999; Manning et al., 2008) while NLI and NER are used in document (Tawfik, 2020) and query (Wen et al., 2019; Cheng et al., 2021)

*Contributed equally to the paper.

†Work done before joining Amazon.




Figure 1: Histograms of L_2 distances between pairs of mBERT *last-layer* representations for randomly sampled related (i.e., aligned) and unrelated word pairs from WikiMatrix (Hi-En): (a) original¹, (b) after cross-lingual adjustment, (c) after fine-tuning on English NLI data, (d) after cross-lingual adjustment and subsequent fine-tuning on English NLI data, (e) after cross-lingual adjustment and subsequent *continual* fine-tuning on English NLI data.

analysis. Named entities are the core of entity-oriented (Balog, 2018) and expert search (Brandesen et al., 2022). Entity-oriented retrieval models can also be used to boost recall of traditional keyword-based retrievers (Shehata et al., 2022; Ensan and Bagheri, 2017).

Furthermore, existing studies typically avoided statistical testing and presented results for a single seed. In contrast, we performed each experiment with five seeds and assessed statistical significance of the difference from a baseline. In addition to the original method of Cao et al. (2020) we evaluated its *continual-learning* extension (Ratcliff, 1990; Robins, 1995; Parisi et al., 2019) where we combined the target task loss with the cross-lingual adjustment loss (Caruana, 1996).

In our study, we ask the following research questions:

- R1 How does cross-lingually adjusted mBERT fine-tuned on English data and zero-shot transferred to a target language perform on various IR/NLP tasks and typologically different languages?
- R2 How does adjustment of mBERT on parallel data and fine-tuning for a specific task affect similarity of contextualized embeddings of semantically related and unrelated words across languages?
- R3 Inspired by our observation (see Fig. 1c-1d) that fine-tuning draws embeddings of *both* related and unrelated words closer to each other, which may negatively affect the cross-lingual transfer, we wonder if continual learning—with an auxiliary cross-lingual adjustment loss—can improve effectiveness of the zero-shot transfer.

Our experiments demonstrated the following:

- The cross-lingual adjustment of mBERT improves NLI in four target languages; NER, XSR, and *cross-lingual* QA in three languages (though some cross-lingual QA gains are not statistically significant). Yet, there is no statistically significant improvement for *mono-lingual* QA and a statistically significant deterioration on three out of eight QA datasets.
- When comparing L_2 distances between contextualized-embeddings of words across languages (Fig. 1b), we see that the cross-lingual adjustment of mBERT decreases the L_2 distance between related words while keeping unrelated words apart, which is in line with prior work (Zhao et al., 2021).
- However, we have found no prior work that inspected histograms obtained after fine-tuning. Quite surprisingly, we observe that fine-tuning of mBERT for a specific task draws embeddings of *both* related and unrelated words much closer to each other (Fig. 1c and Fig. 1d). Thus, fine-tuning causes the model to “forget” some of the cross-lingual information learned during adjustment.

¹To simplify the presentation, parts of the plots in Fig. 1a corresponding to small cumulative probabilities have been cropped in the camera-ready version of the paper.

- In that, continual learning allows the model to learn a target task while maintaining the separation between related and unrelated words (Fig. 1e). Continual learning consistently improves performance on NLI data. Aside from NLI, it improves XSR and NER *only* for Russian and there are no improvements on other QA or NER datasets.

In summary, our study contributes to a better understanding of (1) cross-lingual transfer capabilities of large multilingual language models and of (2) effectiveness of their cross-lingual adjustment across various tasks and languages. We believe our results support a conjecture that the cross-lingual adjustment of Cao et al. (2020) is more beneficial for cross-lingual tasks. Our software is publicly available <https://github.com/pefimov/cross-lingual-adjustment>.

2 Related Work

2.1 Cross-Lingual Zero-Shot Transfer with Multilingual Models

The success of mBERT in cross-language zero-shot regime on various tasks has inspired many papers that attempted to explain its cross-lingual abilities and limitations (Wu and Dredze, 2019; Conneau et al., 2020; K et al., 2020; Libovický et al., 2019; Dufter and Schütze, 2020; Chi et al., 2020; Pires et al., 2019; Artetxe et al., 2020; Chi et al., 2020). These studies showed that the multilingual models learn high-level abstractions common to all languages, which make transfer possible even when languages share no vocabulary. However, the gap between performance on English and a target language is smaller if the languages are cognate, i.e. share a substantial portion of model’s vocabulary, have similar syntactic structures, and are from the same language family (Wu and Dredze, 2019; Lauscher et al., 2020). Moreover, the size of target language data used for pre-training and the size of the model vocabulary allocated to the language also positively impacts cross-lingual learning performance (Lauscher et al., 2020; Artetxe et al., 2020).

Zero-shot transfer of mBERT or other multilingual transformer-based models from English to a different language was applied to many tasks including but not limited to cross-lingual information retrieval, POS tagging, dependency parsing, NER, NLI, and QA (Wu and Dredze, 2019; Wang et al., 2019b; Pires et al., 2019; Hsu et al., 2019; Litschko et al., 2021). XTREME data suite (Hu et al., 2020) and its successor XTREME-R (Ruder et al., 2021) are dedicated collections of tasks and corresponding datasets for evaluation of zero-shot transfer capabilities of large multilingual models from English to tens of languages. XTREME includes XSR, NLI, NER, and QA datasets used in the current study. Although transfer from English is not always an optimal choice (Lin et al., 2019; Turc et al., 2021), English still remains the most popular source language. Furthermore, despite there have been developed quite a few new models that differ in architectures, supported languages, and training data (Doddapaneni et al., 2021), mBERT remains the most popular cross-lingual model.

2.2 Cross-lingual Alignment of Embeddings

Mikolov et al. (2013) demonstrated that vector spaces can encode semantic relationships between words and that there are similarities in the geometry of these vectors spaces across languages. A variety of approaches have been proposed for aligning monolingual representations based on bilingual dictionaries and parallel sentences. The most widely used approach—which requires only a bilingual dictionary—consists in finding a rotation matrix that aligns vectors of two monolingual models (Mikolov et al., 2013). Lample et al. (2018) proposed an alignment method based on adversarial training, which does not require parallel data. Ruder et al. (2019) provide a comprehensive overview of alignment methods for pre-Transformer models.

Schuster et al. (2019) applied rotation to align contextualized ELMo embeddings (Peters et al., 2018) using “anchors” (averaged vectors of tokens in different contexts) and bilingual dictionaries. They showed improved results of cross-lingual dependency parsing using English as source and several European languages as target languages. Wang et al. (2019a) aligned English BERT and mBERT representations using rotation method and Europarl parallel data (Koehn, 2005). They employed the resulting embeddings in a cross-lingual dependency parsing model. The parser with aligned embeddings consistently outperformed zero-shot mBERT on 15 out of 17 target languages. Instead of aligning on a

word level, Aldarmaki and Diab (2019) performed a sentence-level alignment of ELMo embeddings and evaluated this approach on the parallel sentence retrieval task.

Cao et al. (2020) proposed to directly modify the mBERT model by bringing the vectors of semantically related words in different languages closer to each other. This was motivated by the observation that embedding spaces of different languages are not always isometric (Søgaard et al., 2018) and, hence, are not always amenable to alignment via rotation. The authors showed that mBERT simultaneously adjusted on five European languages consistently outperformed other alignment approaches on XNLI data. In the current study, we implement the approach with some modifications.

Liu et al. (2021) showed that combining continual learning with fine-tuning improved zero-shot transfer performance for NER and POS tagging. In that, they used cross-lingual sentence retrieval (XSR) and/or masked-language model (MLM) task as additional tasks. Although XSR can be seen as an alternative to the cross-lingual adjustment of Cao et al. (2020), the authors did not evaluate the effectiveness of zero-shot transfer after adjusting the model with XSR. In contrast, we evaluate the marginal effectiveness of continual learning with respect to already cross-lingually adjusted mBERT.

Kulshreshtha et al. (2020) compared different alignment methods (rotation vs. adjustment) on NER and slot filling tasks. According to their results, rotation-based alignment performs better on the NER task, while model adjustment performs better on slot filling. Zhao et al. (2021) continued this line of research and proposed several improvements of the model adjustment method: 1) z-normalization of vectors and 2) text normalization to make the input more structurally ‘similar’ to English training data. Experiments on XNLI dataset and translated sentence retrieval showed that vector normalization leads to more consistent improvements over zero-shot baseline compared to text normalization. Faisal and Anastasopoulos (2021) applied cross-lingually adjusted mBERT and XLM-R to cross-lingual open-domain QA and obtained improvements both on paragraph and span selection subtasks. However, they trained their models on machine-translated data, which is different from our zero-shot settings.

3 Methods

In this study, we use a multilingual BERT (mBERT) as the main model (Devlin et al., 2019). mBERT is a case-sensitive “base” 12-layer Transformer model (Vaswani et al., 2017) with 178M parameters.² It was trained with a masked language model objective on 104 languages with a shared WordPiece (Wu et al., 2016) vocabulary (using 104 Wikipedias). To balance the distribution of languages, high-resource languages were under-sampled and low-resource languages were over-sampled.³ For a number of NLP tasks, cross-lingual transfer of mBERT can be competitive with training a monolingual model using the training data in the target language.

We align cross-lingual embeddings by directly modifying/adjusting the language model itself, following the approach by Cao et al. (2020). The approach—which differs from finding a rotation matrix—proved to be effective in the NLI task. However, there are some differences in our implementation. In all cases, we work with one pair of languages at a time while Cao et al. (2020) adjusted mBERT for five languages at once.

BERT uses WordPiece tokenization (Wu et al., 2016), which splits sufficiently long words into *subword* tokens. We first word-align parallel data with *fast_align* (Dyer et al., 2013) and then average all subword tokens’ vectors.⁴

Based on alignments in parallel data, we obtain a collection of word pairs (s_i, t_i) : s_i from the source language, t_i from the target one. From these alignments we can obtain their mBERT vector representations $\mathbf{f}(s_i)$ and $\mathbf{f}(t_i)$. Then, we *adjust* the mBERT model on aligned pairs’ vectors using the following loss function:

$$L = \sum_{(s_i, t_i)} \|\mathbf{f}(s_i) - \mathbf{f}(t_i)\|_2^2 + \sum_{s_j} \|\mathbf{f}(s_j) - \mathbf{f}^0(s_j)\|_2^2, \quad (1)$$

²<https://huggingface.co/bert-base-multilingual-cased>

³<https://github.com/google-research/bert/blob/master/multilingual.md>

⁴We also experimented with other options reported in the literature – first/last tokens’ vectors, as well as aligning subword tokens produced by BERT. Although these choices induced some variations in results, there is no single pattern across all tasks and languages.

where the first term “pulls” the embeddings in the source and target language together, while the second (regularization) term prevents source (English) representations from deviating far from their initial values in the ‘original’ mBERT f^0 . Finally, the cross-lingually adjusted mBERT model is fine-tuned for a specific task.

Training neural networks via empirical loss minimization is known to suffer from the “catastrophic forgetting” (McCloskey and Cohen, 1989). The histograms of L_2 distances between embeddings of related and unrelated words in pairs of languages (see Fig. 1 and the discussion in § 5.3), confirm that this is, indeed, the case. Specifically, fine-tuning on a target task—in contrast to the cross-lingual adjustment objective—reduces the separation between related and unrelated words. To counter this effect, we ran an additional experiments in a continual-learning mode (Riabi et al., 2021), which relies on experience replay (Ratcliff, 1990; Robins, 1995).

Technically, this entailed a multi-task training (Caruana, 1996) with a combined loss function:

$$L = L_{target} + \alpha L_{align}, \quad (2)$$

where L_{target} was the loss-function for the target task, e.g., NLI, L_{align} was a cross-lingual loss function given by Eq. 1, and $\alpha > 0$ was a small weight. During training, we iterated over the complete (reshuffled) dataset for the target task: After computing L_{target} for a current batch we randomly sampled a small batch of aligned pairs of words $\{(s_i, t_i)\}$ from the parallel corpus and computed L_{align} .

4 Tasks and Data

4.1 Languages and Parallel Data

In our experiments we transfer models trained on English to four languages: Spanish, Russian, Vietnamese, and Hindi. This set represents four different families (including one non-Indo-European language), three scripts, and two different prevalent word orders (see Table 1). All the languages are among languages that were used to train mBERT.⁵

To align embeddings, we use a sample from the parallel corpus WikiMatrix (Schwenk et al., 2021) that contains sentences in 1,620 different language pairs mined from Wikipedia.

Lang	Family	Script	Word order	Number of English Wiki pages
en	IE/Germanic	Latin	SVO	6.3M
es	IE/Romance	Latin	SVO	1.7M
ru	IE/Slavic	Cyrillic	SVO	1.7M
vi	Austroasiatic	Latin	SVO	1.3M
hi	IE/Indo-Aryan	Devanagari	SOV	150K

IE : Indo-European; Prevalent word order: SVO – subject-verb-object, SOV – subject-object-verb;

Table 1: Language information.

4.2 Natural Language Inference

Natural language inference (NLI) is a task of determining the relation between two ordered sentences (hypothesis and premise) and classifying them into: entailment, contradiction, or “no relation”. English MultiNLI collection consists of 433K multi-genre sentence pairs (Williams et al., 2018). The XNLI dataset complements the MultiNLI training set with newly collected 2.5K development and 5K test English examples (Conneau et al., 2018). They were professionally translated into 15 languages, including all four target languages of the current study. Additionally, for each target language test set, we created a new mixed-language XNLI set by randomly picking either a hypothesis or a premise and replacing it with the original English sentence. Performance on XNLI datasets is evaluated using classification *accuracy*.

4.3 Named Entity Recognition

Named entity recognition (NER) is a task of locating named entities in unstructured text and classifying them into predefined categories such as persons, organizations, locations, etc. In our experiments, we employ the Wikiann NER corpus (Rahimi et al., 2019) that is derived from a larger “silver-standard”

⁵However, Hindi Wikipedia is an order of magnitude smaller compared to other Wikipedias, which may have led to somewhat inferior contextualized embeddings.

collection that was created fully automatically (Pan et al., 2017). Wikiann NER has data for 41 languages, including all languages in the current study. The named entity types include location (LOC), person (PER), and organization (ORG). The English training set contains 20K sentences. Test sets for Spanish, Vietnamese, and Russian have 10K sentences each; for Hindi – 1K sentences. Performance is evaluated using the *token-level micro-averaged F1*.

4.4 Question Answering

Machine reading comprehension (MRC) is a variant of a QA task. Given a question and a text paragraph, the system needs to return a continuous span of paragraph tokens as an answer. The first large-scale MRC dataset is the English Wikipedia-based dataset SQuAD (Rajpurkar et al., 2016), which contains about 100K paragraph-question-answer triples. SQuAD has become a *de facto* standard and inspired creation of analogous resources in other languages (Rogers et al., 2021). We use SQuAD as the source dataset to train MRC models. To test the models, we use XQuAD, MLQA, and TyDi QA datasets. XQuAD (Artetxe et al., 2020) is a professional translation of 240 SQuAD paragraphs and 1,190 questions-answer pairs into 10 languages (including four languages of our study). MLQA (Lewis et al., 2020) data is available for six languages including Spanish, Vietnamese, and Hindi (but it does not have Russian). There are about 5K questions for each of our languages. TyDi QA (Clark and others, 2020) includes 11 typologically diverse languages of which we use only Russian (812 test items).

In addition to monolingual test data, we experimented with parallel/cross-lingual MLQA and XQuAD datasets and explored two directions: (1) question is in a target language, but paragraph is in English; (2) a question is in English, but a paragraph is in a target language.

QA performance is evaluated using a *token-level F1-score*.

4.5 Cross-Lingual Sentence Retrieval

The task of cross-lingual sentence retrieval (XSR) consists in retrieving sentences that are translations of each other. A query is a sentence in one language and a corpus is a set of sentences (translations) in another language (in our case English). For the XSR task, we use a subset of the Tatoeba collection (Artetxe and Schwenk, 2019) covering 93 languages (including four languages in our study). This subset has one thousand English sentences each of which has a translation in each of the 93 languages. Following Ruder et al. (2021), we fine-tune the model on a QA task. A sentence representation is obtained by averaging token embeddings in one of the layers. Representations are compared using the cosine similarity. Because different layers perform differently in different scenarios (original mBERT vs. adjusted vs. adjusted with continual learning), we select the scenario-specific best-performing layer. In that, retrieval performance is measured using *the mean reciprocal rank (MRR)*.

5 Experimental Results and Analysis

5.1 Setup

All experiments were conducted on a single Tesla V100 16GB. For cross-lingual model adjustment we use the Adam optimizer and hyper-parameters provided by Cao et al. (2020). To obtain reliable results we run five iterations (using different seeds) of model adjustment (for each configuration) followed by fine-tuning on downstream tasks. For each run we sample 30K sentences from a set of 250K parallel (WikiMatrix) sentences word-aligned with *fast align*.⁶

The code to fine-tune mBERT on XNLI, SQuAD, and Wikiann is based on HuggingFace sample scripts,⁷ which were modified to support continual learning. We use a standard architecture consisting of a BERT model with a task-specific linear layer (Devlin et al., 2019). We also reuse parameters provided by HuggingFace, except for the weight $\alpha = 0.01$ in the multi-task loss (Eq. 2), which was tuned on

⁶We ran the main experiments with 30K parallel sentences. In addition, we conducted experiments with 5K/10K/30k/100K/250K Ru-En sentence pairs. Increasing amount of parallel data benefits both NLI and NER, whereas QA performance peaks at roughly 5K parallel sentences and further decreases as the number of parallel sentences increases. Due to limited space we don't report detailed results and analysis here.

⁷<https://github.com/huggingface/transformers/tree/master/examples/pytorch>

mBERT	es	ru	vi	hi
Original XNLI				
Original	74.20	67.95	69.58	59.03
Adjusted	74.82*	69.45*	70.88*	61.54*
Adj+cont	75.89**	71.26**	72.79**	63.90**
Mixed-language NLI				
Original	70.93	64.24	62.72	53.53
Adjusted	72.06*	66.56*	66.50*	57.31*
Adj+cont	73.50**	69.09**	69.14**	61.09**

Statistically significant differences from an original and adjusted mBERT are marked with * and **, respectively (p-value threshold 0.05).

Table 2: Performance on original and mixed-language NLI datasets (accuracy).

mBERT	es	ru	vi	hi
Original	73.40	63.43	71.02	65.24
Adjusted	73.28	65.49*	71.99*	68.22*
Adj+cont	72.71**	66.27**	71.35**	66.07**

Statistically significant differences from an original and adjusted mBERT are marked with * and **, respectively (p-value threshold 0.05).

Table 3: Performance on NER task (token-level F1).

mBERT	Spanish		Russian		Vietnamese		Hindi	
	MLQA	XQuAD	TyDi QA	XQuAD	MLQA	XQuAD	MLQA	XQuAD
Original	64.96	75.59	67.05	70.72	59.95	69.18	48.73	57.56
Adjusted	63.11*	73.99*	67.03	70.58	58.46*	68.63	48.47	57.81
Adj+cont	62.76**	73.44	67.63	70.51	57.71**	68.64	48.02**	57.83
Question in target language, paragraph in English								
Original	67.34	75.74	—	71.54	56.08	65.00	42.48	47.83
Adjusted	66.93*	75.65	—	71.68	56.74*	66.75*	44.91*	50.45*
Adj+cont	66.31**	74.88**	—	70.99	54.51**	64.63**	43.88**	50.13
Question in English, paragraph in target language								
Original	67.36	76.71	—	67.31	64.43	68.12	55.32	58.62
Adjusted	66.96*	76.42	—	68.25*	65.01*	68.99	55.63	58.93
Adj+cont	66.68**	76.21	—	68.06	64.36**	68.54	54.74**	58.22

Statistically significant differences from an original and adjusted mBERT are marked with * and **, respectively (p-value threshold 0.05).

Table 4: Effectiveness of QA systems (F1-score).

XNLI validation sets. Also note that batch sizes are 32 (for the main target loss) and 16 (for the auxiliary cross-lingual adjustment loss in the case of continual learning).

All reported results are averages over five runs with different seeds. We compute statistical significance of differences between the original and adjusted mBERT using a paired t-test. For XSR, QA, and NLI we first average metric values for each example over different runs and then carry out a paired t-test using averaged values. For NER we concatenate example-specific predictions for all seeds and run 1,000 iterations of a permutation test for concatenated sequences (Pitman, 1937; Efron and Tibshirani, 1993).

5.2 Main Results

Results for NLI, NER, QA, and XSR tasks are summarized in Tables 2, 3, 4, and 5, respectively. We can observe consistent and statistically significant improvements (up to 2.5 accuracy point) of aligned models over zero-shot transfer on XNLI for all languages. This is in line with Cao et al. (2020) even though we used a set of more diverse languages, parallel data of lower quality, and a slightly different learning scheme (we adjusted models individually for each pair of languages). In that, employing continual learning lead to additional substantial gains (up to 2.4 accuracy points).

We also evaluated models on the (bilingual) mixed-language XNLI test data (see § 4.2). According to the bottom part of Table 2, compared to the original XNLI, we observe bigger gains for all four languages, especially when we employ continual learning. For Hindi, we obtain a 7.5 point gain by using both the adjustment and continual learning.

NER results are somewhat mixed: We observe statistically significant gains (up to 3 points for Hindi) on all languages except Spanish. In that, continual learning is beneficial only for Russian.

When we fine-tuned a cross-lingually adjusted mBERT on *mono-lingual* QA tasks, there were no statistically significant gains. In that, there was a statistically significant decrease for all Spanish datasets and Vietnamese MLQA. Use of continual learning lead to further degradation in nearly all cases. Note

that models were noticeably less accurate on MLQA compared to XQuAD, which is a translation of our training set SQuAD.

Muttenthaler et al. (2020) and van Aken et al. (2019) showed that QA models essentially clustered answer token vectors and separated them from the rest of the paragraph token vectors using a vector representation of the question. Thus, to solve the QA task, the model learns to rely on *mutual similarities* among question and answer tokens (on English QA data) rather than on their actual vector representations. As a consequence, there may be no need to make representations in the target language to be similar to English-language representations. This may *partially* explain why the cross-lingual adjustment was unsuccessful for *mono-lingual* QA.

This hypothesis—together with the observation of stronger performance on XQuAD compared to MLQA—prompted us to explore whether the cross-lingual adjustment could be more useful for *cross-lingual* QA. We explored two directions: (1) question is in a target language, but paragraph is in English; (2) a question is in English, but a paragraph is in a target language. According to results in the lower part of Table 4, we observe improvements in three languages (except Spanish) and most of these improvements are statistically significant. This is in line with cross-lingual QA results by Faisal and Anastasopoulos (2021).

For XSR—another cross-lingual task—using the cross-lingual adjustment lead to substantial and statistically significant improvements in three languages: Spanish, Vietnamese, and Hindi. However, continual learning only marginally helped Russian. Finally, we note that cross-lingual adjustment resulted in bigger gains for the mixed, i.e., cross-lingual XNLI data, than for original one (Table 2). In summary, we believe our results support a conjecture that the cross-lingual adjustment of Cao et al. (2020) is more beneficial for cross-lingual tasks.

mBERT	es	ru	vi	hi
Original	0.80	0.73	0.67	0.49
Adjusted	0.82*	0.72*	0.72*	0.56*
Adj+cont	0.78**	0.74**	0.69**	0.53**

Statistically significant differences from an original and adjusted mBERT are marked with * and **, respectively (p-value threshold 0.05).

Table 5: Performance on XSR task (MRR for best-performing layers).

5.3 Analysis of the Adjusted mBERT for NLI

This section analysis focuses on the NLI task. We calculate L_2 distances between contextualized embeddings in English and other languages.⁸ The embeddings are taken from the last layer output (i.e., no prediction heads are used). To this end we sampled semantically related words from parallel sentences (matched via *fastalign*) and unrelated words from unpaired sentences (nearly always unrelated). For each pair of languages and each task, the sampling process was carried out for: (1) the original mBERT, (2) an adjusted mBERT, (3) the original mBERT fine-tuned for the target task, (4) the adjusted mBERT fine-tuned for the target task, (5) the adjusted mBERT fine-tuned for the target task using *continual* learning.⁹

From Fig. 1 we can see that the cross-lingual adjustment makes embeddings of semantically similar words from different languages closer to each other while keeping unrelated words apart, which is in line with Zhao et al. (2021). However, prior work did not inspect histograms obtained after fine-tuning. Yet, quite surprisingly, fine-tuning of both the original and adjusted mBERT on the English NLI data (Fig. 1c and 1d) makes distributions of related and unrelated words almost fully overlap, i.e. all embeddings become close to each other. Compared to the original mBERT, fine-tuning of the *adjusted* mBERT (Fig. 1d) does result in a better separation of related and unrelated words, but the effect is *quite modest*. We believe this is an example of “catastrophic forgetting” (McCloskey and Cohen, 1989), where fine-tuning the model on a target task causes the model to forget some of the knowledge obtained during cross-lingual adjustment.

Continual learning (Fig. 1e) permits fine-tuning for the target task while maintaining a separation between related and unrelated words, which also consistently improves performance for the NLI task.

⁸Although most prior work uses the cosine similarity instead of L_2 (Rudman et al., 2022), it does not distinguish between vectors with the same direction, but different lengths.

⁹Due to limited space we demonstrate histograms only for Hindi/NLI task, other languages/tasks exhibit similar trends.

However, when we compared histograms for additional tasks (not shown due to space limit) there was no direct relationship between the degree of separation and the success of cross-lingual transfer among all tasks and training regimes. In the case of NER, the biggest separation was achieved for Spanish, but fine-tuning of the adjusted mBERT resulted in a lower accuracy. More generally, fine-tuning with continual learning *always* led to a better separation of related and unrelated words, but this extra separation was beneficial only for the NLI task.

6 Conclusion

We evaluate effectiveness of an existing approach to cross-lingual adjustment of mBERT (Cao et al., 2020) using four typologically different languages (Spanish, Russian, Vietnamese, and Hindi) and four IR/NLP tasks (XSR, QA, NLI, and NER). The original mBERT is being compared to mBERT “adjusted” with a help of a small parallel corpus. The cross-lingual adjustment of mBERT improves NLI in four target languages; NER, XSR, and *cross-lingual* QA in three languages (though some cross-lingual QA gains are not statistically significant). However, in the case of *mono-lingual* QA performance never improves and sometimes degrades. We believe our results support a conjecture that the cross-lingual adjustment of Cao et al. (2020) is more beneficial for cross-lingual tasks.

Inspired by the analysis of histograms of distances, we obtain *additional* improvement on NLI using continual learning. Our study contributes to a better understanding of cross-lingual transfer capabilities of large multilingual language models and identifies limitations of the cross-lingual adjustment in various IR and NLP tasks.

Acknowledgments

This research was supported in part through computational resources of HPC facilities at HSE University (Kostenetskiy et al., 2021). PE is grateful to Yandex Cloud for their grant toward computing resources of Yandex DataSphere.¹⁰ PB acknowledges support by the Russian Science Foundation, grant № 20-11-20166.

References

- Hanan Aldarmaki and Mona Diab. 2019. Context-aware cross-lingual mapping. 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 3906–3911.
- 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.
- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. On the cross-lingual transferability of monolingual representations. In *ACL*, pages 4623–4637.
- Ricardo Baeza-Yates, Berthier Ribeiro-Neto, et al. 1999. *Modern information retrieval*, volume 463. ACM press New York.
- Krisztian Balog. 2018. *Entity-Oriented Search*, volume 39 of *The Information Retrieval Series*. Springer.
- Alex Brandsen, Suzan Verberne, Karsten Lambers, and Milco Wansleeben. 2022. Can BERT dig it? named entity recognition for information retrieval in the archaeology domain. *J. Comput. Cult. Herit.*, 15(3).
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. Multilingual alignment of contextual word representations. In *8th International Conference on Learning Representations, ICLR 2020*.
- Rich Caruana. 1996. Algorithms and applications for multitask learning. In *ICML*, pages 87–95. Morgan Kaufmann.
- Xiang Cheng, Mitchell Bowden, Bhushan Ramesh Bhange, Priyanka Goyal, Thomas Packer, and Faizan Javed. 2021. An end-to-end solution for named entity recognition in ecommerce search. In *AAAI*, pages 15098–15106. AAAI Press.

¹⁰<https://cloud.yandex.com/en/services/datasphere>

- Ethan A Chi, John Hewitt, and Christopher D Manning. 2020. Finding universal grammatical relations in multilingual bert. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5564–5577.
- Jonathan H Clark et al. 2020. TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages. *TACL*, 8:454–470.
- 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.
- Alexis Conneau, Shijie Wu, Haoran Li, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Emerging cross-lingual structure in pretrained language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6022–6034.
- 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, June.
- Sumanth Doddapaneni, Gowtham Ramesh, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh M Khapra. 2021. A primer on pretrained multilingual language models. *arXiv preprint arXiv:2107.00676*.
- Philipp Dufter and Hinrich Schütze. 2020. Identifying elements essential for bert’s multilinguality. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4423–4437.
- 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.
- Abteen Ebrahimi, Manuel Mager, Arturo Oncevay, Vishrav Chaudhary, Luis Chiruzzo, Angela Fan, John Ortega, Ricardo Ramos, Annette Rios, et al. 2021. AmericasNLI: Evaluating zero-shot natural language understanding of pretrained multilingual models in truly low-resource languages. *arXiv preprint arXiv:2104.08726*.
- Bradley Efron and Robert J Tibshirani. 1993. An introduction to the bootstrap. *Monographs on statistics and applied probability*, 57(1).
- Faezeh Ensan and Ebrahim Bagheri. 2017. Document retrieval model through semantic linking. In *WSDM*, pages 181–190. ACM.
- Fahim Faisal and Antonios Anastasopoulos. 2021. Investigating post-pretraining representation alignment for cross-lingual question answering. In *Proceedings of the 3rd Workshop on Machine Reading for Question Answering*, pages 133–148.
- Tsung-Yuan Hsu, Chi-Liang Liu, and Hung-yi Lee. 2019. Zero-shot reading comprehension by cross-lingual transfer learning with multi-lingual language representation model. 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 5933–5940, November.
- 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 generalisation. In *International Conference on Machine Learning*, pages 4411–4421.
- Karthikeyan K, Zihan Wang, Stephen Mayhew, and Dan Roth. 2020. Cross-lingual ability of multilingual BERT: an empirical study. In *ICLR*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, pages 79–86.
- P. S. Kostenetskiy, R. A. Chulkevich, and V. I. Kozyrev. 2021. Hpc resources of the higher school of economics. *Journal of Physics: Conference Series*, 1740(1):012050, jan.
- Saurabh Kulshreshtha, Jose Luis Redondo Garcia, and Ching-Yun Chang. 2020. Cross-lingual alignment methods for multilingual BERT: A comparative study. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 933–942, November.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *International Conference on Learning Representations*.

- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499.
- Patrick Lewis, Barlas Oğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. MLQA: Evaluating cross-lingual extractive question answering. In *ACL*, pages 7315–7330.
- Jindřich Libovický, Rudolf Rosa, and Alexander Fraser. 2019. How language-neutral is multilingual bert? *arXiv preprint arXiv:1911.03310*.
- 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, July.
- Robert Litschko, Ivan Vulić, Simone Paolo Ponzetto, and Goran Glavaš. 2021. Evaluating multilingual text encoders for unsupervised cross-lingual retrieval. *arXiv preprint arXiv:2101.08370*.
- Zihan Liu, Genta Indra Winata, Andrea Madotto, and Pascale Fung. 2021. Preserving cross-linguality of pre-trained models via continual learning. In *RepL4NLP@ACL-IJCNLP*, pages 64–71. Association for Computational Linguistics.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to information retrieval*. Cambridge University Press.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Benjamin Muller, Antonios Anastasopoulos, Benoît Sagot, and Djamé Seddah. 2021. When being unseen from mBERT is just the beginning: Handling new languages with multilingual language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 448–462.
- Lukas Muttenthaler, Isabelle Augenstein, and Johannes Bjerva. 2020. Unsupervised evaluation for question answering with transformers. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 83–90.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. Cross-lingual name tagging and linking for 282 languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958.
- German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237.
- 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.
- Edwin JG Pitman. 1937. Significance tests which may be applied to samples from any populations. *Supplement to the Journal of the Royal Statistical Society*, 4(1):119–130.
- Afshin Rahimi, Yuan Li, and Trevor Cohn. 2019. Massively multilingual transfer for ner. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 151–164.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Roger Ratcliff. 1990. Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review*, 97(2):285.

- Arij Riabi, Thomas Scialom, Rachel Keraron, Benoît Sagot, Djamé Seddah, and Jacopo Staiano. 2021. Synthetic data augmentation for zero-shot cross-lingual question answering. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7016–7030.
- Anthony Robins. 1995. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146.
- Anna Rogers, Matt Gardner, and Isabelle Augenstein. 2021. QA dataset explosion: A taxonomy of nlp resources for question answering and reading comprehension. *arXiv preprint arXiv:2107.12708*.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2019. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*, 65:569–631.
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Dan Garrette, Graham Neubig, and Melvin Johnson. 2021. XTREME-R: Towards more challenging and nuanced multilingual evaluation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10215–10245.
- William Rudman, Nate Gillman, Taylor Rayne, and Carsten Eickhoff. 2022. Isoscore: Measuring the uniformity of embedding space utilization. In *ACL (Findings)*, pages 3325–3339. Association for Computational Linguistics.
- Tal Schuster, Ori Ram, Regina Barzilay, and Amir Globerson. 2019. Cross-lingual alignment of contextual word embeddings, with applications to zero-shot dependency parsing. 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 1599–1613.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361.
- Dahlia Shehata, Negar Arabzadeh, and Charles L. A. Clarke. 2022. Early stage sparse retrieval with entity linking. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22*, page 4464–4469.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. On the limitations of unsupervised bilingual dictionary induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788.
- Noha S. Tawfik. 2020. *Text Mining for Precision Medicine: Natural Language Processing, Machine Learning and Information Extraction for Knowledge Discovery in the Health Domain*. Ph.D. thesis, Utrecht University, Netherlands.
- Iulia Turc, Kenton Lee, Jacob Eisenstein, Ming-Wei Chang, and Kristina Toutanova. 2021. Revisiting the primacy of english in zero-shot cross-lingual transfer. *arXiv preprint arXiv:2106.16171*.
- Betty van Aken, Benjamin Winter, Alexander Löser, and Felix A Gers. 2019. How does BERT answer questions? A layer-wise analysis of transformer representations. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 1823–1832.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Yuxuan Wang, Wanxiang Che, Jiang Guo, Yijia Liu, and Ting Liu. 2019a. 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 5721–5727, Hong Kong, China.
- Zirui Wang, Jiateng Xie, Ruochen Xu, Yiming Yang, Graham Neubig, and Jaime Carbonell. 2019b. Cross-lingual alignment vs joint training: A comparative study and a simple unified framework. *arXiv preprint arXiv:1910.04708*.
- Musen Wen, Deepak Kumar Vasthimal, Alan Lu, Tian Wang, and Aimin Guo. 2019. Building large-scale deep learning system for entity recognition in e-commerce search. In *BDCAT*, pages 149–154. ACM.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122.

- 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.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. 2021. Inducing language-agnostic multilingual representations. In *Proceedings of *SEM 2021: The Tenth Joint Conference on Lexical and Computational Semantics*, pages 229–240.