

Cross-Lingual BERT Transformation for Zero-Shot Dependency Parsing

Yuxuan Wang¹, Wanxiang Che^{1*}, Jiang Guo², Yijia Liu¹, and Ting Liu¹

¹Research Center for Social Computing and Information Retrieval,
Harbin Institute of Technology

²Computer Science and Artificial Intelligence Laboratory, MIT

{yxiwang, car, yjliu, tliu}@ir.hit.edu.cn, jiang-guo@csail.mit.edu

Abstract

This paper investigates the problem of learning cross-lingual representations in a contextual space. We propose Cross-Lingual BERT Transformation (CLBT), a simple and efficient approach to generate cross-lingual contextualized word embeddings based on publicly available pre-trained BERT models (Devlin et al., 2018). In this approach, a linear transformation is learned from contextual word alignments to align the contextualized embeddings independently trained in different languages. We demonstrate the effectiveness of this approach on zero-shot cross-lingual transfer parsing. Experiments show that our embeddings substantially outperform the previous state-of-the-art that uses static embeddings. We further compare our approach with XLM (Lample and Conneau, 2019), a recently proposed cross-lingual language model trained with massive parallel data, and achieve highly competitive results.¹

1 Introduction

One of the most promising directions for cross-lingual dependency parsing, which also remains a challenge, is to bridge the gap of lexical features. Prior works (Xiao and Guo, 2014; Guo et al., 2015) have shown that cross-lingual word embeddings are able to significantly improve the transfer performance compared to delexicalized models (McDonald et al., 2011, 2013). These cross-lingual word embeddings are *static* in the sense that they do not change with the context.²

Recently, contextualized word embeddings derived from large-scale pre-trained language models (McCann et al., 2017; Peters et al., 2017, 2018;

*Email corresponding

¹Our code is released at <https://github.com/WangYuxuan93/CLBT>

²In this paper, we refer to these embeddings as *static* as opposed to *contextualized* ones.

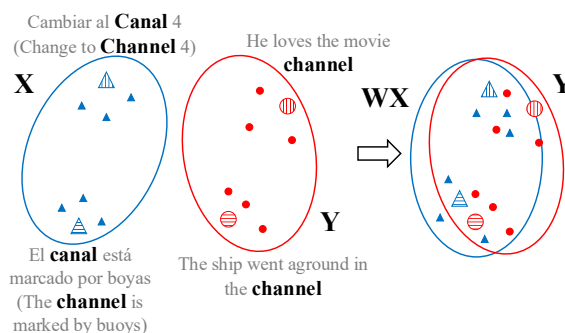


Figure 1: A toy illustration of the method, where contextualized embeddings of the word *canal* from Spanish is transformed to the semantic space of English.

Devlin et al., 2018) have demonstrated dramatic superiority over traditional static word embeddings, establishing new state-of-the-arts in various monolingual NLP tasks (Ilić et al., 2018; Schuster et al., 2018). The success has also been recognized in dependency parsing (Che et al., 2018). The great potential of these contextualized embeddings has inspired us to extend its power to cross-lingual scenarios.

Several recent works have been proposed to learn contextualized cross-lingual embeddings by training cross-lingual language models from scratch with parallel data as supervision, and has been demonstrated effective in several downstream tasks (Schuster et al., 2018; Mulcaire et al., 2019; Lample and Conneau, 2019). These methods are typically resource-demanding and time-consuming.³ In this paper, we propose Cross-Lingual BERT Transformation (CLBT), a simple and efficient off-line approach that learns a linear transformation from contextual word alignments. With CLBT, contextualized embeddings

³For instance, XLM was trained on 64 Volta GPUs (Lample and Conneau, 2019). While the time of training is not described in the paper, we may take the statistics from BERT as a reference, e.g., BERT_{BASE} was trained on 4 Cloud TPUs for 4 days (Devlin et al., 2018).

from pre-trained BERT models in different languages are projected into a shared semantic space. The learned transformation is then used on top of the BERT encodings for each sentence, which are further fed as input to a parser.

Our approach utilizes the semantic equivalence in word alignments, and thus is supposed to be *word sense-preserving*. Figure 1 illustrates our approach, where contextualized embeddings of the Spanish word “canal” are transformed to the corresponding semantic space in English.

Experiments on the Universal Dependencies (UD) treebanks (v2.2) (Nivre et al., 2018) show that our approach substantially outperforms previous models that use static cross-lingual embeddings, with an absolute gain of 2.91% in averaged LAS. We further compare to XLM (Lample and Conneau, 2019), a recently proposed large-scale cross-lingual language model. Results demonstrate that our approach requires significantly fewer training data, computing resources and less training time than XLM, yet achieving highly competitive results.

2 Related Work

Static cross-lingual embedding learning methods can be roughly categorized as *on-line* and *off-line* methods. Typically, on-line approaches integrate monolingual and cross-lingual objectives to learn cross-lingual word embeddings in a joint manner (Klementiev et al., 2012; Kočiský et al., 2014; Guo et al., 2016), while off-line approaches take pre-trained monolingual word embeddings of different languages as input and retrofit them into a shared semantic space (Xing et al., 2015; Lample et al., 2018; Chen and Cardie, 2018).

Several approaches have been proposed recently to connect the rich expressiveness of contextualized word embeddings with cross-lingual transfer. Mulcaire et al. (2019) based their model on ELMo (Peters et al., 2018) and proposed a polyglot contextual representation model by capturing character-level information from multilingual data. Lample and Conneau (2019) adapted the objectives of BERT (Devlin et al., 2018) to incorporate cross-lingual supervision from parallel data to learn cross-lingual language models (XLMs), which have obtained state-of-the-art results on several cross-lingual tasks. Similar to our approach, Schuster et al. (2019) also aligned pre-trained contextualized word embeddings through

linear transformation in an off-line fashion. They used the averaged contextualized embeddings as an *anchor* for each word type, and learn a transformation in the anchor space. Our approach, however, learns this transformation directly in the contextual space, and hence is explicitly designed to be *word sense-preserving*.

3 Cross-Lingual BERT Transformation

This section describes our proposed approach, namely CLBT, to transform pre-trained monolingual contextualized embeddings to a shared semantic space.

3.1 Contextual Word Alignment

Traditional methods of learning static cross-lingual word embeddings have been relying on various sources of supervision such as bilingual dictionaries (Lazaridou et al., 2015; Smith et al., 2017), parallel corpus (Guo et al., 2015) or on-line Google Translate (Mikolov et al., 2013; Xing et al., 2015). To learn contextualized cross-lingual word embeddings, however, we require supervision at word token-level (or context-level) rather than type-level (i.e. dictionaries). Therefore, we assume a parallel corpus as our supervision, analogous to on-line methods such as XLM (Lample and Conneau, 2019).

In our approach, unsupervised bidirectional word alignment is first applied to the parallel corpus to obtain a set of aligned word pairs with their contexts, or *contextual word pairs* for short. For one-to-many and many-to-one alignments, we use the left-most aligned word,⁴ such that all the resulting word pairs are one-to-one. In practice, since WordPiece embeddings (Wu et al., 2016) are used in BERT, all the parallel sentences are tokenized using BERT’s wordpiece vocabulary before being aligned.

3.2 Off-Line Transformation

Given a set of contextual word pairs, their BERT representations $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$ can be easily obtained from pre-trained BERT models,⁵ where $\mathbf{x}_i \in \mathbb{R}^{d_1}$ is the contextualized embedding of token i in the

⁴Preliminary experiments indicate that this way works better than keeping all the alignments.

⁵In this work, we use the English BERT (enBERT) for the source language (English) and the multilingual BERT (mBERT), which is trained on 102 languages without cross-lingual supervision, for all the target languages.

target language, and $\mathbf{y}_i \in \mathbb{R}^{d_2}$ is the representation of its alignment in the source language.

In our experiments, a parser is trained on source language data and applied directly to all the target languages. Therefore, we propose to project the embeddings of target languages to the space of the source language, instead of the opposite direction. Specifically, we aim at finding an appropriate linear transformation \mathbf{W} , such that $\mathbf{W}\mathbf{x}_i$ approximates \mathbf{y}_i .⁶ This can be achieved by solving the following optimization problem:

$$\min_{\mathbf{W}} \sum_{i=1}^n \|\mathbf{W}\mathbf{x}_i - \mathbf{y}_i\|^2,$$

where $\mathbf{W} \in \mathbb{R}^{d_1 \times d_2}$ is a parameter matrix.

Previous works on static cross-lingual embeddings have shown that an orthogonal \mathbf{W} (i.e. $\mathbf{W}^\top \mathbf{W} = \mathbf{I}$) is helpful for the word translation task (Xing et al., 2015). In this case, an analytical solution can be found through singular value decomposition (SVD) of $\mathbf{Y}^\top \mathbf{X}$:

$$\mathbf{W} = \mathbf{V}\mathbf{U}^\top, \text{ where } \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \text{SVD}(\mathbf{Y}^\top \mathbf{X}).$$

Here $\mathbf{X} \in \mathbb{R}^{n \times d}$ and $\mathbf{Y} \in \mathbb{R}^{n \times d}$ are the contextualized embedding matrices, where n is the number of aligned contextual word pairs, d is the dimension of monolingual contextualized embeddings. Each pair of rows $(\mathbf{x}_i, \mathbf{y}_i)$ indicates an aligned contextual word pair.

Although this can be computed in CPUs within several minutes, more memories will be required with the growth of the amount of training data. Therefore, we present an approximate solution, where \mathbf{W} is optimized with gradient decent (GD) and is not constrained to be orthogonal.⁷ This GD-based approach can be trained on a single GPU and typically converges in several hours.

To validate the effectiveness of our approach in cross-lingual dependency parsing, we first obtain the CLBT embeddings with the proposed approach, and then use them as input to a modern graph-based neural parser (described in next section), in replacement of the pre-trained static embeddings. Note that BERT produces embeddings in wordpiece-level, so we only use the left-most wordpiece embedding of each word.⁸

⁶We also investigated non-linear transformation in our experiments, but didn't observe any improvements.

⁷We found the orthogonal constraint doesn't help for GD.

⁸We tried alternative strategies such as averaging, using the middle or right-most wordpiece, but observed no significant difference.

4 Experiments

4.1 Data and Settings

In our experiments, the contextual word pairs are obtained from the Europarl corpora (Koehn, 2005) using the *fast_align* toolkit (Dyer et al., 2010). Only 10,000 sentence pairs are used for each target language. For the parsing datasets, we use the Universal Dependencies(UD) Treebanks (v2.2) (Nivre et al., 2018),⁹ following the settings of the previous state-of-the-art system (Ahmad et al., 2018). From the 31 languages they have analyzed, we select 18 whose Europarl data is publicly available.¹⁰ Statistics of the selected languages and treebanks can be found in the Appendix. We employ the Biaffine Graph-based Parser of Dozat and Manning (2017) and adopt their hyper-parameters for all of our models.

In all the experiments, English is used as the source language, and the other 17 languages as targets. The model is trained on the English treebank and applied directly to target languages with the transformed contextualized embeddings. We train our models using the Adam optimizer (Kingma and Ba, 2015), and most of the them converge within a few thousand epochs in several hours. More implementation details are reported in the Appendix.

4.2 Baseline Systems

We compare our method with the following three baseline models:

- **mBERT** (contextualized). Embeddings generated by the mBERT model are directly used in the training and testing procedures.
- **FT-SVD** (Ahmad et al., 2018, off-line, static). SVD-based transformation (Smith et al., 2017) is applied on 300-dimensional FastText embeddings (Bojanowski et al., 2017) to obtain cross-lingual static embeddings, which represents the previous state-of-the-art. We report results from their paper of the RNNGraph model which used the same architecture as ours.
- **XLM** (Lample and Conneau, 2019, on-line, contextualized). A strong method which learns contextualized cross-lingual embeddings from scratch with cross-lingual data.

⁹hdl.handle.net/11234/1-2837

¹⁰For languages with multiple treebanks, we use the same combinations as they did.

Lan.	Static	Contextualized		
	FT-SVD	mBERT	CLBT (SVD)	CLBT (GD)
en	88.31	90.71	91.03*	
de	59.31	63.41	64.47*	62.14
da	68.81	70.57	71.60*	71.66*
sv	73.49	70.09	73.33*	75.95*
nl	60.11	65.66	65.45	63.86
fr	73.46	72.97	74.70*	76.59*
it	76.23	79.02	79.46	78.98
es	66.91	65.43	67.14*	68.33*
pt	67.98	67.11	69.12*	69.25*
ro	52.11	46.40	55.14*	55.84*
sk	56.98	50.76	59.46*	59.92*
pl	58.59	63.10	65.37*	65.80*
bg	66.68	71.20	70.26	70.75
sl	54.57	56.78	57.42*	57.21*
cs	52.80	45.20	52.20*	52.99*
fi	48.74	49.56	51.00*	52.61*
et	44.40	46.64	47.79*	48.52*
lv	49.59	45.11	48.59*	49.78*
AVG.	60.63	60.53	63.09	63.54

Table 1: Results (LAS%) on test sets. Languages are split by language families with dashed lines. AVG. means the average of results from all target languages. Statistically significant differences between our methods and the mBERT model are marked with *, with p-value < 0.05 under McNemar’s test.

For the **XLM** model, we employ the XNLI-15 model¹¹ they released to generate embeddings and apply them to cross-lingual dependency parsing in the same way as we do with our own model. We compare with them in the 4 overlapped languages both works have researched on.

4.3 Comparison with Off-Line Methods

Results on the test sets are shown in Table 1.¹² Languages are grouped by language families. Overall, our approach with either SVD or GD outperforms both FT-SVD and mBERT by a substantial margin (+2.91% in averaged LAS), among which GD turns out to be slightly better than SVD in most of the languages. When combined with FT-SVD, the performances can be further improved by 0.33% in LAS for the GD method and 0.51% for SVD (see the Appendix for more details). Interestingly, the mBERT model which is trained without any cross-lingual supervision but

¹¹github.com/facebookresearch/XLM

¹²UAS results are listed in the Appendix due to space limit. Note that since we have no access to the parsed files of the FT-SVD model, we only report statistical significant tests between our methods and the mBERT model, which is highly comparable to the FT-SVD model on average.

Lan.	XLM	CLBT (SVD)	CLBT (GD)
en	91.85/89.92	92.81*/91.03*	
de	74.65/65.31	73.72/64.47	71.08/62.14
fr	79.62/73.41	80.01/74.70*	80.85*/76.59*
es	75.41/67.43	75.52/67.14*	75.70*/68.33*
bg	81.07/69.45	82.14*/70.26	81.51/ 70.75*
AVG.	77.69/68.90	77.85/69.14	77.29/ 69.45
Data	0.2-13.1M	10K	

Table 2: Results (UAS%/LAS%) on test sets. The last row shows the training data used in each language by sentence. AVG. means the average of results from 4 target languages. Statistically significant differences between our methods and the XLM are marked with an asterisk, with p-value < 0.05 under McNemar’s test.

using a shared multilingual wordpiece vocabulary works surprisingly well in some languages, especially in those linguistically close to English. Similar observations have also been identified in other works (Pires et al., 2019; Wu and Dredze, 2019).

4.4 Comparison with On-Line Methods

Comparison of our approach and a cross-lingual language model pre-training (XLM) method (Lample and Conneau, 2019) in the 4 overlapped languages is shown in Table 2. CLBT outperforms XLM in 3 out of the 4 languages but lower in German (de). The amount of training data used in each method is also shown in the bottom: the number of parallel sentences used by XLM ranges from 0.2 million (10 million tokens) for Bulgarian to 13.1 million (682 million tokens) for French. In comparison, only 10,000 parallel sentences (0.4 million tokens) are used for each language in CLBT, demonstrating the data-efficiency of our approach. Moreover, given the efficiency in both data and training, CLBT can be readily scaled to new language pairs in hours.

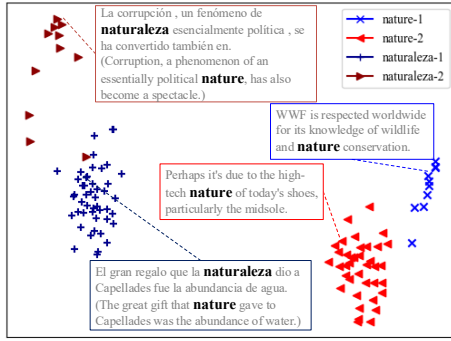
4.5 Analysis

4.5.1 Transformation of Cross-lingual BERT Embedding


In order to investigate the properties of contextualized representations before and after the linear transformation, we employ the SENSEVAL2 data (Edmonds and Cotton, 2001),¹³ where words from different languages are tagged by their word senses in different contexts.

We took contextualized representations of the English word *nature* and its Spanish transla-

¹³www.hipposmond.com/senseval2/



(a)



(b)

Figure 2: t-SNE visualization of the English word *nature* and its Spanish translation *naturaleza* in different contexts by the contextualized representations before (a) and after (b) the linear transformation. Points are colored by word senses. Example contexts are given in (a). Translations of Spanish sentences are in brackets.

tion *naturaleza* in different contexts from pre-trained English and multilingual BERT respectively and visualize their distributions in Figure 2(a), where we can observe obvious clustering of word senses. Specifically, words with sense *nature-1* and *naturaleza-1* mean *the physical world*, whereas *nature-2* and *naturaleza-2* mean *inherent features*. We then apply our GD-based method to embeddings of *naturaleza* and depict the resulting cross-lingual embeddings in Figure 2(b). The distance between embeddings from English and Spanish is effectively reduced after the transformation. And it is apparent that embeddings of Spanish words are closer to those with similar meanings from English, which indicates the effectiveness of our approach.

4.5.2 Effect of Training Data Size

We select several languages from each language family, and investigate the effect of the amount of training data on the performances of zero-shot cross-lingual dependency parsing. Specifically,




Figure 3: Effects of the amount of training data on different languages. (*y*-axis represents the LAS.)

we take the SVD-based approach, since it is faster than the GD-based one, and trained different transformation models with different amount of parallel sentences from Europarl dataset on each of the 13 selected languages.

As shown in Figure 3, for most of the languages, the best performance is achieved with only 5000 parallel sentences. It is also worth noting that for most of Germanic (e.g. German, Danish, Swedish and Dutch) and Romance (e.g. French, Italian, Spanish and Romanian) languages, which are typologically closer to English, a rather small training set of merely 100 sentences is capable of achieving comparative results.

5 Conclusion

We propose the Cross-Lingual BERT Transformation (CLBT) approach for contextualized cross-lingual embedding learning, which substantially outperforms the previous state-of-the-art in zero-shot cross-lingual dependency parsing. By exploiting publicly available pre-trained BERT models, our approach provides a fast and data-efficient solution to learning cross-lingual contextualized embeddings. Compared to the XLM, our method requires much fewer parallel data and less training time, yet achieving comparable performance.

For future work, we are interested in unsupervised cross-lingual alignment, inspired by prior success on static embeddings (Lample et al., 2018; Alvarez-Melis and Jaakkola, 2018), which demands a deeper understanding to the geometry of the multilingual contextualized embedding space.

Acknowledgments

We thank the anonymous reviewers for their valuable suggestions. This work was supported by the National Natural Science Foundation of China (NSFC) via grant 61976072, 61632011 and 61772153.

References

- Wasi Uddin Ahmad, Zhisong Zhang, Xuezhe Ma, Eduard H. Hovy, Kai-Wei Chang, and Nanyun Peng. 2018. [Near or far, wide range zero-shot cross-lingual dependency parsing](#). *CoRR*, abs/1811.00570.
- David Alvarez-Melis and Tommi Jaakkola. 2018. Gromov-Wasserstein alignment of word embedding spaces. In *Proc. of EMNLP*, pages 1881–1890.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Wanxiang Che, Yijia Liu, Yuxuan Wang, Bo Zheng, and Ting Liu. 2018. Towards better ud parsing: Deep contextualized word embeddings, ensemble, and treebank concatenation. In *Proc. of CoNLL*, pages 55–64. Association for Computational Linguistics.
- Xilun Chen and Claire Cardie. 2018. Unsupervised multilingual word embeddings. In *Proc. of EMNLP*, pages 261–270. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *Proc. of ICLR*.
- Chris Dyer, Adam Lopez, Juri Ganitkevitch, Jonathan Weese, Ferhan Ture, Phil Blunsom, Hendra Setiawan, Vladimir Eidelman, and Philip Resnik. 2010. cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *Proc. of the ACL 2010 System Demonstrations*, pages 7–12. Association for Computational Linguistics.
- Philip Edmonds and Scott Cotton. 2001. Senseval-2: Overview. In *The Proc. of the Second International Workshop on Evaluating Word Sense Disambiguation Systems*, SENSEVAL '01, pages 1–5, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. Cross-lingual dependency parsing based on distributed representations. In *Proc. of ACL and IJCNLP*, pages 1234–1244.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *Proc. of AACL*, pages 2734–2740.
- Suzana Ilić, Edison Marrese-Taylor, Jorge Balazs, and Yutaka Matsuo. 2018. Deep contextualized word representations for detecting sarcasm and irony. In *Proc. of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 2–7. Association for Computational Linguistics.
- Diederick P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proc. of ICLR*.
- Alexandre Klementiev, Ivan Titov, and Binod Bhat-tarai. 2012. Inducing crosslingual distributed representations of words. In *Proc. of COLING*, pages 1459–1474.
- Tomáš Kočiský, Karl Moritz Hermann, and Phil Blunsom. 2014. Learning bilingual word representations by marginalizing alignments. In *Proc. of ACL*, pages 224–229. Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proc. of the tenth Machine Translation Summit*, pages 79–86, Phuket, Thailand. AAMT.
- Guillaume Lample and Alexis Conneau. 2019. [Cross-lingual language model pretraining](#). *CoRR*, abs/1901.07291.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proc. of ICLR*.
- Angeliki Lazaridou, Georgiana Dinu, and Marco Baroni. 2015. Hubness and pollution: Delving into cross-space mapping for zero-shot learning. In *Proc. of ACL and IJCNLP*, pages 270–280. Association for Computational Linguistics.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems 30*, pages 6294–6305. Curran Associates, Inc.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. Universal dependency annotation for multilingual parsing. In *Proc. of ACL*, pages 92–97. Association for Computational Linguistics.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proc. of EMNLP*, pages 62–72.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. [Exploiting similarities among languages for machine translation](#). *CoRR*, abs/1309.4168.
- Phoebe Mulcaire, Jungo Kasai, and Noah Smith. 2019. Polyglot contextual representations improve crosslingual transfer. In *Proc. of NAACL*.

- Joakim Nivre, Mitchell Abrams, Željko Agić, and et al. 2018. [Universal dependencies 2.2](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Matthew Peters, Waleed Ammar, Chandra Bhagavatula, and Russell Power. 2017. Semi-supervised sequence tagging with bidirectional language models. In *Proc. of ACL*, pages 1756–1765. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*, pages 2227–2237. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual BERT? In *Proc. of ACL*, pages 4996–5001.
- Sebastian Schuster, Sonal Gupta, Rushin Shah, and Mike Lewis. 2018. [Cross-lingual transfer learning for multilingual task oriented dialog](#). *CoRR*, abs/1810.13327.
- 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 *Proc. of NAACL*, pages 1599–1613.
- Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. In *Proc. of ICLR*.
- Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of bert. In *Proc. of EMNLP*.
- 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](#). *CoRR*, abs/1609.08144.
- Min Xiao and Yuhong Guo. 2014. Distributed word representation learning for cross-lingual dependency parsing. In *Proc. of CoNLL*, pages 119–129. Association for Computational Linguistics.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proc. of NAACL*, pages 1006–1011. Association for Computational Linguistics.