

Unsupervised Multilingual Sentence Embeddings for Parallel Corpus Mining

Ivana Kvapilíková[‡] Mikel Artetxe[§] Gorka Labaka[§]
Eneko Agirre[§] Ondřej Bojar[‡]

[‡]Institute of Formal and Applied Linguistics, Charles University (MFF UK)

{kvapilikova,bojar}@ufal.mff.cuni.cz

[§]Ixa NLP group, University of the Basque Country (UPV/EHU)

{mikel.artetxe,gorka.labaka,e.agirre}@ehu.eus

Abstract

Existing models of multilingual sentence embeddings require large parallel data resources which are not available for low-resource languages. We propose a novel unsupervised method to derive multilingual sentence embeddings relying only on monolingual data. We first produce a synthetic parallel corpus using unsupervised machine translation, and use it to fine-tune a pretrained cross-lingual masked language model (XLM) to derive the multilingual sentence representations. The quality of the representations is evaluated on two parallel corpus mining tasks with improvements of up to 22 F1 points over vanilla XLM. In addition, we observe that a single synthetic bilingual corpus is able to improve results for other language pairs.

1 Introduction

Parallel corpora constitute an essential training data resource for machine translation as well as other cross-lingual NLP tasks. However, large parallel corpora are only available for a handful of language pairs while the rest relies on semi-supervised or unsupervised methods for training. Since monolingual data are generally more abundant, parallel sentence mining from non-parallel corpora provides another opportunity for low-resource language pairs.

An effective approach to parallel data mining is based on multilingual sentence embeddings (Schwenk, 2018; Artetxe and Schwenk, 2019b). However, existing methods to generate cross-lingual representations are either heavily supervised or only apply to static word embeddings. An alternative approach to unsupervised multilingual training is that of Devlin et al. (2018) or Lample and Conneau (2019), who train a masked language model (M-BERT, XLM) on a concatenation of monolingual corpora in different languages to

learn a joint structure of these languages together. While several authors (Pires et al., 2019; Wu and Dredze, 2019; Karthikeyan et al., 2019; Libovický et al., 2019) bring evidence of cross-lingual transfer within the model, its internal representations are not entirely language agnostic.

We propose a method to further align representations from such models into the cross-lingual space and use them to derive sentence embeddings. Our approach is completely unsupervised and is applicable even for very distant language pairs. The proposed method outperforms previous unsupervised approaches on the BUCC 2018¹ shared task, and is even competitive with several supervised baselines.

The paper is organized as follows. Section 2 gives an overview of related work; Section 3 introduces the proposed method; Section 4 describes the experiments and reports the results. Section 5 concludes.

2 Related Work

Related research comprises supervised methods to model multilingual sentence embeddings and unsupervised methods to model multilingual word embeddings which can be aggregated into sentences. Furthermore, our approach is closely related to the recent research in cross-lingual language model (LM) pretraining.

Supervised multilingual sentence embeddings. The state-of-the-art performance in parallel data mining is achieved by LASER (Artetxe and Schwenk, 2019b) – a multilingual BiLSTM model sharing a single encoder for 93 languages trained on parallel corpora to produce language agnostic sentence representations. Similarly, Schwenk and Douze (2017); Schwenk (2018); Espana-Bonet

¹11th Workshop on Building and Using Comparable Corpora

et al. (2017) derive sentence embeddings from internal representations of a neural machine translation system with a shared encoder. The universal sentence encoder (USE) (Cer et al., 2018; Yang et al., 2019) family covers sentence embedding models with a multi-task dual-encoder training framework including the tasks of question-answer prediction or natural language inference. Guo et al. (2018) directly optimize the cosine similarity between the source and target sentences using a bidirectional dual-encoder. These approaches rely on heavy supervision by parallel corpora which is not available for low-resource languages.

Unsupervised multilingual word embeddings. Cross-lingual embeddings of words can be obtained by post-hoc alignment of monolingual word embeddings (Mikolov et al., 2013) and mean-pooled with IDF weights to represent sentences (Litschko et al., 2019). Unsupervised techniques to find a linear mapping between embedding spaces were proposed by Artetxe et al. (2018) and Conneau et al. (2018), using iterative self-learning or adversarial training. Several recent studies (Patra et al., 2019; Ormazabal et al., 2019) criticize this simplified approach, showing that even the embedding spaces of closely related languages are not isometric. Vulić et al. (2019) question the robustness of unsupervised mapping methods in challenging circumstances.

Cross-lingual LM pretraining. Ma et al. (2019); Reimers and Gurevych (2019) derive monolingual sentence embeddings by mean-pooling contextualized word embeddings from BERT. Schuster et al. (2019); Wang et al. (2019b) propose mapping such contextualized embeddings into the multilingual space and report favorable results on the task of dependency parsing. Pires et al. (2019) extract contextualized embeddings directly from unsupervised multilingual LMs and use them for parallel sentence retrieval. Other authors improve the alignment of representations in a multilingual LM using a parallel corpus as an anchor (Cao et al., 2020) or using iterative self-learning (Wang et al., 2019a). None of these works apply multilingual embeddings to mine parallel sentences. Our work is the first in improving unsupervised cross-lingual models using additional unsupervised information.

3 Proposed Method

We propose a method to enhance the cross-lingual ability of a pretrained multilingual model by fine-

tuning it on a small synthetic parallel corpus. The parallel corpus is obtained via unsupervised machine translation (MT) so the method remains unsupervised. In this section, we describe the pretrained model (Section 3.1), the fine-tuning objective (Section 3.2) and the extraction of sentence embeddings (Section 3.3). We provide details on the unsupervised MT system in Section 3.4.

3.1 XLM Pretraining

The starting point for our experiments is a cross-lingual language model (XLM) (Lample and Conneau, 2019) of the BERT family pretrained on concatenated monolingual texts in 100 languages using the masked language model (MLM) training objective (Devlin et al., 2018). The model processes the input in BPE subword units (Sennrich et al., 2016) with a shared vocabulary for all languages. In this work, we use the publicly available pretrained model XLM-100² (Lample and Conneau, 2019) with 16 transformer layers, 16 attention heads and a hidden unit size of 1280. The model was trained on monolingual corpora in 100 languages with the BPE vocabulary of 240k subwords.

3.2 XLM Fine-tuning with a Translation Objective

When parallel data is available, it can be leveraged in training of the multilingual language model using a translation language model loss (TLM) (Lample and Conneau, 2019). Pairs of sentences are concatenated, random tokens are masked from both sentences and the model is trained to fill in the blanks by attending to any of the words of the two sentences. The Transformer self-attention layers thus have the capacity to enrich word representations with the information about their monolingual context as well as their translation counterparts. This explicit cross-lingual training objective further enhances the alignment of the embeddings in the cross-lingual space.

We use this objective to fine-tune the pretrained model on a small synthetic parallel data set obtained via unsupervised MT for one language pair, aiming to improve the overall cross-lingual alignment of the internal representations of the model. In our experiments, we also compare the performance to fine-tuning on small authentic parallel corpora.

²<https://github.com/facebookresearch/XLM>

3.3 Sentence Embeddings

Pretrained language models produce contextual representations capturing the semantic and syntactic properties of word (subword) tokens in their variable context (Devlin et al., 2018). Contextualized embeddings can be derived from any of the internal layer outputs of the model. We tune the choice of the layer on the task of parallel sentence matching and conclude that the best cross-lingual performance is achieved at the 12th (5th-to-last) layer. Therefore, we use the representations from this layer in the rest of this paper. The evaluation across layers is summarized in Figure 1 in Section 4.6.

Aggregating subword embeddings to fixed-length sentence representations necessarily leads to an information loss. We compose sentence embeddings from subword representations by simple element-wise averaging. Even though mean-pooling is a naive approach to subword aggregation, it is often used for its simplicity (Reimers and Gurevych, 2019; Ruiters et al., 2019; Ma et al., 2019) and in our scenario it yields better results than max-pooling.

3.4 Unsupervised Machine Translation

Our unsupervised MT model follows the approach of Lample and Conneau (2019). It is a Transformer model with an encoder-decoder architecture. Both the encoder and the decoder are shared across languages and they are initialized with a pretrained bilingual LM to bootstrap the training. Both the encoder and the decoder have 6 layers, 8 attention heads and a hidden unit size of 768. The system is trained using the unsupervised MT training pipeline of denoising and back-translation (Lample et al., 2018).

4 Experiments & Results

In this section, we empirically evaluate the quality of our cross-lingual sentence embeddings and compare it with state-of-the-art supervised methods and unsupervised baselines. We evaluate the proposed method on the task of parallel corpus mining and parallel sentence matching. We fine-tune two different models using English-German and Czech-German synthetic parallel data.

4.1 Data

The XLM model was pretrained on the Wikipedia corpus of 100 languages (Lample and Conneau, 2019). The monolingual data for fine-tuning was

sampled from NewsCrawl 2018 (10k Czech sentences, 10k German sentences, 10k English sentences).

Monolingual training data for the unsupervised MT models was obtained from NewsCrawl 2007-2008 (5M sentences per language). The text was cleaned and tokenized using standard Moses (Koehn et al., 2007) tools and segmented into BPE units based on 60k BPE splits.

4.2 Experiment Details

To generate synthetic data for fine-tuning, we train two unsupervised MT models (Czech-German, English-German) using the same method and parameters as in Lample and Conneau (2019) on 8 GPUs for 24 hours. We use these models to translate 10k sentences in each language. The translations are coupled with the originals into two parallel corpora of 20k synthetic sentence pairs.

The small synthetic parallel corpora obtained in the first step are used to fine-tune the pretrained XLM-100 model using the TLM objective. We measure the quality of induced cross-lingual embeddings from different layers on the task of parallel sentence matching described in Section 4.5 and observe the best results at the 12th layer after fine-tuning for one epoch with a batch size of 8 sentences and all other pretraining parameters intact. The development accuracy decreases with fine-tuning on a larger data set.

4.3 Baselines

We assess our method against two unsupervised baselines to separately measure the fine-tuning effect on the XLM model and to compare our results to another possible unsupervised approach based on post-hoc alignment of word embeddings.

Vanilla XLM. Contextualized token representations are extracted from the 12th layer of the original XLM-100³ model and mean-pooled into sentence embeddings.

Word Mapping. We use Word2Vec embeddings with 300 dimensions pretrained on NewsCrawl and map them into the cross-lingual space using the unsupervised version of VecMap (Artetxe et al., 2018). As above, word embeddings are aggregated by mean-pooling to represent sentences.⁴

³Using M-BERT model yielded similar results to XLM.

⁴Weighting word embeddings by their sentence frequency (IDF) did not lead to a significant improvement over a simple average.

	en-de	en-fr	en-ru	en-zh	Supervision	
Leong et al. (2018)	-	-	-	56.00	bitext	0.5M sent.
Bouamor and Sajjad (2018)	-	76.00	-	-	bitext	2M sent.
Schwenk (2018)	76.90	75.80	73.80	71.60	9-way parallel	2M sent.
Azpeitia et al. (2018)	85.52	81.47	81.30	77.45	bitext	2-9M sent.
Artetxe and Schwenk (2019b)	96.19	93.91	93.30	92.27	2- or 3-way parallel	223M sent.
Unsup. baseline (Word Mapping)	32.04	32.94	17.68	20.65	none	n/a
Unsup. baseline (Vanilla XLM)*	62.10	64.77	61.65	44.79	none	n/a
Proposed method* (en↔de)	80.06	78.77	77.16	67.04	none	20k sent.**

Table 1: F1 score on the parallel sentence mining task (BUCC test set). The supervised (upper part) and unsupervised (lower part) winners are highlighted in bold. * The model was pretrained on Wikipedia. ** Synthetic translations produced by unsupervised MT.

	en-de	en-fr	en-ru	en-zh	en-kk	cs-zh	de-ru
Artetxe and Schwenk (2019b)	90.30	87.38	94.34	83.92	12.07	73.41	88.39
Unsup. baseline (Word Mapping)	28.45	30.79	17.81	16.04	2.28	10.86	19.55
Unsup. baseline (Vanilla XLM)	72.58	71.92	72.90	59.26	24.00	43.00	58.29
Proposed method (en↔de)	79.32	77.05	80.98	65.49	35.41	48.79	65.91

Table 2: F1 score on the parallel sentence mining task (News test set). The supervised and unsupervised winners are highlighted in bold. Artetxe and Schwenk (2019b) values obtained using the public implementation of the LASER toolkit.

4.4 Evaluation I: Parallel Corpus Mining

We measure the performance of our method on the BUCC shared task of parallel corpus mining where the system is expected to search two comparable non-aligned corpora and identify pairs of parallel sentences. We evaluate on two data sets – the original BUCC 2018 corpus created by inserting parallel sentences into monolingual texts extracted from Wikipedia (Zweigenbaum et al., 2017) and a new BUCC-like data set (News train and test) which we created by shuffling 10k parallel sentence from News Commentary into 400k monolingual sentences from News Crawl. The BUCC and News data sets are comparable in size and contain parallel sentences from the same source, but differ in overall domain.

In order to score all candidate sentence pairs, we use the margin-based approach of Artetxe and Schwenk (2019a) which was proved to eliminate the hubness problem of embedding spaces and yield superior results (Artetxe and Schwenk, 2019b). The score relies on cosine similarity to measure the distance between sentences but it is defined in relative terms to the average cosine similarity between the two sentences and their nearest neighbors. The optimal threshold for filtering the translation pairs is learned by tuning on the train set F1 scores. Tables 1 and 2 show the results of our proposed model on the BUCC and News test sets, resp., comparing them to related work and unsupervised baselines.

When comparing our method to related work, it must be noted that the XLM model was pretrained on Wikipedia and therefore has seen the monolingual BUCC sentences during training. This could result in an advantage over other systems, as the model could exploit the fact that it has seen the non-parallel part of the comparable corpus during training. However, since both the proposed method and the *vanilla XLM* baseline suffer from this, their results remain comparable. We also report results on the News test set which is free from such potential bias (Table 2).

The results reveal that TLM fine-tuning brings a substantial improvement over the initial pretrained model trained only using the MLM objective (*vanilla XLM*). In terms of the F1 score, the gain across four BUCC language pairs is 14.0-22.3 points. Even though the fine-tuning focused on a single language pair (English-German), the improvement is notable for all evaluated language pairs. The largest margin of 21.6 points is observed for the English-Chinese mining task. We observe that using a small parallel data set of authentic translation pairs instead of synthetic ones does not have a significant effect.

The weak results of the *word mapping* baseline can be partially attributed to the superiority of contextualized embeddings for representation of sentences over static ones. Furthermore, word mapping relies on the questionable assumption of isomorphic embedding spaces which weakens its performance especially for distant languages. In

	de-en	cs-en	cs-de	cs-fr	cs-ru	fr-es	fr-ru	es-ru
Artetxe and Schwenk (2019b)	98.78	99.08	99.23	99.37	98.77	99.42	98.60	98.77
Unsup. baseline (Word Mapping)	60.60	55.03	75.35	43.33	79.87	71.07	41.25	53.87
Unsup. baseline (Vanilla XLM)	87.15	79.83	82.87	80.55	85.15	91.07	85.28	85.73
Proposed method (en↔de)	93.97	90.47	90.48	90.07	92.23	94.68	91.80	91.92
Proposed method (cs↔de)	94.43	90.15	90.50	89.48	92.33	94.65	91.72	91.25

Table 3: Accuracy on a parallel sentence matching task (*newstest2012*) averaged over both matching directions.

our proposed model, it is possible that joint training of contextualized representations induces an embedding space with more convenient geometric properties which makes it more robust to language diversity.

Although the performance of our model generally lags far behind the supervised LASER benchmark, it is valuable because of its fully unsupervised nature and it works even for distant languages such as Chinese-Czech or English-Kazakh.

4.5 Evaluation II: Parallel Sentence Matching

To assess the effect of proposed fine-tuning on other language pairs not covered by BUCC, we evaluate our embeddings on the task of parallel sentence matching (PSM). The task entails searching a pool of shuffled parallel sentences to recover correct translation pairs. Cosine similarity is used for the nearest neighbor search.

We first evaluate the pairwise matching accuracy on a *newstest* multi-way parallel data set of 3k sentences in 6 languages.⁵ We use *newstest2012* for development and *newstest2013* for testing. The results in Table 3 show that the fine-tuned model is able to match correct translations in 90-95% of cases, depending on the language pair, which is $\sim 7\%$ more than *vanilla XLM*. It is notable that the model which was only fine-tuned on English-German synthetic parallel data has a positive effect on completely unrelated language pairs as well (e.g. Russian-Spanish, Czech-French).

Since the greatest appeal of parallel corpus mining is to enhance the resources for low-resource languages, we also measure the PSM accuracy on the Tatoeba (Artetxe and Schwenk, 2019b) data set of 0.5–1k sentences in over 100 languages aligned with English. Aside from the two completely unsupervised models, we fine-tune two more models on small authentic parallel data in English-Nepali (5k sentence pairs from the Flores development sets) and English-Kazakh (10k sentence pairs from

News Commentary). Table 4 confirms that the improvement over *vanilla XLM* is present for every language we evaluated, regardless on the language pair used for fine-tuning. Although we initially hypothesized that the performance of the English-German model on English-aligned language pairs would exceed the German-Czech model, their results are equal on average. Fine-tuning on small authentic corpora in low-resource languages exceeds both by a slight margin.

The results are clearly sensitive to the amount of monolingual sentences in the Wikipedia corpus used for XLM pretraining and the matching accuracy of very low-resource languages is significantly lower than we observed for high-resource languages. However, the benefits of fine-tuning are substantial (around 20 percentage points) and for some languages the results even reach the supervised baseline (e.g. Kazakh, Georgian, Nepali).

It seems that explicitly aligning one language pair during fine-tuning propagates through the shared parameters and improves the overall representation alignment, making the contextualized embeddings more language agnostic. The propagation effect could also positively influence the ability of cross-lingual transfer within the model in downstream tasks. A verification of this is left to future work.

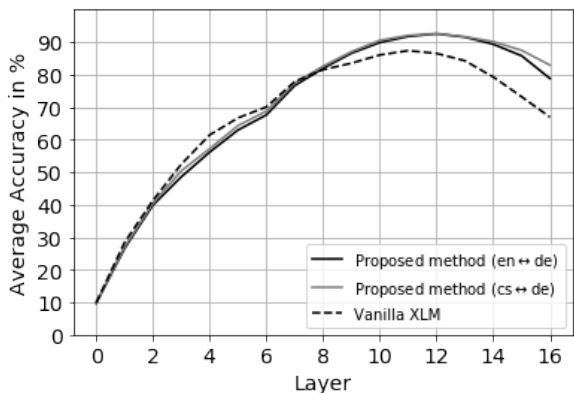


Figure 1: Average PSM accuracy on *newstest2012* before and after fine-tuning from the input embedding layer (0th) to the deepest layer (16th).

⁵Czech, English, French, German, Russian, Spanish

	af	ar	az	be	bg	ca	cs	de	el	eo	et	fi	fy	gl
Sup. baseline	89.5	92.0	66.0	66.2	95.0	95.9	96.5	99.0	95.0	97.2	96.7	96.3	51.7	95.5
Vanilla XLM	38.1	19.9	25.1	33.7	36.2	51.0	31.5	65.0	27.0	45.8	19.8	31.4	37.0	51.4
Proposed method:														
en↔de (synth)	57.3	41.1	46.3	58.4	56.0	66.9	53.5	83.1	51.3	68.0	39.0	47.5	48.6	66.9
cs↔de (synth)	54.2	41.2	44.2	61.8	60.7	68.9	59.9	87.3	53.1	67.4	41.4	49.5	44.8	67.3
en↔kk (auth)	58.4	45.6	51.4	60.2	59.2	72.6	53.9	87.0	54.6	72.1	43.4	51.3	51.7	72.2
en↔ne (auth)	59.9	46.6	54.2	63.1	62.9	71.0	57.6	85.0	51.0	71.2	44.6	52.7	48.6	71.0
	hi	hr	ia	is	id	ja	ka	kk	ku	la	lt	mk	ml	mn
Sup. baseline	94.7	97.2	95.2	95.6	94.5	91.8	35.9	18.6	17.2	58.5	96.2	94.7	96.9	8.2
Vanilla XLM	26.2	47.2	57.3	25.0	46.4	29.5	22.1	17.4	10.6	15.5	22.0	25.8	17.4	12.6
Proposed method:														
en↔de (synth)	53.4	68.2	71.4	43.1	64.9	54.4	41.4	33.6	16.8	24.9	43.9	48.8	51.6	29.0
cs↔de (synth)	51.7	71.8	70.5	43.7	64.1	53.3	39.8	34.7	16.2	27.7	46.2	51.1	44.3	24.5
en↔kk (auth)	60.3	71.3	79.5	45.0	66.4	59.6	44.0	46.1	20.0	28.6	46.2	54.7	54.0	32.7
en↔ne (auth)	59.3	72.1	75.7	47.1	67.8	59.6	47.8	38.4	20.9	30.0	47.7	53.8	56.0	34.9
	mr	ms	ne	nn	oc	sl	sr	sv	ta	te	tl	uk	ur	yi
Sup. baseline	91.5	96.4	20.6	88.3	61.2	95.9	95.3	96.6	69.4	79.7	50.5	94.5	81.9	5.7
Vanilla XLM	15.3	52.0	21.3	49.9	20.0	34.7	35.9	47.2	11.9	14.1	14.6	38.0	19.3	9.9
Proposed method:														
en↔de (synth)	37.3	67.0	32.8	66.8	34.3	54.9	58.6	69.7	40.9	44.7	24.0	66.1	43.7	22.1
cs↔de (synth)	34.2	65.4	31.4	67.5	35.9	59.2	64.8	71.8	31.9	37.8	20.4	70.4	43.8	22.8
en↔kk (auth)	41.9	69.8	37.3	69.2	40.3	58.0	64.3	73.3	42.8	44.0	24.4	71.6	48.2	25.8
en↔ne (auth)	43.5	72.1	42.8	69.2	36.9	58.8	65.0	72.0	41.7	53.2	26.8	71.0	49.9	26.7

Table 4: Accuracy on a parallel sentence matching task (*Tatoeba*) averaged over both matching directions (to and from English). The supervised baseline was obtained using the public implementation of the LASER model (Artetxe and Schwenk, 2019b). Our proposed models were fine-tuned on synthetic parallel data (en↔de, cs↔de) and authentic parallel data (en↔kk, en↔ne).

4.6 Analysis: Representations Across Layers

We derive sentence embeddings from all layers of the model and show PSM results on the development set averaged over all language pairs in Figure 1, both before and after fine-tuning. The accuracy differs substantially across the model depth, the best cross-lingual performance is consistently achieved around the 12th (5th-to-last) layer of the model. The TLM fine-tuning affects especially the deepest layers.

5 Conclusion

We proposed a completely unsupervised method to train multilingual sentence embeddings which can be used for building a parallel corpus with no previous translation knowledge.

We show that fine-tuning an unsupervised multilingual model with a translation objective using as little as 20k synthetic translation pairs can significantly enhance the cross-lingual alignment of its representations. Since the synthetic translations were obtained from an unsupervised MT system, the entire procedure requires no authentic parallel sentences for training.

Our sentence embeddings yield significantly better results on the tasks of parallel data mining and parallel sentence matching than our unsupervised

baselines. Interestingly, targeting only one language pair during the fine-tuning phase suffices to propagate the alignment improvement to unrelated languages. It is therefore not necessary to build a working MT system for every language pair we wish to mine.

The average F1 margin across four language pairs on the BUCC task is ~ 17 points over the original XLM model and ~ 7 on the News dataset where only one of the evaluated language pairs was seen during fine-tuning. The gain in accuracy in parallel sentence matching across 8 language pairs is 7.2% absolute, lagging only 7.1% absolute behind supervised methods.

For the future we would like to apply our model on other cross-lingual NLP tasks such as XNLI or cross-lingual semantic textual similarity.

Acknowledgments

This study was supported in parts by the grants SVV 260 575, 1050119 of the Charles University Grant Agency and 19-26934X of the Czech Science Foundation, by a Facebook Fellowship, the Basque Government excellence research group (IT1343-19), the Spanish MINECO (UnsupMT TIN201791692EXP MCIU/AEI/FEDER, UE) and Project BigKnowledge (Ayudas Fundacin BBVA a equipos de investigacin cientfica 2018).

References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the 56th Annual Meeting of the ACL*, pages 789–798, Melbourne. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019a. [Margin-based parallel corpus mining with multilingual sentence embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019b. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597610.
- Andoni Azpeitia, Thierry Etchegoyhen, and Eva Martinez Garcia. 2018. [Extracting parallel sentences from comparable corpora with stacc variants](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France. European Language Resources Association (ELRA).
- Houda Bouamor and Hassan Sajjad. 2018. [H2@bucc18: Parallel sentence extraction from comparable corpora using multilingual sentence embeddings](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France. European Language Resources Association (ELRA).
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. [Multilingual alignment of contextual word representations](#). In *International Conference on Learning Representations (ICLR 2020)*.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word translation without parallel data](#). In *International Conference on Learning Representations (ICLR 2018)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *arXiv [e-Print archive]*, abs/1810.04805.
- Cristina Espana-Bonet, Adam Csaba Varga, Alberto Barron-Cedeno, and Josef van Genabith. 2017. [An empirical analysis of nmt-derived interlingual embeddings and their use in parallel sentence identification](#). *IEEE Journal of Selected Topics in Signal Processing*, 11(8):13401350.
- Mandy Guo, Qinlan Shen, Yinfei Yang, Heming Ge, Daniel Cer, Gustavo Hernandez Abrego, Keith Stevens, Noah Constant, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. [Effective parallel corpus mining using bilingual sentence embeddings](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 165–176, Brussels, Belgium. Association for Computational Linguistics.
- Kaliyaperumal Karthikeyan, Zihan Wang, Stephen Mayhew, and Dan Roth. 2019. [Cross-lingual ability of multilingual BERT: An empirical study](#). *arXiv [e-Print archive]*, abs/1912.07840.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the ACL*, pages 177–180, Prague. Association for Computational Linguistics.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#). *arXiv [e-Print archive]*, abs/1901.07291.
- Guillaume Lample, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. [Unsupervised machine translation using monolingual corpora only](#). In *International Conference on Learning Representations (ICLR 2018)*.
- Chongman Leong, Derek F. Wong, and Lidia S. Chao. 2018. [Um-palinger: Neural network-based parallel sentence identification model](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Paris, France. European Language Resources Association (ELRA).
- Jindřich Libovický, Rudolf Rosa, and Alexander M. Fraser. 2019. [How language-neutral is multilingual BERT?](#) *arXiv [e-Print archive]*, abs/1911.03310.
- Robert Litschko, Goran Glavaš, Ivan Vulic, and Laura Dietz. 2019. [Evaluating resource-lean cross-lingual embedding models in unsupervised retrieval](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR19*, page 11091112, New York, NY, USA. Association for Computing Machinery.
- Xiaofei Ma, Zhiguo Wang, Patrick Ng, Ramesh Nallapati, and Bing Xiang. 2019. [Universal text representation from bert: An empirical study](#). *arXiv [e-Print archive]*, abs/1910.07973.

- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. [Exploiting similarities among languages for machine translation](#). *arXiv [e-Print archive]*, abs/1309.4168.
- Aitor Ormazabal, Mikel Artetxe, Gorka Labaka, Aitor Soroa, and Eneko Agirre. 2019. [Analyzing the limitations of cross-lingual word embedding mappings](#). In *Proceedings of the 57th Annual Meeting of the ACL*, pages 4990–4995, Florence. Association for Computational Linguistics.
- Barun Patra, Joel Ruben Antony Moniz, Sarthak Garg, Matthew R. Gormley, and Graham Neubig. 2019. [Bilingual lexicon induction with semi-supervision in non-isometric embedding spaces](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 184–193, Florence, Italy. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the ACL*, pages 4996–5001, Florence. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentencebert: Sentence embeddings using siamese bert-networks](#). *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Dana Ruiter, Cristina España-Bonet, and Josef van Genabith. 2019. [Self-supervised neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1828–1834, Florence, Italy. 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](#). *Proceedings of the 2019 Conference of the North American Chapter of the ACL*.
- Holger Schwenk. 2018. [Filtering and mining parallel data in a joint multilingual space](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–234, Melbourne, Australia. Association for Computational Linguistics.
- Holger Schwenk and Matthijs Douze. 2017. [Learning joint multilingual sentence representations with neural machine translation](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 157–167, Vancouver, Canada. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1715–1725, Berlin. Association for Computational Linguistics.
- Ivan Vulić, Goran Glavaš, Roi Reichart, and Anna Korhonen. 2019. [Do we really need fully unsupervised cross-lingual embeddings?](#) 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 4407–4418, Hong Kong, China. Association for Computational Linguistics.
- Shuai Wang, Lei Hou, Juanzi Li, Meihan Tong, and Jiabo Jiang. 2019a. [Learning multilingual sentence embeddings from monolingual corpus](#). In *Chinese Computational Linguistics*, pages 346–357.
- Yuxuan Wang, Wanxiang Che, Jiang Guo, Yijia Liu, and Ting Liu. 2019b. [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 5720–5726, Hong Kong, China. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Multilingual universal sentence encoder for semantic retrieval](#). *arXiv [e-Print archive]*, abs/1907.04307.
- Pierre Zweigenbaum, Serge Sharoff, and Reinhard Rapp. 2017. [Overview of the second BUCC shared task: Spotting parallel sentences in comparable corpora](#). In *Proceedings of the 10th Workshop on Building and Using Comparable Corpora*, pages 60–67, Vancouver, Canada. Association for Computational Linguistics.