

Systematic Investigation of Strategies Tailored for Low-Resource Settings for Low-Resource Dependency Parsing

Jivnesh Sandhan¹, Laxmidhar Behera^{1,3} and Pawan Goyal²

¹IIT Kanpur, ²IIT Kharagpur, ³IIT Mandi

jivnesh@iitk.ac.in, pawang@cse.iitkgp.ac.in

Abstract

In this work, we focus on low-resource dependency parsing for multiple languages. Several strategies are tailored to enhance performance in low-resource scenarios. While these are well-known to the community, it is not trivial to select the best-performing combination of these strategies for a low-resource language that we are interested in, and not much attention has been given to measuring the efficacy of these strategies. We experiment with 5 low-resource strategies for our ensembled approach on 7 Universal Dependency (UD) low-resource languages. Our exhaustive experimentation on these languages supports the effective improvements for languages not covered in pretrained models. We show a successful application of the ensembled system on a truly low-resource language Sanskrit.¹

1 Introduction

Recently, the supervised learning paradigm has dramatically increased the state-of-the-art performance for the dependency parsing task for resource-rich languages (Chen and Manning, 2014; Dyer et al., 2015; Kiperwasser and Goldberg, 2016; Dozat and Manning, 2017; Kulmizev et al., 2019). However, only a handful of resource-rich languages are able to take advantage, and many low-resource languages are far from these benefits (Joshi et al., 2020; More et al., 2019a; Zeman et al., 2018).

In literature, several strategies have been proposed to enhance performance in low-resource scenarios, such as data augmentation (Şahin and Steedman, 2018; Gulordava et al., 2018), cross/monolingual pretraining (Conneau et al., 2020; Peters et al., 2018; Kondratyuk and Straka, 2019), sequential transfer learning (Ruder et al., 2019), multi-task learning (Nguyen and Verspoor, 2018), cross-lingual transfer (Das and Sarkar, 2020) and self-training (Rotman and Reichart, 2019; Clark et al.,

2018). However, not much attention has been given to measuring the efficacy of the existing low-resource strategies well-known to the community for low-resource dependency parsing (Vania et al., 2019). This is essential to assess their utility for low-resource languages (Hedderich et al., 2021) before inventing novel ways to tackle data sparsity.

In this work, we systematically explore 5 pragmatic strategies for low-resource settings on 7 languages. We experiment with low-resource strategies such as data augmentation, sequential transfer learning, cross/mono-lingual pretraining, multi-task learning and self-training. We investigate: (1) How is the trend in performance of each strategy across various languages? Whether the choice of best performing variant of each strategy is language dependent? (2) We integrate the best performing variant of each strategy and call the resulting system as the ensembled system. Do all the strategies contribute towards performance gain in the ensembled system? How well does this ensemble approach generalize across multiple low-resource languages? (3) How far can we push a purely data-driven ensemble system using the best-performing low-resource strategies? Can this simple ensemble approach outperform the state-of-the-art of a low-resource language? We argue that while it may sound like a simple application of techniques well known to the community; it is non-trivial to select the best performing combination for a target low-resource language.

Our exhaustive experimentation empirically establishes the effective generalization ability of the ensembled system on 7 languages and shows average absolute gains of 5.2/6.2 points Unlabelled/Labelled Attachment Score (UAS/LAS) over strong baseline (Dozat et al., 2017). Notably, our ensembled system shows substantial improvements for the languages not covered in pretrained models. Finally, we show a successful application of the ensembled system on a truly low-resource

¹The code and data are available at: <https://github.com/Jivnesh/SanDP>.

language Sanskrit. We find that the ensembled system outperforms the state-of-the-art system (Krishna et al., 2020a) for Sanskrit by 1.2 points absolute gain in terms of UAS and shows comparable performance in terms of LAS (§ 3).

2 Investigation of Strategies Tailored for Low-resource Settings

We explore 5 strategies specially tailored for low-resource settings on 7 languages, and integrate the best performing strategy of each category in our ensembled system (Table 1). We utilize Dozat and Manning (2017) as a base system for all the experiments, henceforth referred to as BiAFF.

Language selection criteria: We choose low-resource languages with less than 2,500 training samples from 4 different typological families such that each language belongs to a unique sub-family. In order to accommodate a low-resource tailored pretraining (Sandhan et al., 2021), we choose languages that have explicit morphological information. Additionally, we divide the set of languages into the languages covered/not-covered in the multilingual language model’s pretraining: (1) Covered: Arabic (ar), Greek (el), Hungarian (hu) (2) Not covered: Wolof (wo), Gothic (got), Coptic (cop) and Sanskrit (san).




Figure 1: Plot for number of training samples vs. UD languages available in UD-2.6.

Figure 1 illustrates the number of training samples available for all UD languages in UD-2.6. Hedderich et al. (2021) ask “How low is low-resource?” and suggest that the threshold of low-resource is task and language-dependent. The low-resource settings can be seen as a continuum of resource availability due to the absence of a hard threshold. More efforts should focus on evaluating low-resource strategies across multiple languages for

a fair comparison between these strategies. Therefore, we select a threshold for the languages with less than 2,500 training samples (Figure 1). We restrict ourselves to the setting where the target low-resource language does not have a high-resource related language that could possibly facilitate the positive cross-lingual or zero-shot transfer (Vulić et al., 2019; Pires et al., 2019; Søgaard et al., 2018; de Lhoneux et al., 2018; Smith et al., 2018). Thus, we do not consider low-resource languages with only a test set available. Also, we do not consider cross-lingual transfer (Duong et al., 2015; Ahmad et al., 2019; Vania et al., 2019; Das and Sarkar, 2020) strategy in our study.

Dataset and metric: For each of these 7 low-resource languages from Universal Dependencies (UD-2.6) (de Marneffe et al., 2021), following Rotman and Reichart (2019), we use 500 data points for training (allocating equal power to each language for fair comparison) and the original dev/test split as dev/test set. Additionally, 1000 morphologically tagged data points (without dependency annotations) are used for self-training and pretraining. We use sentence level macro averaged UAS/LAS metric for evaluation.

Hyper-parameters: For SeqTraL variants, we use the exact same encoder as Ma et al. (2018) with 2 Bi-LSTM layers and decoder with fully connected layer followed by softmax layer. For the ensembled system, we adopt BiAFF’s codebase by Ma et al. (2018) with the hyper-parameters setting as follows: the batch size as 16, training iterations as 100, a dropout rate as 0.33, the number of stacked Bi-LSTM layers as 2, learning rate as 0.002 and the remaining parameters as the same as Ma et al. (2018). We release our codebase publicly under creative-common licence.

Computing Infrastructure Used: We primarily use RTX-2080, 12 GB GPU memory, 4352 GPU Cores computing infrastructure for our experiments.

Sequential Transfer Learning (SeqTraL): Following Sandhan et al. (2021), we pretrain three encoders (similar to BiAFF) on three sequence labelling auxiliary tasks and integrate them with the BiAFF encoder using a gating mechanism. We adapt these pretrained encoders with various optimization schemes, proposed for reducing a catastrophic forgetting (French, 1999; McCloskey and

Strategy	Model	el		ar		hu		got		cop		wo		san	
		UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS	UAS	LAS
	BiAFF	86.61	82.23	79.89	73.13	80.51	75.11	75.66	69.42	87.33	84.48	82.63	77.83	75.47	65.77
Data aug.	Cropping	85.98	81.76	79.42	73.03	77.46	71.62	75.72	68.25	86.05	82.95	80.38	75.39	73.82	63.46
	Rotation	86.39	82.19	79.22	72.79	77.33	71.16	76.19	68.89	86.27	83.27	80.66	76.05	75.58	65.14
	Nonce	87.43	82.60	79.52	73.00	79.56	69.50	76.42	67.74	87.64	83.52	81.97	76.37	77.25	66.30
Pretraining	BiAFF+mBERT	91.41	87.89	83.50	76.30	85.20	77.50	64.20	53.20	33.70	15.60	71.50	61.40	71.40	55.12
	BiAFF+XLM-R	93.61	90.85	86.04	79.55	89.05	83.80	-	-	-	-	-	-	78.43	66.72
	BiAFF+LCM	89.00	85.83	82.49	76.67	83.22	78.49	79.88	74.65	88.79	86.02	85.85	81.66	81.63	73.86
Self-training	SelfTrain	86.78	82.25	80.86	74.45	80.62	75.09	76.96	70.15	87.95	85.33	83.83	78.80	77.53	66.59
	CVT	80.53	77.37	76.21	71.87	75.21	70.01	69.43	63.59	79.32	74.21	73.21	69.50	69.21	56.21
	DCST	88.26	84.09	82.21	75.78	82.85	77.65	79.52	72.91	88.85	85.53	85.51	80.71	78.55	69.10
SeqTraL	SeqTraL-FE	88.43	85.43	81.86	76.60	82.97	78.46	80.15	75.24	88.08	85.57	85.61	81.77	81.20	73.70
	SeqTraL-UF	88.50	85.36	82.52	76.79	83.83	79.24	80.79	75.65	88.87	86.30	85.78	81.54	81.51	73.65
	SeqTraL-DL	89.06	85.88	82.57	76.66	83.36	78.57	80.29	74.89	88.78	86.14	86.25	81.85	81.17	73.10
	SeqTraL-FT	88.80	85.47	82.66	76.83	83.79	78.95	80.13	75.11	88.86	86.31	86.03	81.64	81.84	73.94
Multi-tasking	MTL-Case	86.73	82.47	80.49	74.08	80.73	75.52	-	-	86.45	83.82	82.86	77.46	76.15	65.36
	MTL-Label	86.13	81.55	79.86	72.72	80.07	73.92	75.52	69.30	87.44	84.62	83.08	77.94	76.02	65.20
	MTL-Morph	86.30	82.23	80.02	73.55	80.49	74.70	77.33	71.05	87.00	84.22	83.25	78.75	76.71	66.69
Prop. system	BiAFF	86.61	82.23	79.89	73.13	80.51	75.11	75.66	69.42	87.33	84.48	82.63	77.83	75.47	65.77
	+Pretraining	93.61	90.85	86.04	79.55	89.05	83.80	79.88	74.65	88.79	86.02	85.85	81.66	81.63	73.86
	+MTL	89.99	86.49	82.47	76.24	84.35	79.74	80.33	75.15	88.42	85.94	85.91	81.56	81.30	73.49
	+SeqTraL	90.31	86.70	82.70	76.57	84.58	80.15	80.79	75.65	88.87	86.30	86.05	81.85	81.84	73.94
	+Self-training	89.83	86.09	82.08	75.92	84.12	79.66	80.08	75.24	88.78	86.07	85.73	81.77	79.89	72.28
+Data. aug.	89.11	85.87	82.08	75.92	84.12	79.66	79.56	73.53	88.31	84.67	85.73	81.77	79.52	71.89	
Evaluation on test set	BiAFF	87.10	83.06	80.92	75.02	80.31	74.16	77.73	70.72	88.50	85.32	80.92	75.02	79.33	67.92
	Prop. system	93.66	90.68	86.43	79.88	88.50	82.67	82.52	77.07	89.31	86.38	87.50	82.95	83.59	74.83

Table 1: Evaluation of low-resource strategies on 7 languages. Experiments are first performed on dev set to find best performing combination of strategies for each language. The best results from strategies from each family are bold and statistically significant compared to its peer baselines belonging to the same family as per t-test ($p < 0.01$). The second last block shows ablations when the best variant from each family is added to the ensembled system. For example, +Data. aug. refers to the system with the best variant from all 5 strategies. The best performing system as per dev set is finally compared with BiAFF on the test set. XLM-R is not compatible with 3 languages and case information of Gothic (*got*) language is missing; hence we do not report their results.

Cohen, 1989). **SeqTraL-FE**: We treat newly integrated layers as Feature Extractors (FE) by freezing them. **SeqTraL-UF**: Gradually Unfreeze (UF) these new layers in the top to down order (Howard and Ruder, 2018; Felbo et al., 2017). **SeqTraL-DL**: The discriminative learning rate (DL) is used for newly added layers (Howard and Ruder, 2018), the learning rate is decreased from top-to-bottom layers. **SeqTraL-FT**: The default learning rate is used to fine-tune all newly added layers.

Cross/mono-lingual Pretraining: We experiment with two multilingual pretrained models, namely, the multilingual BERT (Devlin et al., 2019, mBERT) based system (Kondratyuk and Straka, 2019) and the XLM-Roberta (Conneau et al., 2020, XLM-R) based system (Nguyen et al., 2021). We also consider supervised pretraining specially tailored for low-resource dependency parsing (Sandhan et al., 2021, LCM) which essentially combines three sequence labelling auxiliary tasks. We pre-train it on 1,000 morphologically tagged data points

without dependency annotations.

Self-training: Another line of modelling focuses on self-training (Goldwasser et al., 2011; Clark et al., 2018; Rybak and Wróblewska, 2018) to overcome the bottleneck of task-specific labelled data. Earlier attempts failed to prove effectiveness of self-training for dependency parsing (Rush et al., 2012). However, Clark et al. (2018, CVT) and Roman and Reichart (2019, DCST), show successful application, thus, we consider these two systems. Also, we generate dependency data by applying a pretrained BiAFF system on 1000 unlabelled data points. We augment this predicted data with gold data and retrain BiAFF in **Self-Train** setting.

Multi-task Learning: We simultaneously train BiAFF and a sequence labelling based auxiliary task in a multi-task setting (MTL). We experiment with the following auxiliary tasks: prediction of the morphological label (MTL-Morph), dependency relation between a word and its head (MTL-Label) and the case label (MTL-Case).

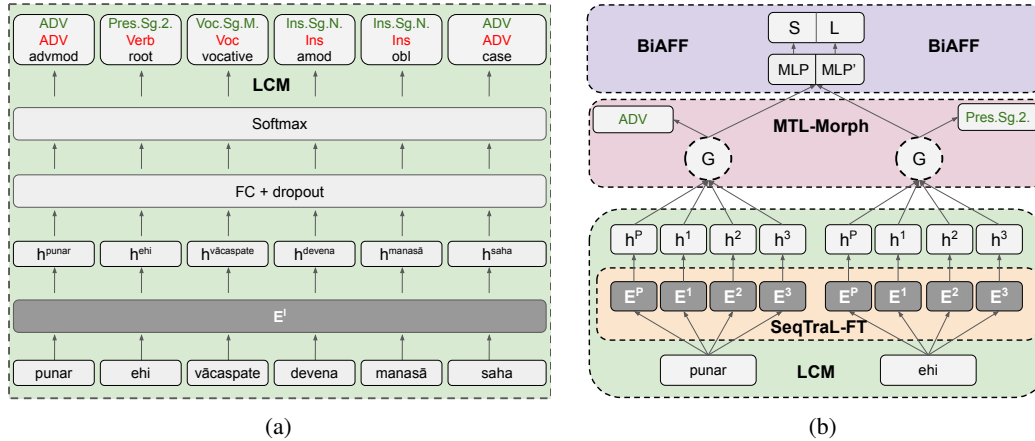


Figure 2: The ensembled system for Sanskrit. Translation: “Oh Vācas pate! Come again with divine mind”.

Data Augmentation: Şahin and Steedman (2018) introduce **Cropping**: delete some parts of a sentence to create multiple short meaningful sentences, and **Rotation**: permute the siblings of headword restricted to a set of relations. Both operations modify a set of words or configurational information; however, they do not change the dependencies. **Nonce:** Gulordava et al. (2018) propose to create nonce sentences by substituting a few words which share the same syntactic labels. For each variant, we use additional 1,000 augmented data points.

Results on multilingual experiments: Table 1 first reports results of all 5 strategies on dev set of 7 languages. Next, the second last block of Table 1 (Prop. system) shows ablations on dev set where the best variant from each family is gradually added into the ensembled system. For example, +*Data.aug.* row refers to the system with the best variant from all 5 strategies. Finally, the best performing system as per dev set is compared with BiAFF on the test set. We observe that (1) the best performing variant from augmentation, SeqTraL and MTL families is language dependent. (2) DCST variant of self-training wins over its peer for all the languages. (3) XLM-R outperforms for the languages which are covered in its pretraining (except Sanskrit²) and LCM outperforms for the rest of the languages which are truly low-resource. (4) Notably, we find effective generalization ability of the proposed approach on languages covered in cross-lingual pretraining (only pretraining helps) and for the rest of the languages (pretraining, MTL and SeqTraL helps).

²Maybe due to limited coverage of corpus for Sanskrit.

3 Application on Sanskrit

Data: We use two standard benchmark datasets available for Sanskrit. We use 1,700, 1,000 and 1,300 sentences (prose domain) from the Sanskrit Treebank Corpus (Kulkarni et al., 2010, STBC) as train, dev and test set, respectively. We also evaluate on the Vedic Sanskrit Treebank (Hellwig et al., 2020, VST) consisting of 1,500, 1,024 and 1,473 sentences (poetry-prose mixed) as train, dev and test data, respectively. For both data, the final results on the test set are reported using systems trained with combined gold train and dev set.

Baselines: We use More et al. (2019b, YAP) and Chang et al. (2016, L2S) from transition-based dependency parsing family. Dozat and Manning (2017, BiAFF) is a graph-based approach with BiAFFINE attention mechanism. Krishna et al. (2020a, MG-EBM) extends Krishna et al. (2020b, Tree-EBM-F) using multi-graph formulation. Systems marked with (*) are hybrid systems which leverage linguistic rules from Pāṇini.

The ensembled system: Figure 2 shows the ensembled system for Sanskrit as per Table 1. It consists of two steps, namely, pretraining (LCM) and integration. As shown in Figure 2a, LCM pretrains three encoders $E^{(1)-(3)}$ using three independent auxiliary tasks, namely, morphological label prediction, case label prediction and relation label prediction. Thereafter, as shown in Figure 2b, these pretrained encoders are integrated with the BiAFF encoder $E^{(P)}$ using a gating mechanism as employed in Sato et al. (2017). We use SeqTraL-FT optimization scheme to update the weights

- Kevin Clark, Minh-Thang Luong, Christopher D. Manning, and Quoc Le. 2018. [Semi-supervised sequence modeling with cross-view training](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1914–1925, Brussels, Belgium. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Ayan Das and Sudeshna Sarkar. 2020. [A survey of the model transfer approaches to cross-lingual dependency parsing](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 19(5).
- Miryam de Lhoneux, Johannes Bjerva, Isabelle Augenstein, and Anders Søgaard. 2018. [Parameter sharing between dependency parsers for related languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4992–4997, Brussels, Belgium. Association for Computational Linguistics.
- Marie-Catherine de Marneffe, Christopher D. Manning, Joakim Nivre, and Daniel Zeman. 2021. [Universal Dependencies](#). *Computational Linguistics*, 47(2):255–308.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- 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.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 845–850.
- Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. [Transition-based dependency parsing with stack long short-term memory](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 334–343, Beijing, China. Association for Computational Linguistics.
- Bjarke Felbo, Alan Mislove, Anders Søgaard, Iyad Rahwan, and Sune Lehmann. 2017. [Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1615–1625, Copenhagen, Denmark. Association for Computational Linguistics.
- Robert M French. 1999. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135.
- Dan Goldwasser, Roi Reichart, James Clarke, and Dan Roth. 2011. [Confidence driven unsupervised semantic parsing](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1486–1495, Portland, Oregon, USA. Association for Computational Linguistics.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. [Colorless green recurrent networks dream hierarchically](#). 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 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Michael A. Hedderich, Lukas Lange, Heike Adel, Janik Strötgen, and Dietrich Klakow. 2021. [A survey on recent approaches for natural language processing in low-resource scenarios](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2545–2568, Online. Association for Computational Linguistics.
- Oliver Hellwig, Salvatore Scarlata, Elia Ackermann, and Paul Widmer. 2020. [The treebank of vedic Sanskrit](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 5137–5146, Marseille, France. European Language Resources Association.
- 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.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and accurate dependency parsing using bidirectional LSTM feature representations](#). *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Dan Kondratyuk and Milan Straka. 2019. [75 languages, 1 model: Parsing Universal Dependencies universally](#). 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 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Amrith Krishna, Ashim Gupta, Deepak Garasangi, Pavankumar Satuluri, and Pawan Goyal. 2020a. [Keep it surprisingly simple: A simple first order graph based parsing model for joint morphosyntactic parsing in Sanskrit](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4791–4797, Online. Association for Computational Linguistics.
- Amrith Krishna, Bishal Santra, Ashim Gupta, Pavankumar Satuluri, and Pawan Goyal. 2020b. [A graph based framework for structured prediction tasks in sanskrit](#). *Computational Linguistics*, 46(4):1–63.
- Amba Kulkarni, Sheetal Pokar, and Devanand Shukl. 2010. [Designing a constraint based parser for sanskrit](#). In *International Sanskrit Computational Linguistics Symposium*, pages 70–90. Springer.
- Artur Kulmizev, Miryam de Lhoneux, Johannes Gontrum, Elena Fano, and Joakim Nivre. 2019. [Deep contextualized word embeddings in transition-based and graph-based dependency parsing - a tale of two parsers revisited](#). 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 2755–2768, Hong Kong, China. Association for Computational Linguistics.
- Xuezhe Ma, Zecong Hu, Jingzhou Liu, Nanyun Peng, Graham Neubig, and Eduard Hovy. 2018. [Stack-pointer networks for dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1403–1414, Melbourne, Australia. Association for Computational Linguistics.
- Michael McCloskey and Neal J. Cohen. 1989. [Catastrophic interference in connectionist networks: The sequential learning problem](#). In Gordon H. Bower, editor, *Psychology of Learning and Motivation*, volume 24, pages 109 – 165. Academic Press.
- Amir More, Amit Seker, Victoria Basmova, and Reut Tsarfaty. 2019a. [Joint Transition-Based Models for Morpho-Syntactic Parsing: Parsing Strategies for MRLs and a Case Study from Modern Hebrew](#). *Transactions of the Association for Computational Linguistics*, 7:33–48.
- Amir More, Amit Seker, Victoria Basmova, and Reut Tsarfaty. 2019b. [Joint transition-based models for morpho-syntactic parsing: Parsing strategies for MRLs and a case study from Modern Hebrew](#). *Transactions of the Association for Computational Linguistics*, 7:33–48.
- Dat Quoc Nguyen and Karin Verspoor. 2018. [An improved neural network model for joint POS tagging and dependency parsing](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 81–91, Brussels, Belgium. Association for Computational Linguistics.
- Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, 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.
- Guy Rotman and Roi Reichart. 2019. [Deep contextualized self-training for low resource dependency parsing](#). *Transactions of the Association for Computational Linguistics*, 7:695–713.
- Sebastian Ruder, Matthew E. Peters, Swabha Swayamdipta, and Thomas Wolf. 2019. [Transfer learning in natural language processing](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Tutorials*, pages 15–18, Minneapolis, Minnesota. Association for Computational Linguistics.

- Alexander Rush, Roi Reichart, Michael Collins, and Amir Globerson. 2012. [Improved parsing and POS tagging using inter-sentence consistency constraints](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1434–1444, Jeju Island, Korea. Association for Computational Linguistics.
- Piotr Rybak and Alina Wróblewska. 2018. Semi-supervised neural system for tagging, parsing and lematization. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual parsing from raw text to universal dependencies*, pages 45–54.
- Gözde Gül Şahin and Mark Steedman. 2018. [Data augmentation via dependency tree morphing for low-resource languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5004–5009, Brussels, Belgium. Association for Computational Linguistics.
- Jivnesh Sandhan, Amrith Krishna, Ashim Gupta, Laxmidhar Behera, and Pawan Goyal. 2021. [A little pretraining goes a long way: A case study on dependency parsing task for low-resource morphologically rich languages](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 111–120, Online. Association for Computational Linguistics.
- Motoki Sato, Hitoshi Manabe, Hiroshi Noji, and Yuji Matsumoto. 2017. [Adversarial training for cross-domain Universal Dependency parsing](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 71–79, Vancouver, Canada. Association for Computational Linguistics.
- Aaron Smith, Bernd Bohnet, Miryam de Lhoneux, Joakim Nivre, Yan Shao, and Sara Stymne. 2018. [82 treebanks, 34 models: Universal Dependency parsing with multi-treebank models](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. [On the limitations of unsupervised bilingual dictionary induction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788, Melbourne, Australia. Association for Computational Linguistics.
- Clara Vania, Yova Kementchedjheva, Anders Søgaard, and Adam Lopez. 2019. [A systematic comparison of methods for low-resource dependency parsing on genuinely low-resource languages](#). 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 1105–1116, Hong Kong, China. 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.
- Daniel Zeman, Jan Hajič, 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, Brussels, Belgium. Association for Computational Linguistics.