

Hexatagging: Projective Dependency Parsing as Tagging


Afra Amini* Tianyu Liu* Ryan Cotterell

{afra.amini, tianyu.liu, ryan.cotterell}@inf.ethz.ch

ETH zürich

Abstract

We introduce a novel dependency parser, the hexatagger, that constructs dependency trees by tagging the words in a sentence with elements from a *finite* set of possible tags. In contrast to many approaches to dependency parsing, our approach is fully parallelizable at training time, i.e., the structure-building actions needed to build a dependency parse can be predicted in parallel to each other. Additionally, exact decoding is linear in time and space complexity. Furthermore, we derive a probabilistic dependency parser that predicts hexatags using no more than a linear model with features from a pretrained language model, i.e., we forsake a bespoke architecture explicitly designed for the task. Despite the generality and simplicity of our approach, we achieve state-of-the-art performance of 96.4 LAS and 97.4 UAS on the Penn Treebank test set. Additionally, our parser’s linear time complexity and parallelism significantly improve computational efficiency, with a roughly 10-times speed-up over previous state-of-the-art models during decoding.

 <https://github.com/rycolab/parsing-as-tagging>

1 Introduction

The combination of parallel computing hardware and highly parallelizable neural network architectures (Vaswani et al., 2017) has enabled the pretraining of language models on increasingly large amounts of data. In order to apply pretrained language models to downstream NLP tasks, many practitioners fine-tune the pretrained model while the task-specific architecture is jointly trained from scratch. Typically, the task-specific architecture is built upon the hidden representations generated by the final layer of a pretrained model. Exploiting pretrained language models in this manner has boosted the performance considerably on many NLP tasks (Devlin et al., 2019; Clark et al., 2020;

*Equal contribution.

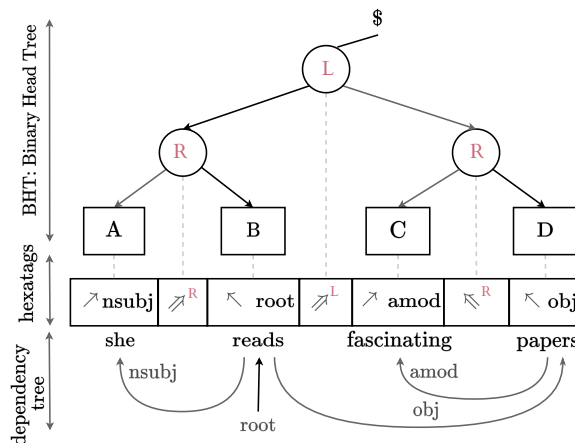


Figure 1: From bottom to top, the figure shows the dependency tree, the hexatags, and the binary head tree for the sentence “She reads fascinating papers.”

Aghajanyan et al., 2021). However, for the end-to-end fine-tuning process to be fully parallelizable, it is also necessary to parallelize the training of the task-specific architecture. Unfortunately, due to the complexity of the output in many structured prediction tasks in natural language, e.g., in dependency parsing, state-of-the-art models still use architectures with limited parallelization during training (Mrini et al., 2020; Yang and Tu, 2022).

In an attempt to develop parsers parallelizable during training, a recent line of work recasts parsing as tagging (Li et al., 2018; Strzyz et al., 2019; Kitaev and Klein, 2020; Amini and Cotterell, 2022). Under this approach, a parse tree is linearized into a sequence of tags.¹ The benefit of such a paradigm is that tagging can be done by only adding a linear classifier on top of a pretrained language model and the tags can, thus, be predicted independently. This leads to a parser that is highly parallelizable and whose training can be easily harmonized with the (parallelizable) fine-tuning of pretrained language models. During decoding, an exact algorithm is used to recover a valid sequence

¹In some tagging-based dependency parsers, the cardinality of the set of tags even grows as a function of the length of the input sequence and, thus, is unbounded.

of tags which is then converted back to a parse tree.

Kitaev and Klein (2020) were the first to propose a parsing-as-tagging scheme with a constant tag space for *constituency parsing* and, additionally, the first to achieve results competitive with the state-of-the-art non-parallelizable constituency parsers using such a tagger. However, for dependency parsing, all dependency parsing-as-tagging schemes in the literature (Li et al., 2018; Strzyz et al., 2019; Vacareanu et al., 2020) have infinite tag sets whose cardinality grows with the length of the input sequence, which limits such parsers’ efficiency and generality (Strzyz et al., 2019). Moreover, in some cases, this growth hinders generalization to sentences longer than the longest training sentence. Furthermore, tagging-based dependency parsers still do not perform competitively with the best-performing parsers in the literature (Li et al., 2018).

In this paper, we propose a novel way of framing projective dependency parsing as a tagging task. Our approach makes use of 6 distinct tags, motivating us to naming the scheme **hexatagger**. In our experiments, hexatagger achieves state-of-the-art performance on the English Penn Treebank (PTB; Marcus et al., 1993) test set. Notably, it outperforms parsers with more computationally expensive training procedures and extra constituency annotations, e.g., the parser developed by Mrini et al. (2020). Furthermore, hexatagger achieves results competitive to Yang and Tu’s (2022) parser on the Chinese Penn Treebank (CTB; Xue et al., 2005) test set and 12 languages on the pseudo-projectivized data from the Universal Dependencies (UD2.2; Nivre et al., 2018) benchmark. In terms of efficiency, our experiments suggest that hexatagger is 10 times faster than previous top-performing parsers, and consumes significantly less memory, despite using an exact dynamic program for decoding.

2 Hexatagging

In this section, we introduce hexatagging, a tagging scheme that consists of 6 unique tag types. We further prove by construction that there exists an injective mapping between valid sequences of hexatags and dependency trees.

2.1 Binary Head Trees

Before going into the details on how to represent dependency trees with a sequence of tags, we introduce **binary head trees** (BHTs), a simple formalism that serves as a useful intermediary between

dependency trees and sequence of hexatags. Intuitively, a BHT is a special form of a constituency tree where each internal node is either labeled \textcircled{L} when the head of the derived constituent is in the left subtree or \textcircled{R} when the head is in the right subtree. See Fig. 1 for a visual depiction of a BHT. In the next theorem, we formally state the relationship between the set of dependency trees and BHTs.

Theorem 1. *There exists a bijective² function that maps every projective dependency tree to a BHT.*

In the following two paragraphs, we sketch a construction that such a function exists, i.e., we describe how to map any dependency tree to a BHT and then how to map back any BHT to a dependency tree and back again.

Projective Dependency Trees to BHTs. To convert a dependency tree to a BHT, we start from the root and do a depth-first traversal of the dependency tree. To avoid spurious ambiguity (Eisner and Satta, 1999), we canonically order arcs of the tree by processing the arcs left to right and inside out.³ Algorithmically, converting a dependency tree to a BHT proceeds as follows. When we first visit a word, we push it onto a stack and proceed with visiting its dependents. When there is no dependent word left to visit, we create a new node (\textcircled{L} or \textcircled{R}) and attach the top two elements in the stack as the left and right child of this node. A step-by-step demonstration of this algorithm is shown in Fig. 2 and pseudocode is provided in Alg. 1.

BHTs to Projective Dependency Trees. To convert a BHT back to the dependency tree we follow Alg. 2. Algorithmically, we process BHT in a depth-first fashion. Upon visiting \textcircled{R} or \textcircled{L} nodes, we combine the top two elements in the stack by creating a dependency arc between them. The direction of the arc is determined by the label of the node (\textcircled{R} or \textcircled{L}). See Fig. 3 for an example.

Once the dependency tree is converted to a BHT, we can linearize it to a sequence of hexatags in a straightforward manner. Theorem 2 states the relationship between BHTs and hexatags formally.

Theorem 2. *There exists a total and injective function that maps every BHT to a valid hexatag sequence, i.e., in other words, every BHT can be*

²We remark that the bijectivity follows from a canonical ordering (left-to-right and inside-out) of a node’s dependents.

³One can also process the right arcs first. In our experiments, however, we observed no significant difference in the performance of the parser, see App. C for more analysis.

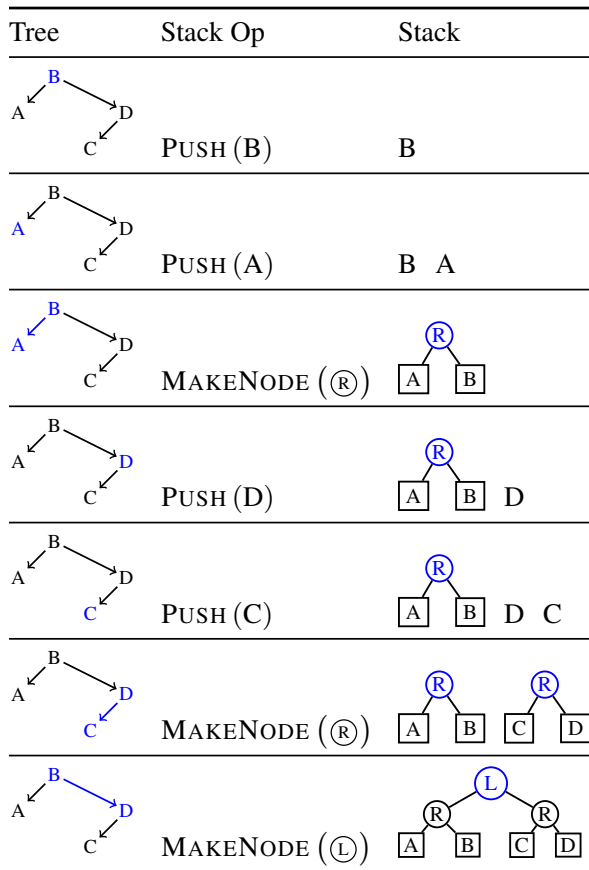


Figure 2: The example shows how to derive the BHT for a dependency tree $\hat{A} \hat{B} \hat{C} \hat{D}$. The top of the stack is on the right.

mapped to a unique hexatag sequence. However, some hexatag sequences do not correspond to BHTs, i.e., the function is not surjective.

In the following subsections, we prove by construction that such a function exists. Throughout the rest of the paper, we refer to those hexatagging sequences that *do* correspond to BHTs as **valid**.

2.2 From BHT to Hexatags

To transform a given BHT to a sequence of hexatags, we enumerate the action sequence that a left-corner shift–reduce parser would take when parsing this BHT (Johnson, 1998). Left-corner parsers have actions that align more closely with the input sequence than top-down or bottom-up shift–reduce actions and, thus, offer a better linearization for tagging tasks (Amini and Cotterell, 2022). A simple explanation of this linearization process is given by Kitaev and Klein (2020, §3.1). Their algorithm involves an in-order traversal of the tree. Upon visiting each node, we generate a tag that includes the direction of the arc that attaches

the node to its parent, i.e., whether that node is a left or a right child of its parent, and the label of the node. When traversing a BHT, this paradigm results in 6 distinct tag types:

- \prec : this terminal node is the right child of its parent;
- \succ : this terminal node is the left child of its parent;
- \bowtie^R (\bowtie^L): this non-terminal node is the right child of its parent and the head of the corresponding constituent is on the right (respectively, left) subtree;
- $\not\bowtie^R$ ($\not\bowtie^L$): this non-terminal node is the left child of its parent and the head of the corresponding constituent is on the right (respectively, left) subtree.

For an input sequence $w = w_1 \cdots w_N$, this process gives us a hexatag sequence of length $2N - 1$. Fig. 1 depicts tree-to-tags transformation through an example.

Labeled Dependency Trees. When converting a *labeled* dependency tree to a sequence of hexatags, the arc labels must be encoded in the tags. Therefore, while reading a terminal node, we concatenate the label of the arc that connects the node to its parent with the hexatag. In this case, the number of distinct tags would be $\mathcal{O}(|\mathcal{A}|)$, where $|\mathcal{A}|$ is the number of unique arc labels. For example, in Fig. 1 the hexatag generated while processing *she* is: $\langle \succ, \text{nsubj} \rangle$.

2.3 From Hexatags to Dependency Trees

To transform a sequence of hexatags back to a dependency tree, we again go through a two-step process. First, we again interpret hexatags as actions in a left-corner shift–reduce transition system to construct a BHT. The actions in such a transition system are as follows:

- \succ : shift the leaf node into the stack;
- $\not\bowtie^R$ ($\not\bowtie^L$): create a new node labeled \textcircled{R} (respectively, \textcircled{L}), attach the top element in the stack as its left child, and attach a dummy node as its right child (\emptyset in step 2 in Fig. 3);
- \prec : pop the subtree on the top of the stack. Replace the dummy node in the subtree with the terminal node. Push the subtree back to the stack;
- \bowtie^R (\bowtie^L): create a new node labeled \textcircled{R} (respectively, \textcircled{L}). Pop the top element of the stack, attach it as the new node’s left child, and set a dummy node as the node’s right child. Pop another subtree of the stack, identify the dummy

node in the subtree and replace it with the newly created subtree. Push the subtree back to the stack (step 6 in Fig. 2);

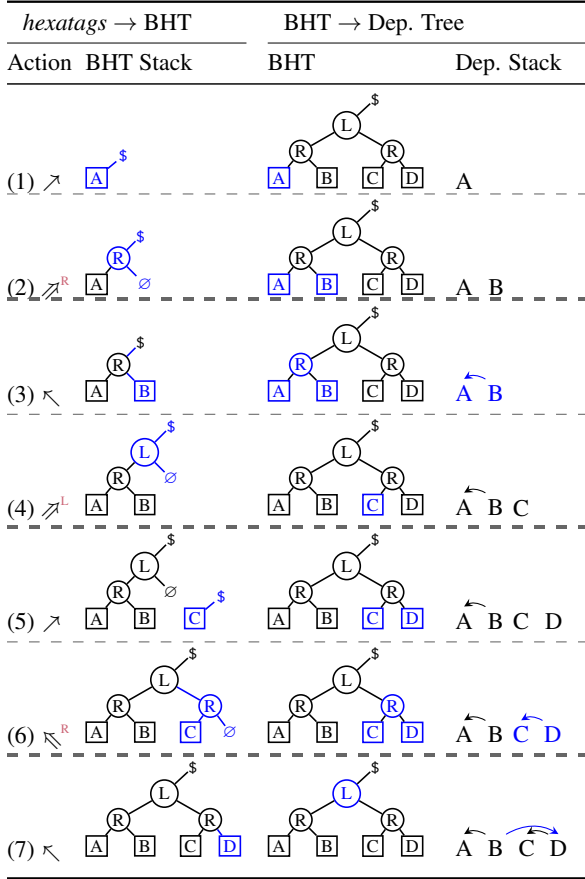


Figure 3: The example shows how to derive BHT from hexatags and how to transform BHT to a dependency tree. The top of the stacks is on the right.

3 Probability Model

In this section, we explain how to predict hexatags in parallel. Our tagging model predicts two hexatags for each word in the input sequence with the exception of that last word, for which we only predict one tag. As discussed in §2.1, a hexatagger produces a sequence of $2N - 1$ tags $\mathbf{t} = [t_1, t_2, \dots, t_{2N-1}]$ for an input sequence of length N , $\mathbf{w} = w_1 w_2 \dots w_N$. Therefore, an intuitive way to match the tag sequence with the input sequence is to assign two tags to each word. We denote a training corpus \mathcal{S} of M tuples of input sequences and tag sequences $\{(\mathbf{w}^m, \mathbf{t}^m)\}_{m=1}^M$.

To learn the scoring function over tags, we follow the same independence assumption as in (Kitaev and Klein, 2020), i.e., the probability of predicting each tag is independent of other tags given the input sequence. This assumption barely harms model performance (see Amini and

Cotterell, 2022, Table 3), but significantly speeds up the training process by enabling each tag to be predicted in parallel, and complexity reduces by a factor of $\mathcal{O}(N)$. The training objective is to minimize the negative log-likelihood of the gold-standard tag sequences, i.e.

$$\mathcal{L}(\theta) = - \sum_{(\mathbf{w}, \mathbf{t}) \in \mathcal{S}} \log p_{\theta}(\mathbf{t} | \mathbf{w}) \quad (1a)$$

$$= - \sum_{(\mathbf{w}, \mathbf{t}) \in \mathcal{S}} \log \prod_{n=1}^{2N-1} p_{\theta}(t_n | \mathbf{w}) \quad (1b)$$

$$= - \sum_{(\mathbf{w}, \mathbf{t}) \in \mathcal{S}} \left(\sum_{n=1}^N \log p_{\theta}(t_{2n-1} | \mathbf{w}) \right. \\ \left. + \sum_{n=1}^{N-1} \log p_{\theta}(t_{2n} | \mathbf{w}) \right) \quad (1c)$$

where θ refers collectively to the parameters of the two linear projections and the parameters of the pretrained model. To obtain $p_{\theta}(t_{2n} | \mathbf{w})$ and $p_{\theta}(t_{2n+1} | \mathbf{w})$, we apply two independent linear projections on the contextualized representation of w_n ⁴ given by a pretrained model and convert that to a probability distribution using softmax.

4 Decoding

Our goal in this section is to develop an efficient algorithm to find the highest-scoring hexatag sequence under the model developed in §3. As stated in Theorem 2, the transformation function between BHTs and hexatag sequences is not surjective, i.e., not all the tag sequences can be transformed back into a BHT. Therefore, we need to find a *valid* hexatag sequence with the maximum probability under the model that *can* be transformed back to a BHT. Once such hexatag sequence is found, we can follow the two-step algorithm described in §2.3 to obtain the corresponding dependency tree.

To find the highest-scoring valid hexatag sequence, we follow the linear-time algorithm developed by Kitaev and Klein (2020). For a hexatag sequence to be valid, we should be able to interpret it as actions in a left-corner shift–reduce transitions system, described in §2.3. Concretely:

- The first action can only be \nearrow because other actions need at least one item in the stack;
- The actions \nwarrow^L , \nwarrow^R can only be performed if there is at least two items in the stack;

⁴If a word consists of more than one subword, we apply the projection to the last subword.

	bg	ca	cs	de	en	es	fr	it	nl	no	ro	ru	Avg.
Zhang et al. (2020)	90.77	91.29	91.54	80.46	87.32	90.86	87.96	91.91	88.62	91.02	86.90	93.33	89.33
Wang and Tu (2020)	90.53	92.83	92.12	81.73	89.72	92.07	88.53	92.78	90.19	91.88	85.88	92.67	90.07
+BERT _{multilingual}													
Wang and Tu (2020)	91.30	93.60	92.09	82.00	90.75	92.62	89.32	93.66	91.21	91.74	86.40	92.61	90.61
Dozat and Manning (2017)	90.30	94.49	92.65	85.98	91.13	93.78	91.77	94.72	91.04	94.21	87.24	94.53	91.82
Yang and Tu (2022)	91.10	94.46	92.57	85.87	91.32	93.84	91.69	94.78	91.65	94.28	87.48	94.45	91.96
Hexatagger	92.87	93.79	92.82	85.18	90.85	93.17	91.50	94.72	91.89	93.95	87.54	94.03	91.86

Table 1: LAS scores on the test sets of 12 languages in UD 2.2. Hexatagger achieves competitive performance in all languages and is state-of-the-art in 4 languages.

Model	PTB		CTB	
	UAS	LAS	UAS	LAS
Zhou and Zhao (2019)*	97.0	95.4	91.2	89.2
Mrini et al. (2020)*	97.4	96.3	94.6	89.3
Chen and Manning (2014)	91.8	89.6	83.9	82.4
Dozat and Manning (2017)	95.7	94.1	89.3	88.2
Yang and Tu (2022)#	97.4	95.8	93.5	92.5
Hexatagger	97.4	96.4	93.2	91.9

Table 2: Results on PTB and CTB. * indicates usage of extra constituency annotation. # is our re-implementation using the same pretrained encoder with hexatagger.

- After performing all the actions, the stack should contain a single element.

The above shows that the validity of a hexatag sequence only depends on the *number* of elements in the stack at each point of the derivation.⁵

5 Experiments

We conduct experiments on the English Penn Treebank (PTB; Marcus et al., 1993), the Chinese Penn Treebank (CTB; Xue et al., 2005), and the Universal Dependencies 2.2 (UD2.2; Nivre et al., 2018). For UD2.2, we adopt the pseudo-projective transformation (Nivre and Nilsson, 2005) to convert non-projective trees into projective trees following previous work (Wang and Tu, 2020; Yang and Tu, 2022). We report dataset statistics in App. E and hyperparameter settings in App. F.

Accuracy. We train the hexatagger model based on XLNet (Yang et al., 2019) and report the results on PTB and CTB in Table 2. Furthermore, we eval-

⁵Specifically, The decoding algorithm can be thought of as constructing a lattice where each node corresponds to the number of elements in the stack for each transition step ($N \times d$ nodes for maximum stack size of d , $d \leq N$). Each transition corresponds to performing a valid action. The score of the tag at step n is set to the negative log probability $-\log p_{\theta}(t_n | \mathbf{w})$ of the corresponding hexatag given by the model. Finally, we remark that our decoding algorithm is essentially a shortest-path dynamic program that finds the highest-scoring valid hexatag sequence. See Amini and Cotterell (2022, §5.1) for a deeper discussion of this point.

Sent length	Speed (sent/s) \uparrow		Memory (GB) \downarrow	
	Hexatagger	Biaffine	Hexatagger	Biaffine
32	2916	493	2.9	4.5
64	3011	328	3.0	10.1
128	2649	202	3.7	30.6
256	3270	98	4.5	56.2*
overall	3176	338	3.0	10.6

Table 3: Comparison of parsing speed and memory consumption on PTB test set. Results are averaged over 3 random runs on the same server with one Nvidia A100-80GB GPU using BERT-large as encoder. We use a batch size of 128 sentences, except for * that uses 64, which otherwise results in an out-of-memory error.

uate hexatagger in a set of 12 topologically diverse languages on UD corpus, where we use Multilingual BERT (Devlin et al., 2019) as the underlying model (see Table 1). In PTB, we observe that hexatagger achieves state-of-the-art results, compared to models with custom architectures and even in some cases with extra annotation. In CTB and UD, hexatagger follows the best performance closely.

Efficiency. We compare the efficiency of hexatagger with biaffine modules,⁶ which are the backbone of many neural graph-based parsers (Kiperwasser and Goldberg, 2016; Dozat and Manning, 2017; Mrini et al., 2020; Yang and Tu, 2022). As depicted in Table 3, we observe that our hexatagger is an order of magnitude faster and consumes less memory. Further analysis is included in App. C.

6 Conclusion

In summary, hexatagging, our novel scheme, offers a parallelizable and efficiently decodable backbone for dependency parsing. Without relying on custom architecture for dependency parsing, the hexatagger achieves state-of-the-art accuracy on several datasets using no more than a pretrained language model and linear classifiers.

⁶By biaffine model we refer to a slim parameterization of a dependency parser that scores the existence of a dependency between w_i and w_j using a biaffine attention layer over the words' contextualized representations.

Limitations

Non-projectivity. The primary theoretical limitation of hexatagger is that it can only produce projective dependency trees. We would like to explore the possibility of extending hexatagger to non-projective parsing for future work.

Interpretability. As a trade-off for efficiency, hexatagger does not model dependency arcs directly. Compared to graph-based models that explicitly score arc scores between pairs of words, it is more difficult to interpret the output of hexatagger.

Ethics Statement

We do not believe the work presented here further amplifies biases already present in the datasets. Therefore, we foresee no ethical concerns in this work.

Acknowledgments

We would like to thank Tim Vieira for his invaluable feedback throughout the process of this paper. Afra Amini is supported by ETH AI Center doctoral fellowship.

References

- Armen Aghajanyan, Anchit Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021. [Muppet: Massive multi-task representations with pre-finetuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5799–5811, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Afra Amini and Ryan Cotterell. 2022. [On parsing as tagging](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8884–8900, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Srinivas Bangalore and Aravind K. Joshi. 1999. [Supertagging: An approach to almost parsing](#). *Computational Linguistics*, 25(2):237–265.
- Danqi Chen and Christopher Manning. 2014. [A fast and accurate dependency parser using neural networks](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 740–750, Doha, Qatar. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations*.
- Shay B. Cohen and Daniel Gildea. 2016. [Parsing Linear Context-Free Rewriting Systems with Fast Matrix Multiplication](#). *Computational Linguistics*, 42(3):421–455.
- Marie-Catherine de Marneffe and Christopher D. Manning. 2008. [The Stanford typed dependencies representation](#). In *Coling 2008: Proceedings of the workshop on Cross-Framework and Cross-Domain Parser Evaluation*, pages 1–8, Manchester, UK. Coling 2008 Organizing Committee.
- 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, Vancouver, Canada. Association for Computational Linguistics.
- Jason Eisner. 1996. [Efficient normal-form parsing for combinatory categorial grammar](#). In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 79–86, Santa Cruz.
- Jason Eisner and Giorgio Satta. 1999. [Efficient parsing for bilexical context-free grammars and head automaton grammars](#). In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 457–464, College Park, Maryland, USA. Association for Computational Linguistics.
- Mark Johnson. 1998. [Finite-state approximation of constraint-based grammars using left-corner grammar transforms](#). In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 619–623, Montreal, Quebec, Canada. Association for Computational Linguistics.
- Eliyahu Kiperwasser and Miguel Ballesteros. 2018. [Scheduled multi-task learning: From syntax to translation](#). *Transactions of the Association for Computational Linguistics*, 6:225–240.

- 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.
- Nikita Kitaev and Dan Klein. 2020. [Tetra-tagging: Word-synchronous parsing with linear-time inference](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6255–6261, Online. Association for Computational Linguistics.
- Taku Kudo and Yuji Matsumoto. 2002. [Japanese dependency analysis using cascaded chunking](#). In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Marco Kuhlmann, Carlos Gómez-Rodríguez, and Giorgio Satta. 2011. [Dynamic programming algorithms for transition-based dependency parsers](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 673–682, Portland, Oregon, USA. Association for Computational Linguistics.
- Kenton Lee, Mike Lewis, and Luke Zettlemoyer. 2016. [Global neural CCG parsing with optimality guarantees](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2366–2376, Austin, Texas. Association for Computational Linguistics.
- Zuchao Li, Jiaxun Cai, Shexia He, and Hai Zhao. 2018. [Seq2seq dependency parsing](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3203–3214, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Ryan McDonald, Fernando Pereira, Kiril Ribarov, and Jan Hajič. 2005. [Non-projective dependency parsing using spanning tree algorithms](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 523–530, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Ryan McDonald and Giorgio Satta. 2007. [On the complexity of non-projective data-driven dependency parsing](#). In *Proceedings of the Tenth International Conference on Parsing Technologies*, pages 121–132, Prague, Czech Republic. Association for Computational Linguistics.
- Khalil Mrini, Franck Dernoncourt, Quan Hung Tran, Trung Bui, Walter Chang, and Ndapa Nakashole. 2020. [Rethinking self-attention: Towards interpretability in neural parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 731–742, Online. Association for Computational Linguistics.
- Joakim Nivre. 2003. [An efficient algorithm for projective dependency parsing](#). In *Proceedings of the Eighth International Conference on Parsing Technologies*, pages 149–160, Nancy, France.
- Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, John Bauer, Sandra Bellato, Kepa Bengoetxea, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Rogier Blokland, Victoria Bobicev, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Giuseppe G. A. Celano, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilarraza, Carly Dickerson, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Tomaž Erjavec, Aline Etienne, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökirmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Matias Grioni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Radu Ion, Elena Irimia, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Tolga Kayadelen, Václava Kettnerová, Jesse Kirchner, Natalia Kotsyba, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phương Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Nikola Ljubešić, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Mackentanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Gustavo Mendonça, Niko Miekka, Anna Missilä, Cătălin Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Shinsuke Mori, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horñiáček, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lương Nguyễn Thị, Huyền Nguyễn Thị Minh, Vitaly Nikolaev, Rattima Nitisaroj,

- Hanna Nurmi, Stina Ojala, Adédayò Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Siyao Peng, Cene-Augusto Perez, Guy Perrier, Slav Petrov, Jussi Piitulainen, Emily Pitler, Barbara Plank, Thierry Poibeau, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Michael Riebler, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roșca, Olga Rudina, Shoval Sadde, Shadi Saleh, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Muh Shohibussirri, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Zsolt Szántó, Dima Taji, Yuta Takahashi, Takaaki Tanaka, Isabelle Tellier, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uriá, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Veronika Vincze, Lars Wallin, Jonathan North Washington, Seyi Williams, Mats Wirén, Tsegay Woldemariam, Tak-sum Wong, Chunxiao Yan, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, Manying Zhang, and Hanzhi Zhu. 2018. [Universal dependencies 2.2](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Joakim Nivre and Jens Nilsson. 2005. [Pseudo-projective dependency parsing](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 99–106, Ann Arbor, Michigan. Association for Computational Linguistics.
- Michalina Strzyz, David Vilares, and Carlos Gómez-Rodríguez. 2019. [Viable dependency parsing as sequence labeling](#). 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 717–723, Minneapolis, Minnesota. Association for Computational Linguistics.
- Robert Vacareanu, George Caique Gouveia Barbosa, Marco A. Valenzuela-Escárcega, and Mihai Surdeanu. 2020. [Parsing as tagging](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5225–5231, Marseille, France. European Language Resources Association.
- 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 *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Xinyu Wang and Kewei Tu. 2020. [Second-order neural dependency parsing with message passing and end-to-end training](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 93–99, Suzhou, China. Association for Computational Linguistics.
- Naiwen Xue, Fei Xia, Fu-dong Chiou, and Marta Palmer. 2005. [The penn chinese treebank: Phrase structure annotation of a large corpus](#). *Natural Language Engineering*, 11(2):207–238.
- Hiroyasu Yamada and Yuji Matsumoto. 2003. [Statistical dependency analysis with support vector machines](#). In *Proceedings of the Eighth International Conference on Parsing Technologies*, pages 195–206, Nancy, France.
- Songlin Yang and Kewei Tu. 2022. [Headed-span-based projective dependency parsing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2188–2200, Dublin, Ireland. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Yu Zhang, Zhenghua Li, and Min Zhang. 2020. [Efficient second-order TreeCRF for neural dependency parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3295–3305, Online. Association for Computational Linguistics.
- Yue Zhang and Stephen Clark. 2008. [A tale of two parsers: Investigating and combining graph-based and transition-based dependency parsing](#). In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 562–571, Honolulu, Hawaii. Association for Computational Linguistics.
- Junru Zhou and Hai Zhao. 2019. [Head-Driven Phrase Structure Grammar parsing on Penn Treebank](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2396–2408, Florence, Italy. Association for Computational Linguistics.
- Ran Zmigrod, Tim Vieira, and Ryan Cotterell. 2020. [Please mind the root: Decoding arborescences for dependency parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4809–4819, Online. Association for Computational Linguistics.

A Algorithms

Algorithm 1 Create a BHT from a dependency tree.

```
1: procedure DEP2TREE(head)
2:   stack.push(head)
3:   for dep in LEFTDEP(head) :
4:     DEP2TREE(dep)
5:     left ← stack.pop()
6:     right ← stack.pop()
7:     stack.push(MAKENODE( $\textcircled{R}$ , left, right))
8:   for dep in RIGHTDEP(head) :
9:     DEP2TREE(dep)
10:    left ← stack.pop()
11:    right ← stack.pop()
12:    stack.push(MAKENODE( $\textcircled{L}$ , left, right))
```

Algorithm 2 Create a dependency tree from a BHT.

```
1: procedure TREE2DEP(node)
2:   if node is leaf :
3:     return node
4:   left ← TREE2DEP(node.left)
5:   right ← TREE2DEP(node.right)
6:   if node is  $\textcircled{L}$  :
7:     return MAKEARC(left  $\overset{\curvearrowright}{\rightarrow}$  right)
8:   return MAKEARC(left  $\overset{\curvearrowleft}{\leftarrow}$  right)
```

B Related Work

Traditionally, approaches to dependency parsing have been taxonomized into graph-based and transition-based parsers. The authors of this paper take the stance that this distinction is misleading because the difference lies not in the models themselves, but rather in whether exact or approximate inference algorithms are employed. For instance, [Kuhlmann et al. \(2011\)](#) gives exact algorithms for transition-based dependency parsers, which exposes the inability to formally distinguish graph-based and transition-based parsers. Thus, we classify our related work into sections: exact and approximate decoding. Further, we review works on tagging-based parsing which is the most relevant line of work to this paper.

Exact Decoding. Most exact algorithms for projective dependency parsing models apply a modified form of the CKY algorithm on nested dependency trees. The best runtime among the commonly deployed algorithms $\mathcal{O}(N^3)$ ([Eisner, 1996](#)), but algorithms based on fast matrix multiplication exist and can achieve a lower runtime bound ([Cohen and Gildea, 2016](#)). However, exact decoding

of *non-projective parsers* is intractable unless under independence assumptions, e.g., edge factored assumption ([McDonald and Satta, 2007](#)). Edge-factored parsers ([McDonald et al., 2005](#); [Dozat et al., 2017](#)) construct graphs by scoring all possible arcs between each pair of words. They then use the maximum spanning tree (MST) finding algorithms for decoding to build the valid dependency trees with maximum score in $\mathcal{O}(N^2)$ ([Zmigrod et al., 2020](#)). The discussed algorithms are exact in inferring the dependency structure, however, they are neither fast nor parallelizable.

Approximate Decoding. Despite not being exact, transition-based parsers offer faster and typically linear-time parsing algorithms ([Kudo and Matsumoto, 2002](#); [Yamada and Matsumoto, 2003](#); [Nivre, 2003](#)). The dependency tree is inferred with a greedy search through transition system actions. Following this approach, actions are not predicted in parallel and the configuration of the transition system (stack and buffer) needs to be modeled with a neural network ([Chen and Manning, 2014](#)), which prevents using pretrained models out of the box.

Tagging-based parsing. Inspired by [Bangalore and Joshi’s \(1999\)](#) seminal work *supertagging*, a recent line of work aims to utilize pretrained models and parse dependency trees by inferring tags for each word in the input sequence. [Li et al. \(2018\)](#); [Kiperwasser and Ballesteros \(2018\)](#) predict the relative position of the dependent with respect to its parent as the tag. They then use beam tree constraints ([Lee et al., 2016](#)) to infer valid dependency trees. [Strzyz et al. \(2019\)](#) provides a framework for analyzing similar tagging schemes. Although these works have demonstrated potential in this area, none achieved state-of-the-art results compared to custom architectures and algorithms developed for dependency parsing. Additionally, the output space, or size of the tag set, is unrestricted, which limits the efficiency of this approach.

C Analysis

LEFT-FIRST vs. RIGHT-FIRST. We examine the effect of the two orders of binarization of Alg. 1 in Table 4. In our experiments, the choice of left-first or right-first order has little to no effect on parsing performance.

Model	PTB		CTB	
	UAS	LAS	UAS	LAS
Right-first	97.2	96.3	93.2	91.9
Left-first	97.4	96.4	93.1	91.9

Table 4: Comparison of left-first and right-first binarization.

D Efficiency Evaluation

For efficiency comparison, we use BERT-large as the base feature encoder for both Hexatagger and Biaffine. We use the English PTB test set and truncate or pad the input sentences to the control length. The results are averaged over 3 random runs on the same server with one Nvidia A100-80GB GPU. The other experimental settings are kept the same (i.e., the version of PyTorch and Transformer, FP32 precision, batching).

E Datasets

Preprocessing. Following previous work (Kiperwasser and Goldberg, 2016; Dozat and Manning, 2017), the dependency annotations are derived by the Stanford Dependency converter v3.3.0 (de Marneffe and Manning, 2008) from the treebank annotations. Punctuation is omitted for evaluation. Gold part-of-speech tags are provided to the model both during training and evaluation following the code released by Mrini et al. (2020).

Some other authors use system-predicted part-of-speech tags (Zhou and Zhao, 2019) or use mixed configurations. E.g., Yang and Tu (2022) uses gold part-of-speech tags on CTB and UD, while not using any on PTB, Dozat and Manning (2017) uses gold part-of-speech tags on CTB but system-predicted ones on PTB. Our preliminary experiments show that removing the usage of part-of-speech information barely affects the UAS metric, and gives us a performance of 97.4 UAS and 95.8 LAS on PTB.

Splits. All the datasets splits are consistent with previous work. For PTB, we follow the standard split of Marcus et al. (1993), resulting in 39,832 sentences for training, 1,700 for development, and 2,416 for testing. For CTB, we follow the split of Zhang and Clark (2008), resulting in 16,091 sentences for training, 803 for development, and 1,910 for testing. For UD2.2, we follow Yang and Tu (2022) and use the standard splits of the following

corpora for experiments: BG-btb, CA-ancora, CS-pdt, DE-gsd, EN-ewt, ES-ancora, FR-gsd, IT-isdt, NL-alpino, NO-rrt, RO-rrt, RU-syntagrus.

Licenses. The PTB and CTB datasets are licensed under LDC User Agreement. The UD2.2 dataset is licensed under the Universal Dependencies License Agreement.

F Hyperparameter Settings

We use the Python NLTK package to process the datasets, i.e., converting CoNLL-U formatted data to dependency trees, extracting dependency arcs from dependency trees for evaluation, implementing Alg. 1 and 2. For UD, we apply MaltParser v1.9.2⁷ to pseudo-projectivize the non-projective trees (Nivre and Nilsson, 2005).

We use xlnet-large-cased⁸ for English PTB, chinese-xlnet-mid⁹ for CTB, and bert-multilingual-cased¹⁰ for UD.

The dimension of POS tag embedding is set to 256 for all experiments. On top of concatenated pretrained representations and POS embedding, we use a 3-layer BiLSTM with a hidden size of 768 for base-sized models (bert-multilingual-cased on UD) and 1024 for large-sized models (xlnet-large-cased on PTB and chinese-xlnet-mid on CTB).

Dropout layers with a rate of 0.33 are applied after the concatenated embedding layer, between LSTM layers, and before the MLP projection layer to hexatags.

For training, we used AdamW with a learning rate of $2e-5$ for pretrained LMs and $1e-4$ for POS embedding, BiLSTM, and MLP. The gradient clipping threshold is set to 1.0. The batch size is set to 32.

⁷<http://www.maltparser.org/download.html>

⁸<https://huggingface.co/xlnet-large-cased>

⁹<https://huggingface.co/hfl/chinese-xlnet-mid>

¹⁰<https://huggingface.co/bert-base-multilingual-cased>

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations
- A2. Did you discuss any potential risks of your work?
Ethics Statement
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and Sec. 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

5, App. D

- B1. Did you cite the creators of artifacts you used?
5
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
App D
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
App D
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
App D
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
App D

C Did you run computational experiments?

5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
App. E

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

App. E

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

5

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

App. F

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.