

Enriched In-Order Linearization for Faster Sequence-to-Sequence Constituent Parsing

Daniel Fernández-González and Carlos Gómez-Rodríguez

Universidade da Coruña, CITIC

FASTPARSE Lab, LyS Group

Depto. de Ciencias de la Computación y Tecnologías de la Información

Campus de Elviña, s/n, 15071 A Coruña, Spain

d.fgonzalez@udc.es, carlos.gomez@udc.es

Abstract

Sequence-to-sequence constituent parsing requires a linearization to represent trees as sequences. Top-down tree linearizations, which can be based on brackets or shift-reduce actions, have achieved the best accuracy to date. In this paper, we show that these results can be improved by using an in-order linearization instead. Based on this observation, we implement an enriched in-order shift-reduce linearization inspired by Vinyals et al. (2015)'s approach, achieving the best accuracy to date on the English PTB dataset among fully-supervised single-model sequence-to-sequence constituent parsers. Finally, we apply deterministic attention mechanisms to match the speed of state-of-the-art transition-based parsers, thus showing that sequence-to-sequence models can match them, not only in accuracy, but also in speed.

1 Introduction

Sequence-to-sequence (seq2seq) neural architectures have proved useful in several NLP tasks, with remarkable success in some of them such as machine translation, but they lag behind the state of the art in others. In constituent parsing, seq2seq models still need to improve to be competitive in accuracy and efficiency with their main competitors: transition-based constituent parsers (Dyer et al., 2016; Liu and Zhang, 2017b; Fernández-González and Gómez-Rodríguez, 2019).

Vinyals et al. (2015) laid the first stone in seq2seq constituent parsing, proposing a linearization of phrase-structure trees as bracketed sequences following a top-down strategy, which can be predicted from the input sequence of words by any off-the-shelf seq2seq framework. While this approach is very simple, its accuracy and efficiency are significantly behind the state of the art in the fully-supervised single-model scenario.

Most attempts to improve this approach focused on modifying the neural network architecture, while keeping the top-down linearization strategy. As exceptions, Ma et al. (2017) and Liu and Zhang (2017a) proposed linearizations based on sequences of transition-based parsing actions instead of brackets. Ma et al. (2017) tried a bottom-up linearization, but they obtained worse results than top-down approaches.¹ Liu and Zhang (2017a) kept the top-down strategy, but using transitions of the top-down transition system of Dyer et al. (2016) instead of a bracketed linearization, achieving a higher performance.

In transition-based constituent parsing, an in-order algorithm has recently proved superior to the bottom-up and top-down approaches (Liu and Zhang, 2017b), but we know of no applications of this approach in seq2seq parsing.

Contributions In this paper, we advance the understanding of linearizations for seq2seq parsing, and improve the state of the art, as follows: (1) we show that the superiority of a transition-based top-down linearization over a bracketing-based one observed by Liu and Zhang (2017a) does not hold when both are tested under the same framework. In fact, we show that the additional information provided by the larger vocabulary in the linearization of Vinyals et al. (2015) is beneficial to seq2seq predictions. (2) We implement a novel in-order transition-based linearization, based on the in-order transition system by Liu and Zhang (2017b), and manage to notably increase parsing accuracy with respect to previous approaches. (3) We enhance the in-order representation of parse trees by adding extra information following the shift-reduce version of the (Vinyals et al., 2015) linearization, obtaining state-of-the-art accuracy among seq2seq parsers

¹We also tested empirically that a bottom-up linearization is not suitable for seq2seq parsing and discarded that option.

and on par with some well-known transition-based approaches. (4) We bridge the remaining gap with transition-based parsers - parsing speed - by applying a new variant of deterministic attention (Kamigaito et al., 2017; Ma et al., 2017) to restrict the hidden states used to compute the attention vector, doubling the system’s speed. The result is a seq2seq parser² that, for the first time, matches the speed and accuracy of transition-based parsers implemented under the same neural framework. (5) Using the neural framework of Dyer et al. (2015) as testing ground, we perform a homogeneous comparison among different seq2seq linearizations and widely-known transition-based parsers.

2 Enriched Linearizations

To cast constituent parsing as seq2seq prediction, each parse tree needs to be represented as a sequence of symbols that can be predicted from an input sentence. Initially, Vinyals et al. (2015) proposed a top-down bracketed linearization of constituent trees, where opening and closing brackets include non-terminal labels and POS tags are normalized by replacing them with a tag XX. An example is shown in linearization *a* of Figure 1.

As an alternative, Liu and Zhang (2017a) presented a shift-reduce linearization based on the top-down transition system defined for constituent parsing by Dyer et al. (2016) (example *b* in Figure 1). This provides three transitions that can be used on a stack and a buffer to build a constituent tree: a Shift transition to push words from the buffer into the stack, a Non-Terminal-X transition to push a non-terminal node X into the stack, and a Reduce transition to pop elements from the stack until a non-terminal node is found and create a new subtree with all these elements as its children, pushing this new constituent into the stack.

Following Vinyals et al. (2015)’s linearization where closing brackets also include the non-terminal label, we define an equivalent shift-reduce variant, where the Reduce transition is also parameterized with the non-terminal on top of the resulting subtree (Reduce-X). In that way, we can one-to-one map opening brackets to Non-Terminal-X transitions, closing brackets to Reduce-X actions and XX-tags to Shift transitions as shown in example *c* of Figure 1. This enriched version will enlarge the vocabulary, but will also add some extra informa-

²Source code available at <https://github.com/danifg/InOrderSeq2seq>.

tion that, as we will see below, improves parsing accuracy.

As an alternative to the top-down parser of (Dyer et al., 2016), Liu and Zhang (2017b) define a transition system based on in-order traversal, as in left-corner parsing (Rosenkrantz and Lewis, 1970): the non-terminal node on top of the tree being built is only considered after the first child is completed in the stack, building each subtree in a bottom-up manner, but choosing the non-terminal node on top before the new constituent is reduced. Transitions are the same as in the top-down algorithm (plus a Finish transition to terminate the parsing process), but the effect of applying a Reduce transition is different: it pops all elements from the stack until the first non-terminal node is found, which is also popped together with the preceding element in the stack to build a new constituent with all of them as children of the non-terminal node.³


This algorithm pushed state-of-the-art accuracies in shift-reduce constituent parsing; and, as we show in Section 4, it can be successfully applied as a linearization method for seq2seq constituent parsing. Sequence *d* in Figure 1 exemplifies in-order linearization.

Similarly to the enriched top-down variant, we also extend the in-order shift-reduce linearization by parametrizing Reduce transitions. Additionally, we can also add extra information to Shift transitions. (Suzuki et al., 2018) leaves POS tags of punctuation symbols out of the normalization proposed by Vinyals et al. (2015) without further explanation, but possibly they consider it can help seq2seq models. We adapt this idea to our novel enriched in-order linearization and lexicalize Shift transitions when a “.” or a “;” are pushed into the stack as “Shift.” and “Shift;”, respectively.⁴ In our experiments, we see that lexicalizing Shift transitions has indeed an impact on parsing performance. In Figure 1 and sequence *e*, we include an example of this linearization technique.

Note that, although we use a transition-based linearization of parse trees, our approach is agnostic to the stack structure and the parsing process is performed by a simple seq2seq model that straightforwardly *translates* input sequences of words into sequences of shift-reduce actions.

³See Appendix A for more details about the top-down and in-order transition systems.

⁴We do not lexicalize Shift transitions on the enriched shift-reduce top-down variant to perform a fair comparison against the original linearization by Liu and Zhang (2017a).



Top-down linearizations

- a) Bracketed: (S (NP XX XX) NP (VP XX (ADVP XX) ADVP (ADJP XX) ADJP) VP XX) S
- b) Shift-reduce: NT_S NT_{NP} SH SH RE NT_{VP} SH NT_{ADVP} SH RE NT_{ADJP} SH RE RE SH RE
- c) Enriched SH-RE: NT_S NT_{NP} SH SH RE_{NP} NT_{VP} SH NT_{ADVP} SH RE_{ADVP} NT_{ADJP} SH RE_{ADJP} RE_{VP}
SH RE_S

In-order linearizations

- d) Shift-reduce: SH NT_{NP} SH RE NT_S SH NT_{VP} SH NT_{ADVP} RE SH NT_{ADJP} RE RE SH RE FI
- e) Enriched SH-RE: SH NT_{NP} SH RE_{NP} NT_S SH NT_{VP} SH NT_{ADVP} RE_{ADVP} SH NT_{ADJP} RE_{ADJP} RE_{VP}
SH. RE_S FI

Figure 1: Top-down and in-order linearizations for a constituent tree taken from English PTB. SH = Shift, NT_X = Non-Terminal-X, RE = Reduce, RE_X = Reduce-X and FI = Finish.

3 Seq2seq Neural Network

Baseline Model In our experiments, we test all proposed linearizations in the seq2seq neural architecture designed by Liu and Zhang (2017a) and implemented on the framework developed by Dyer et al. (2015). This architecture proved to outperform the majority of seq2seq approaches, even without implementing beam search (which penalizes parsing speed). The difference with respect to the vanilla seq2seq configuration (Vinyals et al., 2015) is that two separate attention models are used to cover two different and variable segments of the input. This provides improvements in accuracy, regardless of the linearization method used.

More specifically, Liu and Zhang (2017a) follow the common practice in stack-LSTM-based shift-reduce parsers (Dyer et al., 2015, 2016; Liu and Zhang, 2017b) that uses a concatenation of pre-trained word embeddings ($\mathbf{e}_{w_i}^*$) and randomly initialized word (\mathbf{e}_{w_i}) and POS tag embeddings (\mathbf{e}_{p_i}) to derive (through a ReLU non-linear function) the final representation \mathbf{x}_i of the i th input word:

$$\mathbf{x}_i = \text{relu}(\mathbf{W}_{enc}[\mathbf{e}_{w_i}^*, \mathbf{e}_{w_i}, \mathbf{e}_{p_i}] + \mathbf{b}_{enc})$$

where \mathbf{W}_{enc} and \mathbf{b}_{enc} are model parameters, and w_i and p_i represent the form and the POS tag of the i th input word.

This representation \mathbf{x}_i is fed into the encoder (implemented by a BiLSTM) to output an *encoder hidden state* \mathbf{h}_i :

$$\mathbf{h}_i = [\mathbf{h}_{l_i}; \mathbf{h}_{r_i}] = \text{BiLSTM}(\mathbf{x}_i).$$

As a decoder, a LSTM generates a sequence of *decoder hidden states* from which a sequence of actions is predicted. Concretely, the current decoder hidden state \mathbf{d}_j is computed by:

$$\mathbf{d}_j = \text{relu}(\mathbf{W}_{dec}[\mathbf{d}_{j-1}, \mathbf{l}_{att_j}, \mathbf{r}_{att_j}] + \mathbf{b}_{dec})$$

where \mathbf{W}_{dec} and \mathbf{b}_{dec} are model parameters, \mathbf{d}_{j-1} is the previous decoder hidden state, and \mathbf{l}_{att_j} and \mathbf{r}_{att_j} are the resulting attention vectors over the left and right segments, respectively, of encoder hidden states $\mathbf{h}_1 \dots \mathbf{h}_n$. These two segments of the input are defined by index p , which is initialized to the beginning of the sentence and moves one position to the right each time a Shift transition is applied. Therefore, \mathbf{l}_{att_j} and \mathbf{r}_{att_j} are computed at timestep j as:

$$\mathbf{l}_{att_j} = \sum_{i=1}^p \alpha_{ij} \mathbf{h}_i, \quad \mathbf{r}_{att_j} = \sum_{i=p+1}^n \alpha_{ij} \mathbf{h}_i,$$

where $\alpha_{ij} = \frac{\exp(\beta_{ij})}{\sum_{k=1}^n \exp(\beta_{kj})}$ and

$$\beta_{ij} = \mathbf{U}^T \tanh(\mathbf{W}_{att}[\mathbf{h}_i; \mathbf{d}_{j-1}] + \mathbf{b}_{att})$$

Then, the current token \mathbf{y}_j is predicted from \mathbf{d}_j as:

$$p(\mathbf{y}_j | \mathbf{d}_j) = \text{softmax}(\mathbf{W}_{pred} * \mathbf{d}_j + \mathbf{b}_{pred}),$$

where \mathbf{W}_{att} , \mathbf{b}_{att} , \mathbf{W}_{pred} and \mathbf{b}_{pred} are parameters. In Figure 2, we graphically describe the neural architecture.

Note that current state-of-the-art transition-based parsers, which rely on stack-LSTMs to represent

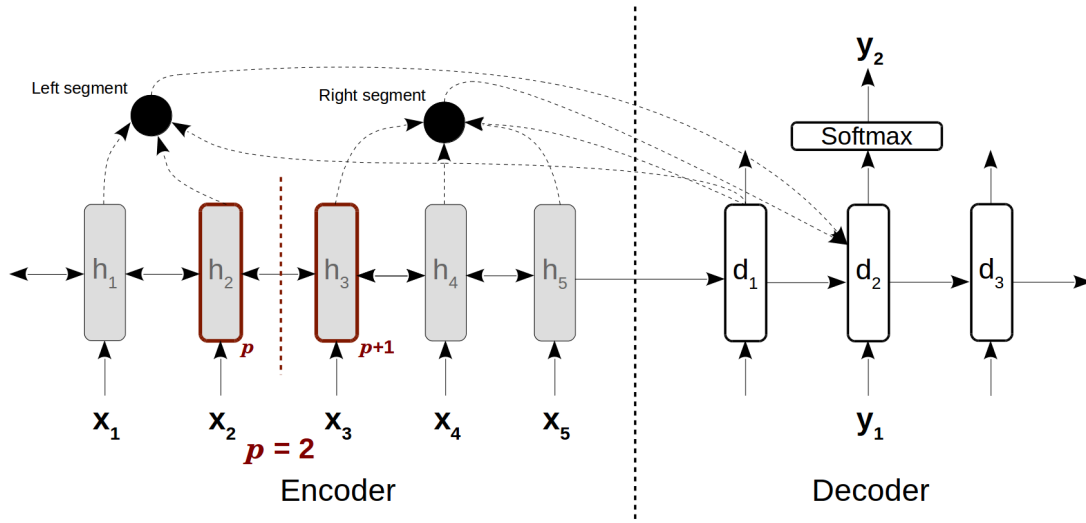


Figure 2: Sequence-to-sequence neural architecture proposed by Liu and Zhang (2017a).

the stack structure, are also implemented under the framework by Dyer et al. (2015) and, therefore, our approach can be fairly compared to them in terms of accuracy and speed.

Deterministic Attention Previous work (Kamigaito et al., 2017; Ma et al., 2017; Liu et al., 2018) claims that using deterministic attention mechanisms instead of the standard probabilistic variant leads to accuracy and speed gains. We propose a simple and effective procedure to implement deterministic attention in the architecture by Liu and Zhang (2017a), substantially reducing the time consumed by the decoder to predict the next token.

Apart from dividing the sequence of encoder hidden states into segments, Liu and Zhang (2017a) provide explicit alignment between the input word sequence and the output transition sequence by keeping the index p that indicates a correspondence between input words and Shift transitions. This information can be used to force the model to focus on those encoder hidden states that are more informative for decoding at each timestep, avoiding going through the whole input to compute the attention vector, and thus considerably reducing decoding time.

To gain some insight on what input words are most relevant, we study on the dev set the attention values assigned by the model to each encoder hidden state and the frequency with which each of them achieves the highest value at each timestep. Surprisingly, we found out that, for the top-down parser, almost 90% of the time the highest attention values were assigned to the words in positions p

and $p + 1$ by a wide margin. For the in-order parser, words in those positions also received considerable attention values, but they were determinant only 75% of the time. Following these results, we propose a computation of \mathbf{l}_{att_j} and \mathbf{r}_{att_j} where only the encoder hidden states in the rightmost position (p) of the left segment and in the leftmost position ($p + 1$) of the right segment are considered:

$$\mathbf{l}_{att_j} = \beta_{pj} \mathbf{h}_p, \quad \mathbf{r}_{att_j} = \beta_{p+1j} \mathbf{h}_{p+1}$$

This change avoids calculating the weight α_{ij} for each encoder hidden state, as needed in probabilistic attention. Attention vectors are computed in constant time, notably reducing running time while keeping the accuracy, as shown in our experiments.

4 Experiments

We test the proposed approaches on the PTB treebank (Marcus et al., 1993) with standard splits.⁵

Table 1 compares parsing accuracy of all linearizations proposed in Section 2 to state-of-the-art fully-supervised transition-based constituent parsing models. The results show that our enriched in-order linearization is the most suitable option implemented so far for seq2seq constituent parsing, outperforming all existing seq2seq approaches (even without beam-search decoding) and matching some transition-based models. We also demonstrate that the enriched top-down variant (equivalent to the bracketed (Vinyals et al., 2015)’s linearization) outperforms the regular top-down approach of Liu and Zhang (2017a). A trend that can also be seen in the

⁵Settings are detailed in Appendix A.3.

Parser	Beam	Strat	F1
Transition-based			
(Cross and Huang, 2016a)	n	bu	90.0
(Cross and Huang, 2016b)	n	bu	91.3
(Liu and Zhang, 2017b)	n	bu	91.3
(Fernández-G and Gómez-R, 2019)*	n	bu	91.7
(Dyer et al., 2016)*	n	td	91.2
(Fernández-G and Gómez-R, 2018)*	n	td	91.7
(Liu and Zhang, 2017b)*	n	in	91.8
(Fernández-G and Gómez-R, 2018)*	n	in	92.0
(Zhu et al., 2013)	y	bu	90.4
(Watanabe and Sumita, 2015)	y	bu	90.7
(Liu and Zhang, 2017c)	y	bu	91.7
(Fried and Klein, 2018)	y	in	92.2
Seq2seq			
(Vinyals et al., 2015)	y	td	88.3
(Ma et al., 2017)	y	bu	88.6
(Kamigaito et al., 2017)	y	td	89.5
(Liu et al., 2018)	y	td	91.2
(Suzuki et al., 2018)	y	td	91.2
(Liu and Zhang, 2017a)* (<i>baseline</i>)	n	td	90.5
Top-down SH-RE w/ det. attention*	n	td	90.7
Enriched top-down SH-RE*	n	td	90.7
In-order SH-RE*	n	in	90.9
Enriched in-order SH-RE*	n	in	91.3
w/o lexicalized SH transition*	n	in	91.2
w/ det. attention*	n	in	91.2
w/ beam-search*	y	in	91.6
Chart-based			
(Stern et al., 2017)	n	bu	91.8
(Gaddy et al., 2018)	n	bu	92.1
(Kitaev and Klein, 2018)	n	bu	93.6

Table 1: Accuracy comparison on PTB test set with greedy (*n*) or beam-search (*y*) decoding and with different strategies followed to parse or to linearize the input sentence (*bu*=bottom-up, *td*=top-down and *in*=in-order). Systems marked with * are implemented under the same framework.

in-order linearization, where the addition of more tokens (parametrized Reduce and lexicalized Shift transitions) to the vocabulary benefits model performance (a gain of 0.4 F-score points), meaning that seq2seq models make use of this additional information. In fact, we analysed the average length of output sequences and noticed that enriched variants with larger vocabulary tend to produce shorter sequences. We hypothesize that the extra information is helping the model to better contextualize tokens in the sequence during training, minimizing the prediction of wrong tokens at decoding time. Finally, we extend the implementation by Liu and Zhang (2017a) with 10-beam-search decoding and increase F-score by 0.3 points.

We also evaluate parsing speeds under the exact same conditions among our approach and the top-down (Dyer et al., 2016) and in-order (Liu and Zhang, 2017b) transition-based constituent parsers, implemented in the framework by Dyer

Parser	sent./s.
Transition-based	
(Dyer et al., 2016) (<i>top-down</i>)	38.78
(Liu and Zhang, 2017b) (<i>in-order</i>)	33.34
Seq2seq	
(Liu and Zhang, 2017a) (<i>top-down SH-RE</i>)	16.65
Top-down SH-RE w/ det. attention	37.93
Enriched in-order SH-RE	16.54
Enriched in-order SH-RE w/ det. attention	35.12

Table 2: Speed comparison on PTB test set.

et al. (2015).⁶ Table 2 shows how the proposed deterministic attention technique doubles the speed of the baseline model, putting it on par with stack-LSTM-based shift-reduce systems, which are considered one of the most efficient approaches for constituent parsing. We can also see from Table 1 that the presented mechanism is more beneficial in terms of accuracy for the top-down algorithm (increasing 0.2 points in F-score) than the in-order variant (suffering a drop of 0.1 points in F-score), as could be expected from our previous analysis of attention vectors.

Finally, at the bottom of Table 1, we show current state-of-the-art chart-based parsers. These approaches, while more accurate, are significantly slower than seq2seq and transition-based parsers, being less appealing for downstream applications where the speed is crucial.

5 Conclusion

We present significant accuracy and speed improvements in seq2seq constituent parsing. The proposed linearization techniques can be used by any off-the-self seq2seq model without building a specific algorithm or structure. In addition, any advances in seq2seq neural architectures or pre-trained transformer-based language models (Devlin et al., 2019) can be directly used to enhance our approach.

Acknowledgments

This work has received funding from the European Research Council (ERC), under the European Union’s Horizon 2020 research and innovation programme (FASTPARSE, grant agreement No 714150), from the ANSWER-ASAP project (TIN2017-85160-C2-1-R) from MINECO, and from Xunta de Galicia (ED431B 2017/01, ED431G 2019/01).

⁶Please note that the implementation by Dyer et al. (2015) is not optimized for speed, but it can be used as a common framework to compare different approaches.

References

- James Cross and Liang Huang. 2016a. Incremental parsing with minimal features using bi-directional LSTM. In *ACL (2)*. The Association for Computer Linguistics.
- James Cross and Liang Huang. 2016b. Span-based constituency parsing with a structure-label system and provably optimal dynamic oracles. In *EMNLP*, pages 1–11. The Association for Computational Linguistics.
- 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, 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, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 334–343.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent neural network grammars. In *HLT-NAACL*, pages 199–209. The Association for Computational Linguistics.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2018. [Dynamic oracles for top-down and in-order shift-reduce constituent parsing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1303–1313, Brussels, Belgium. Association for Computational Linguistics.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2019. [Faster shift-reduce constituent parsing with a non-binary, bottom-up strategy](#). *Artificial Intelligence*, 275:559 – 574.
- Daniel Fried and Dan Klein. 2018. [Policy gradient as a proxy for dynamic oracles in constituency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 469–476.
- David Gaddy, Mitchell Stern, and Dan Klein. 2018. [What’s going on in neural constituency parsers? an analysis](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 999–1010.
- Hidetaka Kamigaito, Katsuhiko Hayashi, Tsutomu Hirao, Hiroya Takamura, Manabu Okumura, and Masaaki Nagata. 2017. [Supervised attention for sequence-to-sequence constituency parsing](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 7–12, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). Cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015.
- Nikita Kitaev and Dan Klein. 2018. [Constituency parsing with a self-attentive encoder](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 2675–2685.
- Jiangming Liu and Yue Zhang. 2017a. [Encoder-decoder shift-reduce syntactic parsing](#). In *Proceedings of the 15th International Conference on Parsing Technologies, IWPT 2017, Pisa, Italy, September 20-22, 2017*, pages 105–114.
- Jiangming Liu and Yue Zhang. 2017b. [In-order transition-based constituent parsing](#). *Transactions of the Association for Computational Linguistics*, 5:413–424.
- Jiangming Liu and Yue Zhang. 2017c. [Shift-reduce constituent parsing with neural lookahead features](#). *TACL*, 5:45–58.
- Lemao Liu, Muhua Zhu, and Shuming Shi. 2018. [Improving sequence-to-sequence constituency parsing](#). In *AAAI Conference on Artificial Intelligence*.
- Chunpeng Ma, Lemao Liu, Akihiro Tamura, Tiejun Zhao, and Eiichiro Sumita. 2017. [Deterministic attention for sequence-to-sequence constituent parsing](#). In *AAAI Conference on Artificial Intelligence*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics*, 19:313–330.
- D. J. Rosenkrantz and P. M. Lewis. 1970. [Deterministic left corner parsing](#). In *Proceedings of the 11th Annual Symposium on Switching and Automata Theory (Swat 1970)*, SWAT ’70, pages 139–152, Washington, DC, USA. IEEE Computer Society.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017. [A minimal span-based neural constituency parser](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 818–827.

Jun Suzuki, Sho Takase, Hidetaka Kamigaito, Makoto Morishita, and Masaaki Nagata. 2018. [An empirical study of building a strong baseline for constituency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 612–618, Melbourne, Australia. Association for Computational Linguistics.

Oriol Vinyals, Lukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. [Grammar as a foreign language](#). In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15*, pages 2773–2781, Cambridge, MA, USA. MIT Press.

Taro Watanabe and Eiichiro Sumita. 2015. [Transition-based neural constituent parsing](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics, ACL 2015, 26-31 July 2015, Beijing, China, Volume 1: Long Papers*, pages 1169–1179.

Muhua Zhu, Yue Zhang, Wenliang Chen, Min Zhang, and Jingbo Zhu. 2013. [Fast and accurate shift-reduce constituent parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*, pages 434–443.

A Appendices

A.1 Top-down Transition System

In the top-down transition system defined by [Dyer et al. \(2016\)](#), parser configurations have the form $c = \langle \Sigma, B \rangle$, where Σ is a *stack* of constituents and B is the *buffer* that contains words from the input sentence. The top-down algorithm also provides three transitions (described in [Figure 3](#)) that can be used on the stack and the buffer (that initially contains the whole unparsed sentence) to build the final constituent tree. Concretely:

- a *Shift* transition is used to push words from the buffer into the stack,
- a *Non-Terminal-X* transition to push a non-terminal node X into the stack,
- and a *Reduce* transition to pop elements from the stack until a non-terminal node is found and create a new subtree with all these elements as its children, pushing this new constituent into the stack.

A.2 In-order Transition System

[Liu and Zhang \(2017b\)](#) define a transition system that builds a phrase structure tree in an in-order traversal order: the non-terminal node on top of

Shift: $\langle \Sigma, w_i | B \rangle \Rightarrow \langle \Sigma | w_i, B \rangle$

NT-X: $\langle \Sigma, B \rangle \Rightarrow \langle \Sigma | X, B \rangle$

Reduce: $\langle \Sigma | X | s_k | \dots | s_0, B \rangle \Rightarrow \langle \Sigma | X_{s_k \dots s_0}, B \rangle$

Figure 3: Transitions available in a top-down transition system (NT-X = Non-Terminal-X).

Shift: $\langle \Sigma, w_i | B, false \rangle \Rightarrow \langle \Sigma | w_i, B, false \rangle$

NT-X: $\langle \Sigma, B, false \rangle \Rightarrow \langle \Sigma | s_0 | X, B, false \rangle$

Reduce: $\langle \Sigma | s_k | X | s_{k-1} | \dots | s_0, B, false \rangle$
 $\Rightarrow \langle \Sigma | X_{s_k \dots s_0}, B, false \rangle$

Finish: $\langle \Sigma, B, false \rangle \Rightarrow \langle \Sigma, B, true \rangle$

Figure 4: Transitions available in an in-order transition system (NT-X = Non-Terminal-X).

the tree being built is only considered after the first child node is completed in the stack, building each subtree in a bottom-up manner, but choosing the non-terminal node on top before the new constituent is reduced. This transition system has parser configurations with the stack-buffer form $c = \langle \Sigma, B \rangle$ and uses the following actions (described in [Figure 4](#)):

- a *Shift* transition to move words from the buffer to the stack,
- a *Non-Terminal-X* transition to push a non-terminal node X into the stack as long as the first child of the future constituent is on top of the stack,
- a *Reduce* transition to pop all elements from the stack until the first non-terminal node is found, which is also popped together with the preceding element in the stack to build a new constituent with all of them as children of the non-terminal node,
- and, finally, a *Finish* transition to terminate the parsing process.

The in-order transition system is a combination of the classic bottom-up and the new top-down algorithms, providing advantages of both of them: the access to information from partial parses from the bottom-up approach, and the non-local outlook of the top-down approach.

Hyper-parameters	
BiLSTM encoder layers	2
BiLSTM encoder input size	100
BiLSTM encoder hidden size	200
LSTM decoder layers	1
LSTM decoder hidden size	400
POS tag embedding dimension	6
Pretrained word embedding dimension	100
Word embedding dimension	64
Label embedding dimension	20
Action embedding dimension	40
Attention hidden size	50
Initial learning rate	0.001
β_1, β_2	0.9
λ	10^{-6}

Table 3: Model hyper-parameters.

A.3 Data and Settings

Following common practice, we test the proposed approaches on the Wall Street Journal sections of the English Penn Treebank (Marcus et al., 1993) with standard splits: sections 2-21 are used as training data, section 22 for development and section 23 for testing.

We adopt stochastic gradient descent with Adam (Kingma and Ba, 2014) and hyper-parameter selection as (Liu and Zhang, 2017a), detailed in Table 3. In addition, we use predicted POS tags and pre-trained word embeddings (generated on the AFP portion of English Gigaword) as (Dyer et al., 2016; Liu and Zhang, 2017a,b).

All neural models are trained by minimizing the following cross-entropy loss objective with an l_2 regularization term:

$$\mathcal{L}(\theta) = - \sum_i \sum_j \log p_{y_{ij}} + \frac{\lambda}{2} \|\theta\|^2$$

where θ is the set of parameters, $p_{y_{ij}}$ is the probability of the j th token in the i th training example given by the model and λ is a regularization hyper-parameter. For further details about the neural architecture, the reader can refer to (Liu and Zhang, 2017a).

For our executions, we report the average accuracy and speed over 3 runs with random initialization and on a single CPU core.