

Zero-shot Dependency Parsing with Pre-trained Multilingual Sentence Representations

Ke Tran*

Amazon Alexa AI
trnke@amazon.com

Arianna Bisazza†

University of Groningen
a.bisazza@rug.nl

Abstract

We investigate whether off-the-shelf deep bidirectional sentence representations (Devlin et al., 2018) trained on a massively multilingual corpus (multilingual BERT) enable the development of an unsupervised universal dependency parser. This approach only leverages a mix of monolingual corpora in many languages and does not require any translation data making it applicable to low-resource languages. In our experiments we outperform the best CoNLL 2018 language-specific systems in all of the shared task’s six truly low-resource languages while using a single system. However, we also find that (i) parsing accuracy still varies dramatically when changing the training languages and (ii) in some target languages zero-shot transfer fails under all tested conditions, raising concerns on the ‘universality’ of the whole approach.

1 Introduction

Pretrained sentence representations (Howard and Ruder, 2018; Radford et al., 2018; Peters et al., 2018; Devlin et al., 2018) have recently set the new state of the art in many language understanding tasks (Wang et al., 2018). An appealing avenue for this line of work is to use a mix of training data in several languages and a shared subword vocabulary leading to general-purpose multilingual representations. In turn, this opens the way to a number of promising cross-lingual transfer techniques that can address the lack of annotated data in the large majority of world languages.

In this paper, we investigate whether deep bidirectional sentence representations (Devlin et al., 2018) trained on a massively multilingual corpus

(m-BERT) allow for the development of a universal dependency parser that is able to parse sentences in a diverse range of languages without receiving *any supervision* in those language. Our parser is fully lexicalized, in contrast to a successful approach based on delexicalized parsers (Zeman and Resnik, 2008; McDonald et al., 2011). Building on the delexicalized approach, previous work employed additional features such as typological properties (Naseem et al., 2012), syntactic embeddings (Duong et al., 2015), and cross-lingual word clusters (Täckström et al., 2012) to boost parsing performance. More recent work by Ammar et al. (2016); Guo et al. (2016) requires translation data for projecting word embeddings into a shared multilingual space.

Among lexicalized systems in CoNLL18, the top system (Che et al., 2018) utilizes contextualized vectors from ELMo. However, they train each ELMo for each language in the shared task. While their approach achieves the best LAS score on average, for low resource languages, the performance of their parser lags behind other systems that do not use pre-trained models (Zeman et al., 2018). By contrast, we build our dependency parser on top of general-purpose context-dependent word representations pretrained on a multilingual corpus. This approach does not require any translation data making it applicable to truly low-resource languages (§3.3). While m-BERT’s training objective is inherently monolingual (predict a word in language ℓ given its sentence context, also in language ℓ), we hypothesize that cross-lingual syntactic transfer occurs via the shared subword vocabulary and hidden layer parameters. Indeed, on the challenging task of universal dependency parsing from raw text, we outperform by a large margin the best CoNLL18 language-specific systems (Zeman et al., 2018) on the shared task’s truly low-resource

* Work done prior to joining Amazon.

† Work done while at Leiden University. Both authors contributed equally.

languages while using a single system.

The effectiveness of m-BERT for cross-lingual transfer of UD parsers has also been demonstrated in concurrent work by Wu and Dredze (2019) and Kondratyuk (2019). While the former utilizes only English as the training language, the latter trains on a concatenation of all available UD treebanks. We additionally experiment with three different sets of training languages beyond English-only and make interesting observations on the resulting large, and sometimes unexplicable, variation of performance among test languages.

2 Model

We use the representations produced by BERT (Devlin et al., 2018) which is a self-attentive deep bidirectional network (Vaswani et al., 2017) trained with a masked language model objective. Specifically we use BERT’s multilingual cased version¹ which was trained on the 100 languages with the largest available Wikipedias. Exponentially smoothed weighting was applied to prevent high-resource languages from dominating the training data, and a shared vocabulary of 110k shared WordPieces (Wu et al., 2016) was used. For parsing we employ a modification of the graph-based dependency parser of Dozat and Manning (2016). We use deep biaffine attention to score arcs and their label from the head to its dependent. While our label prediction model is similar to that of Dozat and Manning (2016), our arc prediction model is a globally normalized model which computes partition functions of non-projective dependency structures using Kirchhoff’s Matrix-Tree Theorem (Koo et al., 2007).

Let $\mathbf{x} = w_1, w_2, \dots, w_n$ be an input sentence of n tokens, which are given by the gold segmentation in training or by an automatic tokenizer in testing (§3.1). To obtain the m-BERT representation of \mathbf{x} , we first obtain a sequence $\mathbf{t} = t_1, \dots, t_m$ of $m \geq n$ subwords from \mathbf{x} using the WordPiece algorithm. Then we feed \mathbf{t} to m-BERT and extract the representations $\mathbf{e}_1, \dots, \mathbf{e}_m$ from the last layer. If word w_i is tokenized into (t_j, \dots, t_k) then the representation \mathbf{h}_i of w_i is computed as the mean of $(\mathbf{e}_j, \dots, \mathbf{e}_k)$.

¹<https://github.com/google-research/bert/blob/master/multilingual.md>

The arc score is computed similar to Dozat and Manning (2016):

$$\mathbf{s}^{(\text{arc})} = \text{DeepBiaffine}(\mathbf{H}^{(\text{arc-head})}, \mathbf{H}^{(\text{arc-dep})}) \quad (1)$$

The log probability of the dependency tree \mathbf{y} of \mathbf{x} is given by

$$\log p(\mathbf{y} | \mathbf{x}) = \sum_{(h,c) \in \mathbf{y}} \mathbf{s}^{(\text{arc})}[h, c] - \log Z(\mathbf{x}) \quad (2)$$

where $Z(\mathbf{x})$ is the partition function. Our objective function for predicting dependency arcs therefore is globally normalized. We compute $Z(\mathbf{x})$ via matrix determinant (Koo et al., 2007). In our experiments, we find that training with a global objective is more stable if the score $\mathbf{s}^{(\text{arc})}[h, c]$ is locally normalized² such that $\sum_h \exp(\mathbf{s}^{(\text{arc})}[h, c]) = 1$. During training, we update both m-BERT and parsing layer parameters.

3 Experiments

While most previous work on parser transfer, including the closely related (Duong et al., 2015) relies on gold tokenization and POS tags, we adopt the more realistic scenario of parsing from *raw text* (Zeman et al., 2018) and adopt the automatic sentence segmenter and tokenizer provided as baselines by the shared task organizers.

3.1 Data

We use the UDpipe-tokenized test data³ (Straka and Straková, 2017) and the CoNLL18 official script for evaluation. Gold tokenization is only used for the training data, while POS information is never used. All of our experiments are carried out on the Universal Dependencies (UD) corpus version 2.2 (Nivre et al., 2018) for a fair comparison with previous work.

While our sentence representations are always initialized from m-BERT, we experiment with four sets of parser training (i.e. fine-tuning) languages, namely: *expEn* only English (200K words); *expLatin* a mix of four Latin-script European languages: English, Italian, Norwegian, Czech (50K each, 200K in total); *expSOV* a mix of two

²We use $\log_{\text{softmax}}(\mathbf{s}^{(\text{arc})})$ in place of $\mathbf{s}^{(\text{arc})}$ in equation 2.

³Preprocessed data available at <http://hdl.handle.net/11234/1-2899>

SOV languages: Hindi and Korean (100K each, 200K in total); *expMix* a larger mix of eight languages including different language families and scripts: English, Italian, Norwegian, Czech, Russian, Hindi, Korean, Arabic (50K each, 400K in total). For high resource languages that have more than one treebank, we choose the treebank that has the best LAS score in ConLL18 for training and the lowest LAS score for zero-shot evaluation.

3.2 Training details

Similar to Dozat et al. (2017), we use a neural network output size of 400 for arc prediction and 100 for label prediction. We use the Adam optimizer with learning rate $5e^{-6}$ to update the parameters of our models. The model is evaluated every 500 updates and we stop training if the score LAS does not increase in ten consecutive validations.

3.3 Results

To put our results into perspective, we report the accuracy of the best CoNLL18 system for each language and that of the Stanford system submitted at the same evaluation (Qi et al., 2018). The latter is also based on the deep biaffine parser of Dozat and Manning (2016), it does not use ensembles and was ranked 2nd on official evaluation metric LAS⁴. Both these parsers receive supervision in most of the languages, therefore comparison to our parser is only fair for the low-resource languages where training data is not available (or negligible, i.e. less than 1K tokens).

Results for a subset of UD languages are presented in Table 1. Beside common European languages, we choose languages with different writing scripts than those presented in the parser training data. We also include SOV (e.g., Korean, Persian) and VSO (e.g., Arabic, Breton) languages. Parser training languages for each experiment are highlighted in grey in Table 1.

In the high resource setting, there is a considerable gap between zero-shot and supervised parsers with Swedish as an exception (slightly better than Stanford’s parser and 2 points below CoNLL18). By contrast, the benefit of multilingual transfer becomes evident in the low resource setting. Here,

⁴Updated results at <https://stanfordnlp.github.io/stanfordnlp/performance.html> March, 2018

most CoNLL18 systems including Stanford’s use knowledge of each target language to customize the parser, e.g., to choose the optimal training language(s). Nevertheless, our single parser trained on the largest mix of languages (*expMix*) beats the best CoNLL18 language-specific systems on all six languages, even though three of these languages are not represented in m-BERT’s training data⁵. This result highlights the advantage of multilingual pre-trained model in the truly low resource scenario.

We notice the poor performance of our parser on spoken French in comparison to other European languages. While there is sufficient amount of Wikipedia text for French, it seems that zero-shot parsing on a different domain remains a challenge even with a large pre-trained model.

4 Analysis

By varying the set of parser training languages we analyze our results with respect to two factors: parser training language diversity and word order similarity.

4.1 Training language diversity

Increasing language diversity (*expEn*→*expLatin* and *expLatin*→*expMix*) leads to improvements in most test languages, even when the total amount of training data is fixed (*expEn*→*expLatin*). The only exceptions are the languages for which training data is reduced (English in *expLatin*) or becomes a smaller proportion of the total training data (Czech, Italian, Norwegian in *expMix*), which confirms previous findings (Ammar et al., 2016). Swedish and Upper Sorbian being related to Norwegian and Czech respectively also lose some accuracy in *expMix*. On the other hand, newly included languages (Czech, Italian, Norwegian in *expLatin* and Arabic, Hindi, Korean, Russian in *expMix*) show the biggest improvements, which was also expected.

More interestingly, some large gains are reported for languages that are unrelated from all training languages of *expLatin*. We hypothesize that such languages (Arabic, Armenian, Hungarian) may benefit from an exposure of the parser to a

⁵This is possible because their sub-words are in BERT’s vocabulary due other similar languages in training data.

target	tbk-code	m-BERT based				State of the art		
		<i>expEn</i>	<i>expLatin</i>	<i>expSOV</i>	<i>expMix</i>	Stanford	CoNLL18	#TrWrds
Russian	<i>ru_syntagrus</i>	59.53	73.13	34.44	81.91	91.20	92.48	872 K
Hindi	<i>hi_hdtb</i>	32.94	33.75	88.51	85.66	91.65	92.41	281 K
Italian	<i>it_isdt</i>	75.45	89.59	25.95	89.44	90.51	92.00	276 K
Norwegian	<i>no_nynorsk</i>	72.09	86.01	33.93	85.11	89.58	90.99	245 K
Czech	<i>cs_pdt</i>	59.97	84.91	34.31	84.36	89.63	89.63	1,173 K
Finnish	<i>fi_tdt</i>	50.65	61.13	40.12	62.29	86.33	88.73	163 K
Persian	<i>fa_seraji</i>	44.34	56.39	24.77	56.92	86.55	88.11	121 K
Korean	<i>ko_kaist</i>	33.67	38.87	84.39	81.73	86.58	86.91	296 K
English	<i>en_ewt</i>	84.64	82.38	30.03	81.65	83.80	84.57	205 K
Urdu	<i>ur_udtb</i>	23.46	23.94	65.21	63.06	82.58	83.39	109 K
Japanese	<i>ja_gsd</i>	12.92	12.65	19.25	24.10	78.48	83.11	162 K
Hungarian	<i>hu_szeged</i>	52.72	61.11	39.65	61.11	78.58	82.66	20 K
German	<i>de_gsd</i>	68.30	70.93	36.30	70.93	79.17	80.36	264 K
Swedish	<i>sv_pud</i>	76.02	78.71	37.58	78.70	78.39	80.35	– K
Arabic	<i>ar_padt</i>	34.55	50.20	12.26	68.20	76.99	77.06	224 K
French	<i>fr_spoken</i>	54.12	59.70	16.06	59.54	69.56	75.78	15 K
Vietnamese	<i>vi_vtb</i>	29.72	30.09	14.13	29.71	47.56	55.22	20 K
Tamil	<i>ta_ttb</i>	18.09	25.79	29.64	32.78	–	–	5 K
Telugu	<i>te_mtg</i>	54.47	63.06	61.68	64.03	–	–	5 K
Faroese*	<i>fo_ofst</i>	58.28	61.71	36.27	61.98	41.54	49.43	0 K
Upper Sorbian*	<i>hsb_ufal</i>	36.66	49.90	23.90	49.74	23.61	46.42	0 K
Breton	<i>br_keb</i>	45.16	51.85	22.49	52.62	11.25	38.64	0 K
Armenian	<i>hy_armtdp</i>	40.20	55.44	41.91	58.95	31.47	37.01	1 K
Kazakh	<i>kk_ktb</i>	33.56	40.18	40.18	44.56	26.25	31.93	1 K
Buryat*	<i>bxr_bdt</i>	19.19	20.90	22.94	23.11	12.47	19.53	0 K
avg(lowRes)		39.41	47.26	31.28	48.45	24.43	37.16	

Table 1: LAS scores of our parser in the raw text setup. Languages not in m-BERT’s training corpus are marked with *. *SVO* and *SOV* languages are indicated by purple and green respectively. Stanford and CoNLL18’s best submitted systems are provided as representative state-of-the-art supervised systems. #TrWrds = Total training data made available at CoNLL18. The amount of training used in each experiment is specified in §3.1. Training languages for each experiment are highlighted in grey.

more diverse set of word orders (§4.2). For instance, Arabic being head initial is closer to Italian than to English in terms of word order.

Actual language relatedness does not always play a clear role: For instance, Upper Sorbian seems to benefit largely from its closeness to Czech in *expLatin* and *expMix*, while Faroese (related to Norwegian) does not improve as much.

In summary, language diversity in training is clearly a great asset. However, there is a large variation in gains among test languages, for which language family relatedness can only offer a partial explanation.

4.2 Training language typology

Training on languages with similar typological features has been shown beneficial for parsing target languages in the delexicalized setup. In particular, word order similarities have been proved beneficial to select source languages for parsing model transfer (Naseem et al., 2012; Duong et al., 2015).

Indeed, when Hindi and Korean are presented in *expSOV*, we report better LAS scores in various *SOV* languages (Japanese, Tamil, Urdu, Buryat) however some other *SOV* languages (Persian and Armenian) perform much worse than in *expLatin* showing that word order is not a reliable criterion for training language selection.

Given these observations, we construct our largest training data (*expMix*) by merging all the languages of *expEn*, *expLatin*, and *expSOV* and adding two more languages with diverse word order profiles for which large treebanks exist, namely Russian and Arabic.

Concurrently to this work, Lin et al. (2019) have proposed an automatic method to choose the optimal transfer languages in various tasks including parsing, based on a variety of typological but also data-dependent features. We leave adoption of their method to future work.⁶

⁶Unfortunately at the time of writing we have not yet managed to use their released implementation.

4.3 Towards explaining transfer performance

Even when keeping the training languages fixed, for instance in *expMix*, we observe a large variation of zero-shot parsing transfer accuracy among test languages which does not often correlate with supervised parsing accuracy. As an attempt to explain this variation we look at the overlap of test vocabulary with (i) parser’s training data vocabulary τ and (ii) m-BERT’s training data vocabulary. Because m-BERT uses a subword vocabulary that also includes characters we resort to measuring the unsegmented word score η :

$$\tau = 100 \times |\text{type}_w(D_{\text{test}}) \cap \text{type}_w(D_{\text{train}})| / |\text{type}_w(D_{\text{test}})|$$

$$\eta = 100 \times |\text{token}_g(D_{\text{test}})| / |\text{token}_w(D_{\text{test}})|$$

where $\text{type}_w(D)$ and $\text{token}_w(D)$ are sets of WordPieces types and tokens in dataset D respectively, and $\text{token}_g(D)$ is the set of gold tokens in D before applying WordPieces. A higher η indicates a less segmented text.

To account for typological features, we also plot the average syntactic similarity $\bar{\sigma}$ of each test language to the eight *expSOV* training languages as computed by the URIEL database⁷ (Littell et al., 2017).




Figure 1: Relationship between parsing accuracy (*expMix*), parser training-test vocabulary overlap τ , m-BERT unsegmented word score η , and average typological syntactic similarity $\bar{\sigma}$. Purple bar indicates there is no language that belongs to the same family presented in training data. Languages in the training set of *expMix* are not shown.

We observe a correlation between LAS, η and τ for test languages that have a relative in the training data, like Urdu and Hindi. For test languages

⁷Specifically, we compute $1 - d$ where d is the pre-computed *syntactic* distance in `lang2vec`.

that belong to a different family than all training languages, no correlation appears. A similar observation is also reported by Pires et al. (2019): namely, they find that the performance of cross-lingual named entity recognition with m-BERT is largely independent of vocabulary overlap.

Although typological features have been shown to be useful when incorporated into the parser (Naseem et al., 2012; Ammar et al., 2016), we do not find a clear correlation between $\bar{\sigma}$ and LAS in our setup. Thus none of our investigated factors can explain transfer performance in a systematic way.

4.4 Language outliers

While massively pre-trained language models promise a more inclusive future for NLP, we find it important to note that cross-lingual transfer performs very badly for some languages.

For instance, in our experiments, Japanese and Vietnamese stand out as strikingly negative outliers. Wu and Dredze (2019) also report a very low performance on Japanese in their zero-shot dependency parsing experiments.⁸ In (Lin et al., 2019) Japanese is completely excluded from the parsing experiments because of unstable results. Japanese and Vietnamese are *language isolates* in an NLP sense, meaning that they do not enjoy the presence of a closely related language among the high-resourced training languages.⁹ For this class of languages, transfer performance is overall very inconsistent and hard to explain.

	UDpipe	Gold
ko→ja	14.96	20.04
ja→ko	37.44	37.45

Table 2: LAS scores when transferring between Korean and Japanese in two tokenization conditions.

The case of Japanese is particularly interesting for its relation to Korean. Family relatedness between these two languages is very controversial but their syntactic features are extremely similar. To put our

⁸They do not report parsing results for Vietnamese.

⁹The original definition of language isolate in linguistics is actually stronger: “a language that has no known relatives, that is, that has no demonstrable phylogenetic relationship with any other language” (Campbell, 2017)

parser in optimal transfer conditions, we perform one last experiment by training only on Korean (all available data) and testing on Japanese, and vice versa. As shown in Table 2, Japanese performance becomes even lower in this setup. We can also see that transferring in the opposite direction leads to a much better result, despite the fact that state-of-the-art supervised systems in these two languages achieve similar results (Japanese: 83.11, Korean: 86.92 by the best CoNLL18 systems). To rule out the impact of unsupervised sentence and token segmentation, which may be performing particularly poorly on some languages, we retrain the parser with gold segmentation and find that it explains only a small part of the gap.

While Pires et al. (2019) hypothesize word order is the main culprit for the poor zero-shot performance for Japanese when transferring a POS-tagger from English, our experiments with Korean and Japanese show a different picture.

5 Conclusions

We have built a Universal Dependency parser on top of deep bidirectional sentence representations pre-trained on a massively multilingual corpus (m-BERT) without any need for parallel data, treebanks or other linguistic resources in the test languages.

Evaluated in the challenging scenario of parsing from raw text, our best parser trained on a mix of languages representing both language family and word order diversity outperforms the best CoNLL18 language-specific systems on the six truly low-resource languages presented at the shared task.

Our experiments show that language diversity in the training treebank is a great asset for transfer to low-resource languages. Moreover, the massively multilingual nature of m-BERT does not neutralize the impact of transfer languages on parsing accuracy, which is only partially explained by language relatedness and word order similarity.

Finally we have raised the issue of language outliers that perform very poorly in all our tested conditions and that, given our analysis, are unlikely to benefit even from automatic methods of transfer language selection (Lin et al., 2019).

Acknowledgements

Arianna Bisazza was funded by the the Netherlands Organization for Scientific Research (NWO) under project number 639.021.646.

References

- Joakim Nivre 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.
- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. Many languages, one parser. *Transactions of the Association for Computational Linguistics*, 4:431–444.
- Lyle Campbell. 2017. *"Language Isolates"*. Routledge.
- 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 *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 55–64, Brussels, Belgium. 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. *arXiv e-prints*, page arXiv:1810.04805.
- Timothy Dozat and Christopher D. Manning. 2016. Deep biaffine attention for neural dependency parsing. *CoRR*, abs/1611.01734.
- Timothy Dozat, Peng Qi, and Christopher D. Manning. 2017. Stanford’s graph-based neural dependency parser at the conll 2017 shared task. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 20–30. Association for Computational Linguistics.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. Cross-lingual transfer for unsupervised dependency parsing without parallel data. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 113–122, Beijing, China. Association for Computational Linguistics.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pages 2734–2740. AAAI Press.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.

- Daniel Kondratyuk. 2019. 75 languages, 1 model: Parsing universal dependencies universally. *CoRR*, abs/1904.02099.
- Terry Koo, Amir Globerson, Xavier Carreras, and Michael Collins. 2007. Structured prediction models via the matrix-tree theorem. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 141–150, Prague, Czech Republic. Association for Computational Linguistics.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastopoulos, 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, Florence, Italy. Association for Computational Linguistics.
- Patrick Littell, David R Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. 2017. Uriel and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, volume 2, pages 8–14.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 62–72, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 629–637, Jeju Island, Korea. 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 *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, New Orleans, Louisiana. 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 Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Peng Qi, Timothy Dozat, Yuhao Zhang, and Christopher D. Manning. 2018. Universal dependency parsing from scratch. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text*
- to Universal Dependencies*, pages 160–170, Brussels, Belgium. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning. Technical report, OpenAI.
- Milan Straka and Jana Straková. 2017. Tokenizing, pos tagging, lemmatizing and parsing ud 2.0 with udpipes. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99, Vancouver, Canada. Association for Computational Linguistics.
- Oscar Täckström, Ryan McDonald, and Jakob Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 477–487. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT. *CoRR*, abs/1904.09077.
- 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.
- Daniel Zeman, Hajič Jan, Martin Popel, Martin Potthast, Milan Straka, Filip Ginter, Joakim Nivre, and Slav Petrov. 2018. Conll 2018 shared task: Multilingual parsing from raw text to universal dependencies. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–21. Association for Computational Linguistics.
- Daniel Zeman and Philip Resnik. 2008. Cross-language

parser adaptation between related languages. In *Proceedings of the IJCNLP-08 Workshop on NLP for Less Privileged Languages*.