

Improving BERT Pretraining with Syntactic Supervision

Giorgos Tziafas

University of Groningen
g.t.tziafas@rug.nl

Konstantinos Kogkalidis

Aalto University
kokos.kogkalidis@aalto.fi

Gijs Wijnholds

Leiden Institute of Advanced Computer Science
g.j.wijnholds@liacs.leidenuniv.nl

Michael Moortgat

Utrecht University
m.j.moortgat@uu.nl

Abstract

Bidirectional masked Transformers have become the core theme in the current NLP landscape. Despite their impressive benchmarks, a recurring theme in recent research has been to question such models' capacity for syntactic generalization. In this work, we seek to address this question by adding a supervised, token-level supertagging objective to standard unsupervised pretraining, enabling the explicit incorporation of syntactic biases into the network's training dynamics. Our approach is straightforward to implement, induces a marginal computational overhead and is general enough to adapt to a variety of settings. We apply our methodology on Lassy Large, an automatically annotated corpus of written Dutch. Our experiments suggest that our syntax-aware model performs on par with established baselines, despite Lassy Large being one order of magnitude smaller than commonly used corpora.

1 Introduction

In recent years, the advent of Transformers (Vaswani et al., 2017) has paved the way for high-performing neural language models, with BERT (Devlin et al., 2019) and its many variants being the main exemplar (Liu et al., 2019; Sanh et al., 2019; Lan et al., 2020). BERT-like models achieve state-of-the-art scores in most major NLP benchmarks via a two-step process. First, they are trained on massive-scale, minimally processed raw text corpora by employing the so-called masked language modeling (*MLM*) objective. Task-specific refinements are then obtained by fine-tuning the pretrained model on labeled corpora, usually orders of magnitude smaller in size.

This pipeline, despite its attested performance, suffers from two key limitations. On the one hand, training a BERT-like model from scratch requires an often prohibitive amount of data and computational resources, barring entry to research projects that lack access to either. On the other hand, a

naturally emerging question is whether such models develop an internal notion of syntax. Discovery of structural biases is hindered by their distributed, opaque representations, requiring manually designed *probing* tasks to extract evidence of syntactic awareness (Hewitt and Manning, 2019; Tenney et al., 2019; Kim et al., 2020; Clark et al., 2019a; Goldberg, 2019; Hu et al., 2020). Alternatively, when syntactic evaluation becomes the focal point, it is usually deferred to downstream tasks (Kitaev et al., 2019; Zhang et al., 2020a), owing both to the lack of sufficiently large labeled corpora as well as the computational bottleneck imposed by hard-to-parallelize operations.

In this work, we seek to alleviate both points by considering them in tandem. Contrary to prior work, we consider the case of introducing explicit syntactic supervision during the pretraining process and investigate whether it can allow for a reduction in the data needs of a BERT-like language model. To facilitate this, we couple the standard unsupervised MLM task with a supervised task, mapping each distinct word to a *supertag*, an abstract syntactic descriptor of its functional role within the context of its surrounding phrase. In essence, this amounts to simple token-level classification, akin to traditional supertagging (Bangalore and Joshi, 1999), except for parts of the input now being masked. In employing both objectives, we ensure that our model is syntax-aware by construction, while incurring only a negligible computational overhead. We evaluate the trained model's performance in a variety of downstream tasks and find that it performs on par with established models, despite being trained on a significantly smaller corpus. Our preliminary experiments suggest an improvement to pretraining robustness and offer a promising direction for cheaper and faster training of structure-enhanced language models. Reflecting on the added objective, we call our model *tagBERT*.

2 Background

Embedding structural biases in neural language models has been a key theme in recent research. Most syntax-oriented models rely on computationally intensive, hard-to-parallelize operations that constrain their integrability with the state of the art in unsupervised language modeling (Tai et al., 2015; Dyer et al., 2016; Kim et al., 2019). This can be ameliorated by either asynchronous pretraining, relying on accurate but slow oracles (Kuncoro et al., 2019), or multi-task training, where the system is exposed to a syntactic task for only part of its training routine (Clark et al., 2018, 2019b). In the BERT setting, there have been attempts at modifying the architecture by either overlaying syntactic structure directly on the attention layers of the network (Wang et al., 2019b) or imposing shallow syntactic cues and/or semantic information in a multi-task setting (Zhang et al., 2020b; Zhou et al., 2020). While such a setup allows for efficient parallel pretraining, the rudimentary nature of the utilized annotations typically forfeits fine aspects of sentential structure, such as function-argument relations.

In this paper, we adopt lexicalism in the categorial grammar tradition (Ajdukiewicz, 1935; Lambek, 1958; Buszkowski et al., 1988; Steedman, 1993; Moortgat, 1997), according to which (most of) the grammatical structure of a language is encoded in its lexicon via an algebra of types that governs the process of phrasal composition. Under such a regime, the parse tree underlying a sentence can be partially (or even fully, in the case of an adequately “strict” grammar) recovered from its constituent words and their respective types alone. In applied terms, the lexical nature of categorial grammars provides us with the opportunity of capturing syntax in a fully-parallel fashion that is straightforward to incorporate with the masked language modeling objective of BERT-like architectures, a fact so far generally overlooked by machine learning practitioners. This perspective is in line with recent insights arguing for the necessity of explicit supervision for syntactic acquisition (Bailly and Gábor, 2020).

The only prerequisite for our methodology is an adequately sized, categorially annotated corpus. Even though gold standard corpora exist for a variety of languages and grammars (Chen and Shanker, 2004; Hockenmaier, 2006; Hockenmaier and Steedman, 2007; Tse and Curran, 2010; Am-

bati et al., 2018; Kogkalidis et al., 2020b), their size is generally insufficient for training a parameter-rich neural language model. This limiting factor can be counteracted by either lexicalizing existing silver-standard corpora of a larger size, or by using an off-the-shelf, high-performance supertagger to annotate the source data prior to pretraining. In both cases the trained system is likely to inherit common errors of the data-generating teacher; the question is whether the added structural biases facilitate faster training of more general language models, despite potential tagging inaccuracies.

3 Methodology

3.1 Data

To facilitate both the data needs of the neural language model and the added supertagging objective we employ Lassy Large (van Noord et al., 2013), a corpus of written Dutch, automatically parsed using the Alpino parser (Bouma et al., 2001). The dataset is comprised of a selection of smaller corpora from varying sources, ranging from excerpts from conventional and modern media to spoken transcripts, enumerating a total of almost 800 million words. Lassy’s syntactic analyses take the form of directed acyclic graphs, with nodes corresponding to words or phrases marked with their part-of-speech as well as syntactic category labels and edges denoting dependency relations. To make the analyses applicable for our setup, we lexicalize them using the type extraction algorithm of Kogkalidis et al. (2020b). The algorithm traverses a parse graph and encodes its structure in a linear logic proof, under the general paradigm of categorial type logics (Moortgat, 1997), simultaneously capturing function-argument and dependency structure. Words, i.e. fringe nodes in the graph, are assigned *types*, abstract syntactic signs that encode a considerable portion of the full structure.

Applying the extraction algorithm, we obtain a collection of around 66 million sentences, represented as sequences of word-type pairs. We drop about 20 million of these in a sanitation step, due to either being duplicates or overlapping with any of the evaluation tasks. We tokenize words using a preconstructed WordPiece (Schuster and Nakajima, 2012) vocabulary of 30 000 tokens based on a larger collection of written Dutch corpora (Vries et al., 2019). Further, we keep the 2 883 most frequent types, which suffice to cover 95% of the type occurrences in the dataset, and replace the filtered

out types with an UNK token. We finally discard sentences lying in the 5%-tail of the length distribution, and train with 45 million sentences spanning less than 100 sub-word tokens.

3.2 Model

Our model is a faithful replica of BERT_{BASE}, except for having a hidden size of 1 536 instead of 3 072 for the intermediate fully-connected layers, reducing our total parameters from 110 to 79 million. We further employ a linear projection from the model’s dimensionality to the number of types in our vocabulary, which we attach to the output of a prespecified encoder block. The projection can be separably applied on the encoder’s intermediate representations, allowing us to optionally query the model for a class weighting over types for each input token.

This addition accounts to a mere 2.5% of the model’s total parameter count and only incurs a negligible computational overhead if explicitly enabled, as it does not interfere with the forward pass when the system is run solely as a contextualization model. If the type classification layer is enabled during pretraining, it introduces a clear error signal that updates all network weights up to the connected encoder block, bolstering the correct acquisition of syntax in the bottom part of the encoding pipeline.

3.3 Pretraining

To train our model, we feed it partially masked sentences following the methodology of Liu et al. (2019); we dynamically mask continuous spans of tokens belonging to the same word and drop the next sentence prediction task, training on single sentences instead. Attaching the type classification layer at the fourth encoder block, we end up with two output streams.¹ One is a prediction over the subword vocabulary for each masked token, as in vanilla BERT, whereas the other comes from the type classifier, yielding a prediction over the type vocabulary for every token, masked or otherwise.² We obtain a loss function by summing the cross-entropy between predictions and truths for each output stream.

¹The choice of depth for the type classifier is due to preliminary experiments where we let a trainable layer weighter freely select from the range of encoder blocks. In the vast majority of runs, most of the importance was interestingly assigned to the fourth layer.

²Masking entire words for the supertagging task can be seen as a severe form of regularization, à la channel dropout.

To deal with the misalignment between subword units and types, we associate every type with the first token of its corresponding word, and mask out predictions spanning subsequent tokens when performing the loss computation. Similarly, we do not penalize predictions over types discarded by the occurrence count filtering (UNK types). For regularization purposes, we randomly replace output types 1% of the time (Wu et al., 2019).

Following standard practices, we optimize using AdamW (Loshchilov and Hutter, 2019) with a batch size of 256, shuffling and iterating the dataset 8 times. The learning rate is gradually increased to 10^{-4} over 10 000 steps and then decayed to zero using a linear warm-up and decay schedule.

4 Evaluation

To evaluate the trained model, we measure its performance on the below selection of downstream tasks, after fine-tuning. We keep our fine-tuning setup as barebones as possible, using Adam (Kingma and Ba, 2014) with a batch size of 32 and a learning rate of 3×10^{-5} . We apply model selection based on the validation-set performance and report test-set results (averaged over three runs) against the available baselines of each task in Table 1. In order to provide fair comparisons, we replicate the evaluation of other models using the same experimental setup.

Lassy Small is a gold-standard syntactically annotated corpus for written Dutch (van Noord et al., 2013). We fine-tune a POS tagger on the subset of the corpus that has been converted to Universal Dependency format (Bouma and van Noord, 2017).

SoNaR-1 is a curated subset of Lassy Small that includes several layers of manually added annotations (Delaere et al., 2009). We employ the named entity recognition (NER), part-of-speech (POS), semantic role labeling (SRL) and spatio-temporal relation tags (STR) that come packed with the corpus and treat their classification as downstream tasks. NER contains approximately 60 000 samples and 6 class labels encoded in the IOB scheme. POS tagging contains about 16 000 samples and comes in two varieties: coarse (12 classes) and fine-grained (241 classes, out of which only 223 appear in the training data, many just once). SRL comes also in two varieties: a) predicate-argument structures, encoded with IOB scheme, and b) modifiers, enhanced with modified phrase labels, for a total of

	<i>SoNaR-1</i>			<i>Lassy UD</i>	<i>CoNLL</i>	<i>Æthel</i>	
	POS-coarse	POS-fine	NER	POS	NER	Supertags	Parse
<i>BERTje</i> (Vries et al., 2019)	98.8	97.5	87.4	96.4	90.6	85.5	56.9
<i>RobBERT</i> (Delobelle et al., 2020)	98.5	97.2	84.8	96.2	85.9	86.3	56.8
<i>tagBERT</i> (ours)	98.8	97.4	87.0	96.7	89.9	86.6	58.3

Table 1: Comparative performance for a selection of downstream tasks. We report test set accuracy (%) on all tasks except NER, where we report F1 scores (%) as produced by the CoNLL evaluation script (Tjong Kim Sang, 2002). For a fair comparison, we replicate the fine-tuning process on all pretrained baselines, including truncation of the maximum token length to 100.

	<i>SoNaR-1</i>			<i>Europarl</i>	<i>DBRD</i>
	SRL-pred	SRL-mod	STR	die/dat	sentiment
<i>BERTje</i> (Vries et al., 2019)	85.3	67.2	57.3	95.0	93.0
<i>RobBERT</i> (Delobelle et al., 2020)	—	—	—	98.7	95.1
<i>tagBERT</i> (ours)	86.5	67.8	68.0	99.1	93.8

Table 2: Comparative performance on higher-level downstream tasks. Scores are F1 (%) for SRL/STR, and test set accuracy (%) for *die/dat* disambiguation and sentiment analysis. —: no results available.

30 000 samples. STR contains a total of 58 000 spatio-temporal tags, including geolocations and use of past verbe tense.

CoNLL-2002 is a named entity recognition dataset from the corresponding shared task (Tjong Kim Sang, 2002). The dataset contains 4 class labels, also encoded in the IOB scheme, with a total size of approximately 24 000 samples.

Æthel is a typological derivation dataset, generated by applying the type extraction algorithm to Lassy Small (Kogkalidis et al., 2020b). We replicate the experiments of Kogkalidis et al. (2020a) to train a typological grammar parser, but instantiate the encoder part with the baselines of Table 1, and report token-level supertagging accuracy as well as full sentential parsing accuracy in the greedy setting. We note that even though our model is exposed to types during pretraining, their representation format is vastly different during the fine-tuning process; rather than being classification outputs for each word, they are broken down to their primitive symbols and transduced from the input sequence with auto-regressive *seq2seq* decoding. In that sense, this task helps us assess the generality of the learned representations.

Dutch Europarl is a sanitized subset of transcripts of the European Parliament in Dutch (Koehn, 2005), used for zero-shot evalua-

tion of the task of relative pronoun disambiguation. The task revolves around picking the most likely between the Dutch relative pronouns *die* and *dat*. While the two agree in their syntactic function (and grammatical category), the former selects exclusively for gendered nouns, whereas the latter selects for neuter ones. As such, the task measures our model’s capacity to resolve morphosyntactic constraints in the presence of grammatical category invariants. This corpus enumerates a total of 1.56M sentences, 90.7% of which contain at least one relative pronoun.

110k Dutch Book Reviews Dataset. The Dutch Book Reviews Dataset (DBRD) is a sentiment analysis benchmark which comprises around 110k Dutch book reviews taken from hebban.nl, out of which 22 252 are manually labeled as either positive or negative (Van der Burgh and Verberne, 2019) and segmented into 90% training and 10% testing splits. Unlike previous tasks, sentiment analysis is done on the sentential, rather than token level, serving as a measure of the model’s semantic understanding.

5 Discussion

Our model performs on par across all tasks considered, indicating pretraining robustness comparable to the heavy weight baselines of BERT- (Devlin et al., 2019) and RoBERTa-based (Liu et al., 2019)

models.³ Considering the non-ideal nature of the silver-standard tags, the significantly smaller size of our corpus compared to competing models, and the ca. 30% reduced parameter count (79M against 119M for BERTje and 117M for RobBERT), our results can be seen as strong evidence in favor of explicitly encoding structural biases in the pretraining process of neural language models. Opting for a lexicalized representation of structure allows for a seamless and cost-efficient integration with BERT’s core architecture, essentially removing the computational bottleneck of alternating between tensor optimization and structure manipulation.

6 Conclusion

We introduced tagBERT, a variation of BERT that is biased towards syntax through coupling the standard MLM loss with a supertagging objective. We trained tagBERT on a modestly sized, silver-standard corpus of written Dutch – after first lexicalizing its annotations – and evaluated the trained model on a number of downstream NLP tasks after fine-tuning. Despite a reduced parameter count and the corpus’ modest size, our method is achieving performance comparable to established state-of-the-art models. This result is contrary to the ongoing trend of utilizing increasingly more data and augmenting model capacity, instead suggesting potential benefits from incorporating richer annotations in convenient representation formats. Our work aims towards a syntactically-transparent, cost-efficient language model that combines both the rigor of formal linguistic theories and the representational power of large-scale unsupervised learning.

Retroactive Placement The current work is situated in the historical landscape where probing for syntactic awareness and the possibility of injecting syntactic structures in the network were still novel enterprises. Alongside and following our original endeavour, many more studies have investigated the role of syntax and ways to incorporate it within large language models, with the end-goal of either jointly acquiring the two, or of using explicit syntactic guidelines to constrain language generation (Zanzotto et al., 2020; Sartran et al., 2022; Li et al., 2021; Bai et al., 2021; Song et al., 2022; Xie et al., 2021; Li et al., 2023). Other than technical differences on the neural front, our

work diverges in opting for a linearized representation of syntax through categorial grammars. This choice stands out for its elegance and formal coherence, setting it apart from more widely used alternatives like constituency trees and dependency arcs. Indeed, categorial grammars (regardless of the particular flavor adopted) offer the means for an expressive, yet fully lexicalized, modeling of syntax and compositional meaning. Their integration into the modern NLP toolkit is facilitated by this inherent flexibility, offering the potential for intricate interplay between structure and form – a potential that still remains, for the most part, untapped.

Future work Given the embedding of this paper within the landscape concerning syntactic awareness and large language models, future work would be based on recent developments in neurosymbolic approaches to lexicalized grammar formalisms. For instance, recent developments in neural supertagging could be exploited, for instance by updating the supertagging from the discriminative setting to a constructive one (Prange et al., 2021; Kogkalidis et al., 2023).

Besides developments relevant to model architecture, several novel evaluation tasks for Dutch have been developed that may shed light on the distinction between vanilla Transformer-based models and syntactically informed ones. For example, the two-sentence classification task of Natural Language Inference (NLI) is a typical task that tests for lexical, syntactic, and sentence-level understanding, for which two Dutch benchmarks exist (Wijnholds and Moortgat, 2021; Wijnholds, 2023). A comparison between a vanilla model, the syntactically informed tagBERT, and the neurosymbolic approach of Abzianidze and Kogkalidis (2021) is in place to put the relationship of syntax and NLI in perspective. Further on the Dutch front, we would be keen to test the model’s ability to understand discontinuous verb-subject dependencies as in Kogkalidis and Wijnholds (2022), or to disambiguate relative clauses as in Wijnholds and Moortgat (2023).

Finally, we invite and look forward to different research directions, such as experimentation with different languages and grammar formalisms, integration with existing pre-trained models in an intermediate-training fashion (Wang et al., 2019a) and exploring architectural adjustments that would allow a two-way dependence or a stronger interfacing between the lexical and syntactic modalities.

³Implementation code is available at <https://github.com/gtziafas/type-enhanced-language-modeling>.

7 Acknowledgments

The authors warmly express their support from the Dutch Research Council (NWO) due to which three of the authors were operational under the scope of the project “A composition calculus for vector-based semantic modelling with a localization for Dutch” (360-89-070), at the time of execution of this research project.

References

- Lasha Abzianidze and Konstantinos Kogkalidis. 2021. [A logic-based framework for natural language inference in dutch](#). *Computational Linguistics in the Netherlands Journal*, 11:35–58.
- Kazimierz Ajdukiewicz. 1935. Die syntaktische Konnexität. *Studia philosophica*, 1:1–27. English translation “Syntactic Connexion” by H. Weber in McCall, S. (Ed.) *Polish Logic*, pp. 207–231, Oxford University Press, Oxford, 1967.
- Bharat Ram Ambati, Tejaswini Deoskar, and Mark Steedman. 2018. [Hindi CCGbank: A CCG treebank from the Hindi dependency treebank](#). *Language Resources and Evaluation*, 52(1):67–100.
- Jiangang Bai, Yujing Wang, Yiren Chen, Yaming Yang, Jing Bai, Jing Yu, and Yunhai Tong. 2021. [Syntax-BERT: Improving pre-trained transformers with syntax trees](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3011–3020, Online. Association for Computational Linguistics.
- Raphaël Bailly and Kata Gábor. 2020. [Emergence of syntax needs minimal supervision](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 477–487, Online. Association for Computational Linguistics.
- Srinivas Bangalore and Aravind K. Joshi. 1999. [Supertagging: An approach to almost parsing](#). *Computational Linguistics*, 25(2):237–265.
- Gosse Bouma and Gertjan van Noord. 2017. [Increasing return on annotation investment: The automatic construction of a universal dependency treebank for Dutch](#). In *Proceedings of the NoDaLiDa 2017 Workshop on Universal Dependencies (UDW 2017)*, pages 19–26, Gothenburg, Sweden. Association for Computational Linguistics.
- Gosse Bouma, Gertjan Van Noord, and Robert Malouf. 2001. [Alpino: Wide-coverage computational analysis of Dutch](#). In *Computational linguistics in the Netherlands 2000*, pages 45–59. Brill Rodopi.
- Wojciech Buszkowski, Witold Marciszewski, and Johan van Benthem. 1988. *Categorical grammar*. John Benjamins Publishing.
- John Chen and Vijay K Shanker. 2004. [Automated extraction of TAGs from the Penn treebank](#). In *New developments in parsing technology*, pages 73–89. Springer.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019a. [What does BERT look at? An analysis of BERT’s attention](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Urvashi Khandelwal, Christopher D. Manning, and Quoc V. Le. 2019b. [BAM! born-again multi-task networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5931–5937, Florence, Italy. Association for Computational Linguistics.
- 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.
- Isabelle Delaere, Veronique Hoste, and Paola Monachesi. 2009. [Cultivating trees: Adding several semantic layers to the Lassy treebank in SoNaR](#). In *7th International workshop on Treebanks and Linguistic Theories (TLT-7)*, pages 135–146. LOT (Landelijke Onderzoeksschool Taalwetenschap).
- Pieter Delobelle, Thomas Winters, and Bettina Berendt. 2020. [RobBERT: a Dutch RoBERTa-based language model](#). *arXiv preprint arXiv:2001.06286*.
- 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.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. [Recurrent neural network grammars](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.
- Yoav Goldberg. 2019. [Assessing BERT’s syntactic abilities](#). *arXiv preprint arXiv:1901.05287*.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Julia Hockenmaier. 2006. [Creating a CCGbank and a wide-coverage CCG lexicon for German](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 505–512, Sydney, Australia. Association for Computational Linguistics.
- Julia Hockenmaier and Mark Steedman. 2007. [CCGbank: A corpus of CCG derivations and dependency structures extracted from the Penn Treebank](#). *Computational Linguistics*, 33(3):355–396.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A systematic assessment of syntactic generalization in neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Taeuk Kim, Jihun Choi, Daniel Edmiston, and Sang-goo Lee. 2020. [Are pre-trained language models aware of phrases? Simple but strong baselines for grammar induction](#). In *International Conference on Learning Representations*.
- Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019. [Unsupervised recurrent neural network grammars](#). 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 1105–1117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Diederik P Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- Nikita Kitaev, Steven Cao, and Dan Klein. 2019. [Multilingual constituency parsing with self-attention and pre-training](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3499–3505, Florence, Italy. Association for Computational Linguistics.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Konstantinos Kogkalidis, Michael Moortgat, and Richard Moot. 2020a. [Neural proof nets](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 26–40, Online. Association for Computational Linguistics.
- Konstantinos Kogkalidis, Michael Moortgat, and Richard Moot. 2020b. [Æthel: Automatically extracted typological derivations for Dutch](#). In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 5257–5266, Marseille, France. European Language Resources Association.
- Konstantinos Kogkalidis, Michael Moortgat, and Richard Moot. 2023. [SPINDLE: Spinning raw text into lambda terms with graph attention](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 128–135, Dubrovnik, Croatia. Association for Computational Linguistics.
- Konstantinos Kogkalidis and Gijs Wijnholds. 2022. [Discontinuous constituency and BERT: A case study of Dutch](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3776–3785, Dublin, Ireland. Association for Computational Linguistics.
- Adhiguna Kuncoro, Chris Dyer, Laura Rimell, Stephen Clark, and Phil Blunsom. 2019. [Scalable syntax-aware language models using knowledge distillation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3472–3484, Florence, Italy. Association for Computational Linguistics.
- Joachim Lambek. 1958. [The mathematics of sentence structure](#). *The American Mathematical Monthly*, 65(3):154–170.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. [ALBERT: A lite BERT for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Yafu Li, Leyang Cui, Jianhao Yan, Yongjing Yin, Wei Bi, Shuming Shi, and Yue Zhang. 2023. [Explicit syntactic guidance for neural text generation](#).
- Zhongli Li, Qingyu Zhou, Chao Li, Ke Xu, and Yunbo Cao. 2021. [Improving BERT with syntax-aware local attention](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 645–653, Online. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A robustly optimized BERT pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Michael Moortgat. 1997. [Categorial type logics](#). In Johan van Benthem and Alice ter Meulen, editors, *Handbook of logic and language*, chapter 2, pages 93–177. Elsevier/MIT Press.
- Jakob Prange, Nathan Schneider, and Vivek Sriku-mar. 2021. [Supertagging the Long Tail with Tree-Structured Decoding of Complex Categories](#). *Transactions of the Association for Computational Linguistics*, 9:243–260.

- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *arXiv preprint arXiv:1910.01108*.
- Laurent Sartran, Samuel Barrett, Adhiguna Kuncoro, Miloš Stanojević, Phil Blunsom, and Chris Dyer. 2022. [Transformer Grammars: Augmenting Transformer Language Models with Syntactic Inductive Biases at Scale](#). *Transactions of the Association for Computational Linguistics*, 10:1423–1439.
- Mike Schuster and Kaisuke Nakajima. 2012. [Japanese and Korean voice search](#). In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5149–5152. IEEE.
- Jian Song, Di Liang, Rumei Li, Yuntao Li, Sirui Wang, Minlong Peng, Wei Wu, and Yongxin Yu. 2022. [Improving semantic matching through dependency-enhanced pre-trained model with adaptive fusion](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 45–57, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mark Steedman. 1993. [Categorial grammar](#). *Lingua*, 90(3):221–258.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. [Improved semantic representations from tree-structured long short-term memory networks](#). 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 1556–1566, Beijing, China. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019. [What do you learn from context? Probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Erik F. Tjong Kim Sang. 2002. [Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition](#). In *COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002)*.
- Daniel Tse and James R. Curran. 2010. [Chinese CCG-bank: extracting CCG derivations from the Penn Chinese treebank](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 1083–1091, Beijing, China. Coling 2010 Organizing Committee.
- Benjamin Van der Burgh and Suzan Verberne. 2019. [The merits of universal language model fine-tuning for small datasets—a case with dutch book reviews](#). *arXiv preprint arXiv:1910.00896*.
- Gertjan van Noord, Gosse Bouma, Frank Van Eynde, Daniël de Kok, Jelmer van der Linde, Ineke Schuurman, Erik Tjong Kim Sang, and Vincent Vandeghinste. 2013. *Large Scale Syntactic Annotation of Written Dutch: Lassy*, pages 147–164. Springer Berlin Heidelberg.
- 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, pages 5998–6008. Curran Associates, Inc.
- Wietse de Vries, Andreas van Cranenburgh, Arianna Bisazza, Tommaso Caselli, Gertjan van Noord, and Malvina Nissim. 2019. [BERTje: A Dutch BERT model](#).
- Alex Wang, Jan Hula, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019a. [Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.
- Yaoshian Wang, Hung-Yi Lee, and Yun-Nung Chen. 2019b. [Tree transformer: Integrating tree structures into self-attention](#). 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 1061–1070, Hong Kong, China. Association for Computational Linguistics.
- Gijs Wijnholds. 2023. [Assessing monotonicity reasoning in Dutch through natural language inference](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1494–1500, Dubrovnik, Croatia. Association for Computational Linguistics.
- Gijs Wijnholds and Michael Moortgat. 2021. [SICK-NL: A dataset for Dutch natural language inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1474–1479, Online. Association for Computational Linguistics.
- Gijs Wijnholds and Michael Moortgat. 2023. [Structural ambiguity and its disambiguation in language model based parsers: the case of dutch clause relativization](#).
- Liwei Wu, Shuqing Li, Cho-Jui Hsieh, and James L Sharpnack. 2019. [Stochastic shared embeddings: Data-driven regularization of embedding layers](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 24–34. Curran Associates, Inc.
- Yikuan Xie, Wenyong Wang, Mingqian Du, and Qing He. 2021. [Transformer with syntactic position encoding for machine translation](#). In *Proceedings of*

the International Conference on Recent Advances in Natural Language Processing (RANLP 2021), pages 1536–1544, Held Online. INCOMA Ltd.

Fabio Massimo Zanzotto, Andrea Santilli, Leonardo Ranaldi, Dario Onorati, Pierfrancesco Tommasino, and Francesca Fallucchi. 2020. **KERMIT: Complementing transformer architectures with encoders of explicit syntactic interpretations**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 256–267, Online. Association for Computational Linguistics.

Yu Zhang, Houquan Zhou, and Zhenghua Li. 2020a. **Fast and accurate neural CRF constituency parsing**. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 4046–4053. International Joint Conferences on Artificial Intelligence Organization. Main track.

Zhuosheng Zhang, Yuwei Wu, Hai Zhao, Zuchao Li, Shuailiang Zhang, Xi Zhou, and Xiang Zhou. 2020b. **Semantics-aware BERT for language understanding**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9628–9635.

Junru Zhou, Zhuosheng Zhang, Hai Zhao, and Shuailiang Zhang. 2020. **LIMIT-BERT: Linguistics informed multi-task BERT**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4450–4461, Online. Association for Computational Linguistics.