

Span-based discontinuous constituency parsing: a family of exact chart-based algorithms with time complexities from $\mathcal{O}(n^6)$ down to $\mathcal{O}(n^3)$

Caio Corro*

Université Paris-Saclay, CNRS, LIMSIS, 91400, Orsay, France.

caio.corro@limsi.fr

Abstract

We introduce a novel chart-based algorithm for span-based parsing of discontinuous constituency trees of block degree two, including ill-nested structures. In particular, we show that we can build variants of our parser with smaller search spaces and time complexities ranging from $\mathcal{O}(n^6)$ down to $\mathcal{O}(n^3)$. The cubic time variant covers 98% of constituents observed in linguistic treebanks while having the same complexity as continuous constituency parsers. We evaluate our approach on German and English treebanks (NEGRA, TIGER and DISCPTB) and report state-of-the-art results in the fully supervised setting. We also experiment with pre-trained word embeddings and BERT-based neural networks.

1 Introduction

Syntactic parsing aims to recover the latent syntactic relations between words in a sentence, expressed in a given syntactic formalism. In this paper, we focus on constituency trees where the syntactic structure is described by the means of a hierarchical structure composed of nodes: words are leaf nodes whereas internal nodes represent labeled constituents or phrases, see Figure 1. Constituency trees can broadly be classified into two categories. On the one hand, in a *continuous* constituent tree, each node must dominate a contiguous sequence of words.¹ On the other hand, in a *discontinuous* constituent tree, a node can dominate a non-contiguous sequence of words. It has been argued that modeling discontinuity is unavoidable, see for example McCawley (1982) and Bunt et al. (1987) for English and Müller (2004) for German.

Phrase-structure grammars have been proposed to parse and generate constituency trees. For example, Context-Free Grammars (CFG) are able to process continuous constituent trees whereas Multiple Context Free Grammars (Seki et al., 1991, MCFG) and Linear Context-Free Rewriting System (Vijay-Shanker et al., 1987, LCFRS) are able to process their discontinuous counterpart. CFGs have been widely studied for practical parsing due to the availability of time-efficient parsing algorithms based on *chart-based algorithms* (i.e. dynamic programming): parsing a sentence of length n is a $\mathcal{O}(gn^3)$ problem where g is a grammar related constant (Kasami, 1966; Younger, 1967; Cocke, 1969). However, parsing algorithms for MCFGs and LCFRSs are deemed to be impractical despite their polynomial-time complexity (see Section 2). Therefore, most of the experimental work in this field has been limited to parsing short sentences, e.g. sentences that contains less than 40 words (Kallmeyer and Maier, 2010; Evang and Kallmeyer, 2011; Maier et al., 2012; Kuhlmann and Nivre, 2006).

Advances in machine learning led to the development of constituency parsers that are not based on phrase-structure grammars. Instead, the prediction step only ensures the well-formedness of the resulting structure and does not enforce compliance of the syntactic content represented by the structure. For example, a verbal phrase is not constrained to contain a verb. As such, they can be assimilated to the mainstream approach to bi-lexical dependency parsing where one consider candidate outputs only in a restricted class of graphs: non-projective (McDonald et al., 2005), projective (Eisner, 1997) or bounded block degree and well-nested spanning aborescences (Gómez-Rodríguez et al., 2009; Gómez-Rodríguez

* Work partially done while the author was a postdoc at University of Amsterdam with Ivan Titov.

¹The set of words that a node dominates is the set of leaf nodes in the subtree for which this node is the root.

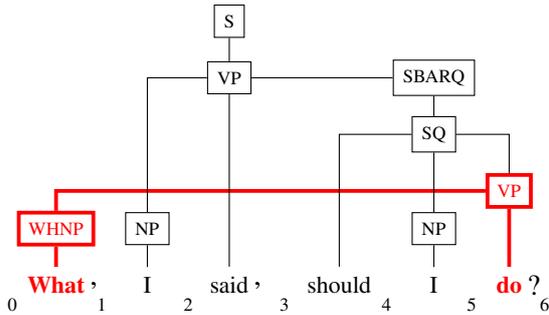


Figure 1: Example of discontinuous constituency tree. The bold red VP node dominates two sequences of words: "What" and "do". All other nodes are continuous. Numbers below the sentence are interstice indices used in the algorithm description.

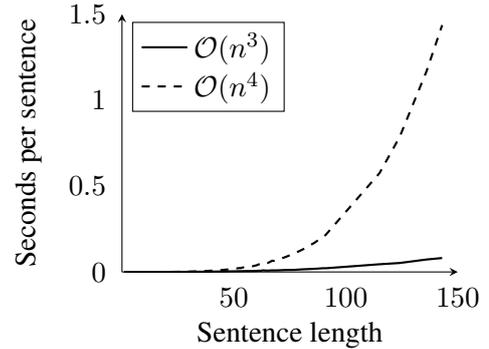


Figure 2: Execution time per sentence length of the chart-based algorithm for the $O(n^3)$ (solid line) and $O(n^4)$ (dashed lines) variants.

et al., 2011; Corro et al., 2016), among others (Kuhlmann and Nivre, 2006; Satta and Kuhlmann, 2013; Pitler et al., 2012). These approaches assume that intricate relations in the syntactic content can be implicitly learned by the scoring function.

Span-based parsing is a grammarless approach to constituency parsing that decomposes the score of a tree solely into the score of its constituents, originally proposed for continuous constituency parsing (Hall et al., 2014; Stern et al., 2017; Cross and Huang, 2016).² Recovering the tree of highest score can be done exactly using a slightly updated CYK algorithm or using inexact³ methods like top-down or transition based algorithms. This approach has obtained state-of-the-art results for continuous constituency parsing (Stern et al., 2017; Kitaev and Klein, 2018; Kitaev et al., 2019). In this work, we propose the first span-based parser with an exact decoding algorithm for discontinuous constituent parsing. To this end, we introduce a novel exact chart-based algorithm based on the parsing-as-deduction formalism (Pereira and Warren, 1983) that can parse constituent trees with a block degree of two, including ill-nested structures (see Section 3). which have been argued to be unavoidable to model natural languages (Chen-Main and Joshi, 2010). Despite its $O(n^6)$ time-complexity, where n is length of the input sentence, the algorithm is reasonably fast: all treebanks can be parsed without removing long sentences. Moreover, we observe that several deduction rules are of little use to retrieve trees present in treebanks. Therefore, we experiment with variants of the algorithm where we remove specific deduction rules. This leads to parsing algorithms with lower asymptotic complexity that experimentally produce accurate parses. Importantly, we show that a specific form of discontinuity can be parsed in $O(n^3)$, that is with the same asymptotic complexity as continuous constituency parsing.

Even with a constraint on the block degree, there are $O(n^4)$ prospective constituents that all have to be scored as we rely on exact decoding without further assumption and/or filtering. This would be too expensive in practice.⁴ Transition-based models address this problem by only scoring constituents that are requested during beam search. Although this is appealing on CPU, this lazy computation of scores cannot fully benefit from modern GPU architectures to parallelize computation at test-time. In this work, we propose to decompose the score of a constituent into independent parts leading to a quadratic number of scores to compute. As such, we can rely on efficient dimension broadcasting and operation batching available on modern GPUs.

Our main contributions can be summarized as follows:

- we propose a new span-based algorithm for parsing discontinuous constituency trees of block degree two with exact decoding and reasonable execution time;

²In contrast, for example, to several transition systems that can incorporate scores related to actions that were executed during derivation, or to split point decision and left-right span scores in the parser of Stern et al. (2017).

³The term inexact refers to the fact that these methods are not guaranteed to recover the highest scoring structure.

⁴For example, in preliminary experiments we found that the neural network computing scores could not fit on a 12GB GPU.

- we propose a cubic-time algorithm that can parse a significant portion of discontinuous constituents in different corpora while having the same theoretical complexity as continuous constituency parsers;
- we report state-of-the-art parsing results on these treebanks in a fully supervised setting and experiment with pre-trained word embeddings, including BERT based models.

We release the C++/Python implementation of the parser.⁵

2 Related Work

Phrase-structure grammars: The LCRFS formalism has been widely used in the context of discontinuous constituency parsing, although MCFG and Simple Range Concatenation Grammars (Boullier, 1998) have been shown to be equivalent, see Seki et al. (1991) and Boullier (2004). Kallmeyer and Maier (2010) introduced the first practical chart-based LCRFS parser for German, which was subsequently applied to English (Evang and Kallmeyer, 2011). However, they restrict their data to sentences that contains fewer than 25 words. To improve parsing time, Maier et al. (2012) proposed an experimentally faster parser based on the A^* search algorithm together with a block degree two restriction. However, they still limit the sentence size to a maximum of 40 words. A single sentence of 40 words takes around 3 minute to be parsed, an impressive improvement over the parser of Kallmeyer and Maier (2010) that needs several hours, but still prohibitively slow for large scale parsing.

Graph based parsing: A different line of work proposed to explore constituency parsing as a dependency parsing problem.⁶ In other words, even if it is straightforward to represent constituency trees as hierarchical phrase structures, the same syntactic content can be represented with different mathematical objects (Rambow, 2010), including directed graphs commonly used for dependency parsing. Fernández-González and Martins (2015) reduced the (lexicalized) constituency parsing task to dependency parsing where the constituency structure is encoded into arc labels. Then, discontinuous constituency parsing is reduced to the labeled *Spanning Arborescence* problem which can be solved in quadratic time. The same reduction has also been used in a sequence-to-sequence framework (Fernández-González and Gómez-Rodríguez, 2020). Corro et al. (2017) proposed a joint supertagging and dependency parsing reduction where vertices represents supertags⁷ and labeled arcs encode combination operations (substitution and adjunction). The problem is then reduced to the labeled *Generalized Spanning Arborescence* problem which is a known NP-hard optimization problem (Myung et al., 1995). One benefit of these approaches is that they do not assume any restriction on the constituency structure: they can parse ill-nested structures and have no block degree restriction. However, they cannot impose such constraints which may be beneficial or required and they factor the score of a tree into dependency, supertag and/or label scores, which means that the learning objective is not directly related to the evaluation metric which focuses on constituents. Moreover, factorization rely on possibly erroneous heuristics (head-percolation tables) to lexicalize the original structure if the information is not present in the treebank. On the contrary, in this work, we directly score parts of the syntactic content (i.e. labeled constituents). Therefore, at training time we can optimize an objective directly related to the end-goal evaluation.

Transition systems: Lastly, transition-based parsers have been proposed, based on the idea of the SWAP transition for non-projective dependency parsing (Nivre, 2009), see Versley (2014) and following work based on shift-reduce strategy (Maier, 2015; Maier and Lichte, 2016; Stanojević and G. Alhama, 2017). These systems rely on the fact that a discontinuous tree can be transformed into a continuous one by changing word order in the input sentence. They do not require strong independence assumption on the scoring model which can be useful to encode richer information, especially for long-distance relationships. However, the number of transitions required to parse discontinuities can impact prediction accuracy. To alleviate this problem, two different approaches have been explored: Coavoux and Crabbé

⁵<https://github.com/FilippoC/disc-span-parser-release>

⁶Note that opposite line of work also exists, that is reducing dependency parsing to constituency parsing, see for example Maier and Kallmeyer (2010).

⁷A supertag is an elementary tree that encodes the sequence of lexicalized constituents for which a given word is the head, see Bangalore and Joshi (1999)

(2017) introduced a two stacks system coupled with a GAP transition and Maier and Lichte (2016) proposed the SHIFT-I transition to access non-local elements directly, therefore reducing the number of transitions. In exchange for a rich parameterization, transition systems lose optimality guarantees with respect to the scoring model and rely on greedy or beam-search decoding that can return sub-optimal solutions. These approaches achieve state-of-the-art results while being fast at test time (Coavoux and Cohen, 2019; Coavoux et al., 2019). On the contrary, our approach is exact with respect to the scoring model, i.e. it will always return the highest scoring structure.

3 Parsing Algorithm

We describe our algorithm using the *parsing-as-deduction* framework (Pereira and Warren, 1983). As such, our description is independent of the value one wants to compute, whether it be the (k-)best derivation(s), partition function or span marginals (Goodman, 1999).⁸ However, we will focus on argmax decoding.

We are interested in constituency trees of block degree two, including ill-nested trees. The block degree two constraint is satisfied if each node dominates at most two disjoint sequences of words. Let $w_1 \dots w_n$ be a sentence. A constituency tree over this sentence is ill-nested if it contains two nodes dominating disjoint sets of words $W^{(1)}$ and $W^{(2)}$ such that there exists $w_i, w_j \in W^{(1)}$ and $w_k, w_l \in W^{(2)}$ such that $i < k < j < l$ or $k < i < l < j$.

We first highlight some properties of span-based parsers:

- **Filtering:** Contrary to CFGs and LCFRS CKY-style parsers, there is no side-condition constraining allowed derivations in span-based parsers. The label of a constituent is independent of the label of its children.
- **Binarization:** Interestingly, span-based parsers do not require explicit binarization of the constituency structure. Although grammar based parsers require binarization of the grammar production rules and therefore of the constituency structure to ensure tractable complexity, span-based parsers can take care of binarization *implicitly* by introducing a supplementary constituency label with a fixed null score.
- **Unary rules:** We follow Stern et al. (2017) and merge unary chains into a single constituent with a new label, e.g. the chain $\text{SBARQ} \rightarrow \text{SQ}$ will result in a single constituent labeled SBARQ_SQ .

3.1 Items

Let \mathcal{N} be the set of non-terminals (labels) and n the length of the input sentence. We define spans with interstice indices instead of word indices, see Figure 1. Items manipulated by our deduction rules are 5-tuples $[A, i, k, l, j]$ where $A \in \mathcal{N} \cup \{\emptyset\}$ is the constituent label with value \emptyset indicating an empty span used for implicit binarization. Given that each item represent a constituent, we will use the same notation to refer to the chart item and to the linguistic structure interchangeably. Indices $i, j \in \mathbb{N}, k, l \in \mathbb{N} \cup \{-\}$ defines the constituent span as follows:

- if the constituent is continuous, then $k = l = -$ and $0 \leq i < j \leq n$;
- otherwise, the constituent is discontinuous (with a single gap) and $0 \leq i < k$ and $l < j \leq n$, with $k < l$, define its left and right spans, respectively. For example, the tree in Figure 1 contains the discontinuous constituent $[VP, 0, 1, 5, 6]$.

3.2 Axioms and goal

Axiom items are word level constituents, i.e. items of the form $[A, i, -, -, i + 1]$ with $0 \leq i < n$ and $A \in \mathcal{N} \cup \{\emptyset\}$. In our experiments, axioms can have a null label, i.e. $A = \emptyset$, because we do not include part-of-speech tags as leaves of the constituency tree.

⁸Note that parsing without grammatical constraints results in all sentences having a non-empty parse forest, therefore the recognition problem is ill-defined.

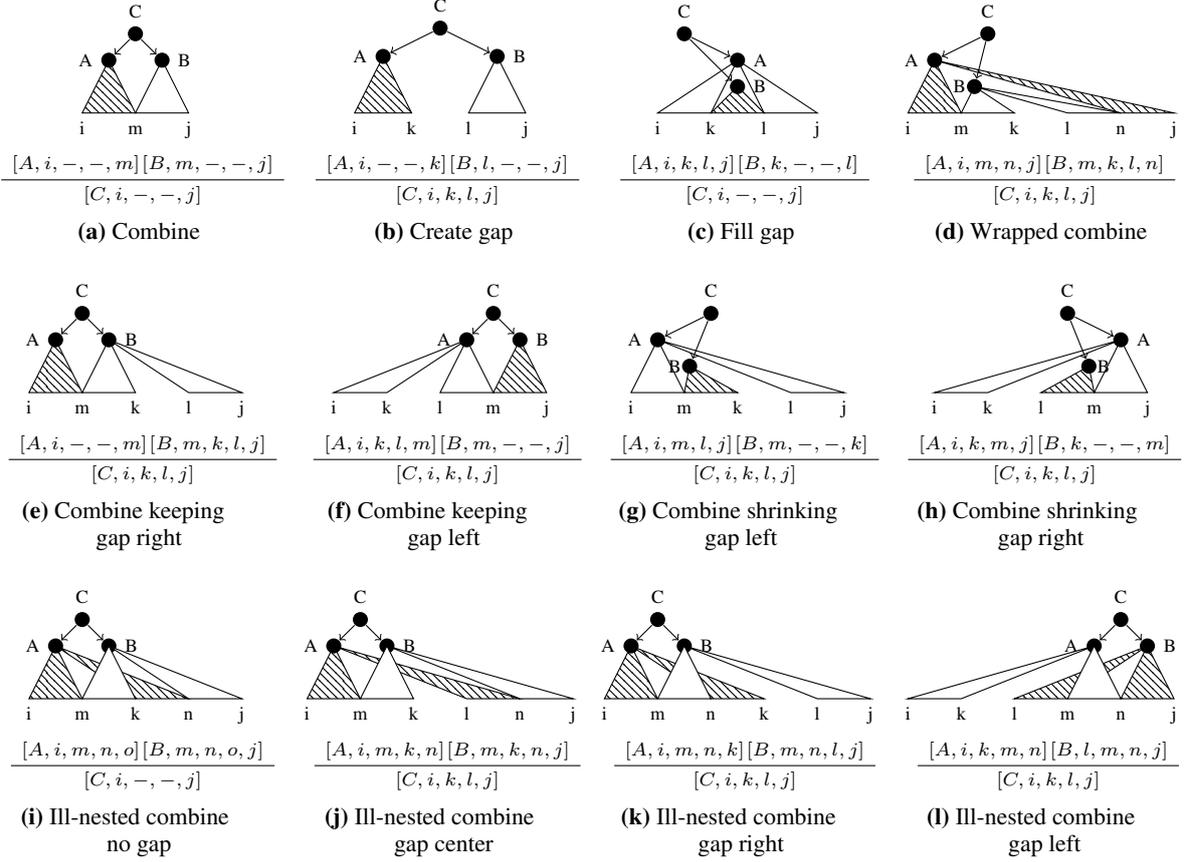


Figure 3: Deduction rules of our algorithm.

The goal item is defined as $[A, 0, -, -, n]$ with $A \in \mathcal{N} \cup \{\emptyset\}$. Similarly, in our experiments, the goal can have a null label, so we can parse empty trees and disconnected structures present in treebanks without further pre/post-processing steps.

3.3 Deduction rules

The deduction rules used to derive the goal from axioms are listed on Figure 3. Each rule takes exactly two antecedents. Note that rule (a) is the single rule used by span-based continuous constituency parsers.

Rule (b) creates a discontinuous constituent from two continuous constituents. The set of rules (e)-(f)-(g)-(h) (resp. (c)) allow to combine one discontinuous and one continuous constituent to produce a discontinuous one (resp. a continuous one).

Finally, there are rules that combine two discontinuous antecedents. Rule (d) is the only such rule that is allowed for building well-nested trees. The other four rules (i)-(j)-(k)-(l) are used for the construction of ill-nested trees. As such, it is easy to control whether ill-nested structures are permitted or not by including or excluding them.

3.4 Soundness and completeness

On the one hand, the algorithm is sound by definition because:

- items cannot represent constituents with a gap degree strictly greater to two;
- every rule deduces an item representing a constituent spanning a greater number of words, therefore they cannot construct invalid trees where the parent of constituent spans fewer words than one of its children.

On the other hand, completeness can be proved by observing that every possible binary parent-children combination can be produced by one of the rule. For the non-binary case, completeness follows the fact

the fact a constituent with 3 or more children can be built by first deriving intermediary constituents with label \emptyset .

Note that due to implicit binarization, a non-binary tree can be constructed by different sequences of deductions. Therefore, special care must be taken for computing partition function and marginal probabilities. As we are not interested by these values in this work, we do not dwell into this issue.

3.5 Complexity

The space and time complexity can be inferred from item structures and deduction rules: the space complexity is $\mathcal{O}(|\mathcal{N}|n^4)$ and time complexity is $\mathcal{O}(|\mathcal{N}|^3n^6)$. In practice, we decompose the score of a tree into the sum of the score of its constituents only and there are no constraints between antecedents and consequent labels. Therefore, we can build intermediary unlabeled items as follow:

$$\frac{[A, i, k, l, j]}{[i, k, l, j]}$$

which replace antecedents in every rule in Figure 3. With this update, the time complexity is linear in the number of labels, that is, $\mathcal{O}(|\mathcal{N}|n^6)$.

We instantiate variants of the algorithm than cannot parse the full family of block degree two trees but that can still fit most actual linguistic structures present in treebanks, with a lower time complexity. By using only rules **(a)**, **(b)** and **(c)** we can build a parser with a $\mathcal{O}(n^4)$ time complexity. In the next section, we show that this specific variant can be optimized into a $\mathcal{O}(n^3)$ time parser. By adding rules **(e)**, **(f)**, **(g)**, **(h)** and **(i)** we build a $\mathcal{O}(n^5)$ parser. Finally, we construct $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ well-nested parsers by excluding rules **(i)**, **(j)**, **(k)** and **(l)**.

3.6 Cubic time discontinuous constituency parser

A specific variant uses only deduction rules **(a)**, **(b)** and **(c)** from Figure 3, leading to a $\mathcal{O}(n^4)$ space and time complexity. In this setting, there is no way to combine two items representing discontinuous constituents or to have a discontinuous constituent that has a discontinuous child in the resulting parse tree. In this section, we prove that the family of trees induced by this variant of the parser can actually be parsed with a $\mathcal{O}(n^3)$ time complexity, that is equivalent to continuous constituency parsers.

The intuition goes as follows. We could replace rules **(b)** and **(c)** with the single rule **(m)** in Figure 4 where the right hand side condition $D \in \mathcal{N}$ induces the existence of a discontinuous constituent with label D ,⁹ with its left part spanning words i to k and right part spanning words l to j . However, observe that this new rule realizes two tests that could be done independently:

1. the right span boundary of the first antecedent must match the left span boundary of the second one;
2. the right span boundary of the second antecedent must match the left span boundary of the third antecedent.

Therefore, we can break the deduction into two sequential deductions, first testing the "k" boundary then the "l" boundary.¹⁰

To this end, we build a parser based on 4-tuple items $[A, \tau, i, j]$ where $\tau \in \{\top, \perp\}$ indicates whether the item represents a continuous constituent ($\tau = \top$) or an incomplete discontinuous constituent ($\tau = \perp$). More precisely, an item $[A, \perp, i, j]$ represents a partial discontinuous constituent who would be represented as $[A, i, ?, j, ?]$ in the previous formalization. The right boundary of its two spans are unknown: the one of the left span has been "forgotten" and the one on the right span is yet to be determined. The deduction rules of this new parser are listed on Figure 4, with axioms $[A, \tau, i, i + 1]$, $0 \leq i < n$, and goal $[A, \tau, 0, n]$.

⁹Without loss of generality, we assume the label D is not null. Although it could be without changing the overall idea, we would just add an extra way to do implicit binarization that can already be handled with rule **(a)**.

¹⁰This idea of breaking up simultaneous tests in a deduction rule has been previously proposed for improving time complexity of lexicalized grammar parsers (Eisner and Satta, 1999; Eisner and Satta, 2000)

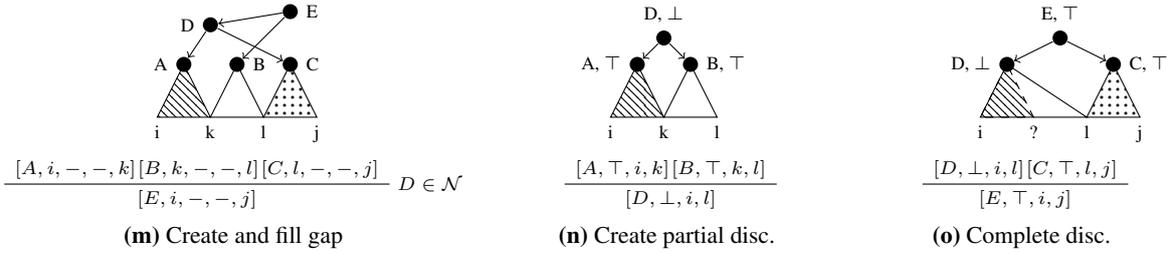


Figure 4: (m) The create gap and fill gap rules can be merged into a single rule if there are no other rule with discontinuous antecedents in the parser. (n)-(o) Rules for the cubic time discontinuous constituency parser. The rule to combine two continuous constituents follows the previous one.

Note that this cubic time algorithm imposes an additional restriction for weighted parsing: the score of discontinuous constituent must be divided into smaller sub-parts, which we do in practice for all deduction systems due to computational reasons.

We report the running time per sentence length for the $\mathcal{O}(n^4)$ and $\mathcal{O}(n^3)$ parsers in Figure 2. As expected, the running time of the cubic time parser is way lower for long sentence.

4 Experiments

We experiment on the Discontinuous Penn Treebank (Marcus et al., 1993; Evang and Kallmeyer, 2011, DISCPTB) with standard split, the TIGER treebank (Brants et al., 2002) with the SPMRL 2014 shared task split (Seddah et al., 2014) and the NEGRA treebank (Skut et al., 1997) with the split proposed by Dubey and Keller (2003).

4.1 Data coverage

One important question is whether our parser has a good coverage of the dataset as we can only retrieve constituents of block degree one and two. We report the maximum recall that our parser can achieve in its different variants in Table 1.

First, we observe that our cubic time parser can recover 98% of all constituents in the three treebanks, or around 80% of constituents of block degree of exactly two. Second, the $\mathcal{O}(n^5)$ variant of the parser can recover more than 99% of all treebanks, and, interestingly, there is almost no coverage change when moving to the full deduction system. If we consider the parsers with well-nested restriction, the $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ variants have the same coverage in German datasets and the later can only recover 2 additional constituents in the English treebanks. If we include ill-nested construction, the difference is either 2 (DISCPTB and NEGRA) or 8 (TIGER) constituents. In practice, we observed that both $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ variants predict the same analysis.

4.2 Neural parameterization

We use a neural architecture based on bidirectional LSTMs detailed in Appendix 6.

Constituent scores Even with the block degree two restriction, there is a larger number (quartic!) of constituent scores to compute. In early experiments, we observed that weighting such a number of constituents without further decomposition blow up the neural network memory usage and was prohibitively slow. Therefore, we introduce a score decomposition that results in a quadratic number of scores to compute and that can be efficiently parallelized on GPU using batched matrix operations.¹¹

We decompose the score of a constituent $[A, i, k, l, j]$ as the sum of a score associated with its outer boundaries (i.e. indices i and j) and one with its gap boundaries (i.e. indices k and l). The score of

¹¹Technical details about the implementation are out of the scope of this paper. However, they are described in our implementation that will be publicly available.

		Continuous	$\mathcal{O}(n^3)$	$\mathcal{O}(n^5)$ / WN	$\mathcal{O}(n^5)$	$\mathcal{O}(n^6)$ / WN	$\mathcal{O}(n^6)$
DPTB	All	98.16	99.46	99.81	99.83	99.81	99.83
	BD \leq 2	98.32	99.63	99.98	99.99	99.98	100.00
	BD = 2	0.00 (0)	78.27 (10713)	99.15 (13572)	99.98 (13685)	99.17 (13574)	100.00 (13687)
TIGER	All	94.51	98.61	99.37	99.49	99.37	99.49
	BD \leq 2	94.99	99.11	99.88	99.99	99.88	100.00
	BD = 2	0.00 (0)	82.39 (15324)	97.65 (18161)	99.95 (18590)	97.65 (18161)	100.00 (18598)
NEGRA	All	94.37	98.59	99.32	99.46	99.32	99.46
	BD \leq 2	94.87	99.12	99.85	99.99	99.85	100.00
	BD = 2	0.00 (0)	82.91 (6106)	97.24 (7161)	99.97 (7362)	97.24 (7161)	100.00 (7364)

Table 1: Maximum constituent recall that be can obtained using a continuous constituency parser and all the variant of our parser in three settings: considering all constituents, considering constituents with a block degree less or equal to two and exactly two. For the last case, we also report the number of constituents. We do not remove punctuation. The analysis is done with the concatenation of train, development and test sets.

	NEGRA		TIGER		DISCPTB	
	F1	Disc. F1	F1	Disc. F1	F1	Disc. F1
Fully supervised						
Fernández-González and Martins (2015)	77.0				77.3	
Versley (2016)			79.5			
Gebhardt (2018)			75.1			
Corro et al. (2017)					89.2	
Coavoux and Crabbé (2017)			79.3			
Coavoux et al. (2019)	83.2	54.6	82.7	55.9	91.0	71.3
Coavoux and Cohen (2019)	83.2	56.3	82.5	55.9	90.9	67.3
Fernández-González and Gómez-Rodríguez (2020)	83.7	54.7	84.6	57.9		
This work, $\mathcal{O}(n^3)$	86.2	54.1	85.5	53.8	92.7	64.2
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$, well-nested	84.9	46.1	84.8	50.4	92.6	62.6
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$	84.9	46.2	84.9	51.0	92.6	62.9
+ gold part-of-speech tags						
Maier (2015)	77.0	19.8	74.7	18.8		
Coavoux and Crabbé (2017)	82.2	50.0	81.6	49.2		
Corro et al. (2017)			81.6		90.1	
Semi-supervised: pre-trained word embeddings						
Stanojević and Alhama (2017)			77.0			
Fernández-González and Gómez-Rodríguez (2020), with pred tags	85.4	58.8	85.3	59.1		
Fernández-González and Gómez-Rodríguez (2020), without pred tags	85.7	58.6	85.7	60.4		
This work, $\mathcal{O}(n^3)$	86.3	56.1	85.2	51.2	92.9	64.9
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$, well-nested	85.6	52.9	84.9	50.4	92.6	59.4
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$	85.6	53.0	84.9	51.0	92.6	59.7
+ gold POS tags						
Stanojević and Alhama (2017)	82.9		81.6			
Fernández-González and Gómez-Rodríguez (2020)	86.1	59.9	86.3	60.7		
Semi-supervised: BERT						
This work, $\mathcal{O}(n^3)$	91.6	66.1	90.0	62.1	94.8	68.9
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$, well-nested	90.5	58.8	89.3	57.8	94.5	64.5
This work, $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$	90.6	59.6	89.3	58.7	94.5	64.7

Table 2: Discontinuous constituency parsing results on the three test sets. The $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ variants produced exactly the same results in all settings.

	NEGRA			TIGER			DISCPTB		
	D. recall	D. prec.	D. F1	D. recall	D. prec.	D. F1	D. recall	D. prec.	D. F1
$\mathcal{O}(n^3)$	42.0	76.0	54.1	40.6	79.8	53.8	49.7	90.8	64.2
$\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$, wn	47.0	45.2	46.1	46.9	54.3	50.4	63.8	61.4	62.6
$\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$	47.3	45.2	46.2	47.8	54.8	51.0	64.0	61.8	62.9

Table 3: Detailed discontinuous constituency parsing results for the fully supervised model.

	nn	$\mathcal{O}(n^3)$	$\mathcal{O}(n^4)$	$\mathcal{O}(n^5)$, wn	$\mathcal{O}(n^5)$	$\mathcal{O}(n^6)$, wn	$\mathcal{O}(n^6)$
NEGRA	1.74	0.35	1.10	3.73	4.48	8.82	22.58
TIGER	7.73	2.81	12.96	98.44	133.22	507.84	1841.95
DISCPTB	4.67	2.13	6.70	19.35	22.71	43.00	105.98

Table 4: Total time in seconds to parse the full test sets: the nn column corresponds to the time taken by the forward pass of the neural network (maximum 5000 words per batch on a NVIDIA Tesla V100 SXM2 32 Go), each supplementary column is the time take by each variant of the chart-based algorithm (without any parallelization, one sentence at a time).

constituent is defined as:¹²

$$W_{A,i,k,l,j} = \begin{cases} \mathbf{S}_{A,i+1,j}^{(c. label)} + S_{i+1,j}^{(c. span)} & \text{if } i = k = -, \\ \mathbf{S}_{A,i+1,j}^{(o. label)} + S_{i+1,j}^{(o. span)} + \mathbf{S}_{A,k+1,l}^{(g. label)} + S_{k+1}^{(g. span)} & \text{otherwise.} \end{cases}$$

where tensors $\mathbf{S}^{(c. label)}$, $\mathbf{S}^{(o. label)}$, $\mathbf{S}^{(g. label)} \in \mathbb{R}^{|\mathcal{V}| \times n \times n}$ and matrices $\mathbf{S}^{(c. span)}$, $\mathbf{S}^{(o. span)}$, $\mathbf{S}^{(g. span)} \in \mathbb{R}^{n \times n}$ are computed using the deep biaffine attention mechanism (Dozat and Manning, 2016). The tensor \mathbf{W} is never explicitly built: during the dynamic program execution we lazily compute requested constituent scores.

Training loss Span-based continuous constituency parsers are usually trained using a decomposable margin-based objective (Stern et al., 2017; Kitaev and Klein, 2018; Kitaev et al., 2019). This approach requires to repeatedly perform loss-augmented inference during training (Taskar et al., 2005), which can be prohibitively slow even when tractable. A current trend in dependency parsing is to ignore the tree structure and rely on negative log likelihood for head selection for each modifier word independently (Dozat and Manning, 2016; Zhang et al., 2017). We rely on a similar approach and use as training objective the negative log-likelihood loss independently for each span (continuous, outer and gap), adding a null label with a fixed 0 weight as label for spans that do not appear in the gold annotation.

4.3 Evaluation

We evaluate our parser on the test sets of the three treebanks. We report F-measure and discontinuous F-measure as computed using the disco-dop tool¹³ with the standard parameters in Table 2.

First, we observed that the $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ variants of our parsers produced exactly the same results in all settings. This may be expected as their cover of the original treebanks where almost similar. Second, surprisingly, the $\mathcal{O}(n^3)$ parser produced better results in term of F-measure than other variants in all cases. We report labeled discontinuous constituent recall and precision measures for the fully supervised model in Table 3. We observe that while the $\mathcal{O}(n^5)$ and $\mathcal{O}(n^6)$ have a better recall than the $\mathcal{O}(n^3)$ parser, their precision is drastically lower. This highlights a benefit of restricting the search space: the parser can retrieve less erroneous constituents leading to an improved overall precision.

Finally, in almost all cases, we achieve a novel state-of-the-art for the task in term of labeled F-measure. However, we are slightly lower when evaluating discontinuous constituent only. We suspect that this is due to the fact that our best parser is the one with the smallest search space.

4.4 Runtime

The runtime on the test sets of our approach is reported on Table 4. In all cases, the runtime is reasonably fast and we do not need to remove long sentences. Interestingly, most of the time is spent for computing scores with the neural network with the cubic time parser, even if we use batches to benefit from the GPU architecture while our chart-based algorithm is not parallelized on CPU.

¹²The +1 in tensor indices result of the fact that we use interstice indices for constituents but that the neural network layers focus on word indices.

¹³<https://github.com/andreasvc/disco-dop>

5 Conclusion

We proposed a novel family of algorithms for discontinuous constituency parsing achieving state-of-the-art results. Importantly, we showed that a specific set of discontinuous constituency trees can be parsed in cubic time while covering most of the linguistic structures observed in treebanks. Despite being based on chart-based algorithms, our approach is fast at test time and we can parse all sentences without pruning or filtering long sentences. Future research could explore neural architecture and training losses tailored to this approach for discontinuous constituency parsing.

Acknowledgments

We thank Djamé Seddah, Maximin Coavoux and Vlad Niculae for their comments and suggestions. We thank Serhii Havrylov for the help with the code. We thank Ivan Titov for comments on the early draft of this paper. We thank Laura Kallmeyer and Kilian Evang for providing us with the script for the discontinuous PTB. We thank Maximin Coavoux for the help with pre-processing the data. The project was supported by the Dutch National Science Foundation (NWO VIDI 639.022.518) and European Research Council (ERC Starting Grant BroadSem 678254). This work benefited from the Jean-Zay cluster.

References

- Srinivas Bangalore and Aravind K. Joshi. 1999. Supertagging: An approach to almost parsing. *Computational Linguistics*, 25(2):237–265.
- Pierre Boullier. 1998. Proposal for a natural language processing syntactic backbone.
- Pierre Boullier. 2004. Range concatenation grammars. In *New developments in parsing technology*, pages 269–289. Springer.
- Sabine Brants, Stefanie Dipper, Silvia Hansen, Wolfgang Lezius, and George Smith. 2002. The tiger treebank. In *Proceedings of the workshop on treebanks and linguistic theories*, volume 168.
- Harry Bunt, Jan Thesingh, and Ko van der Sloot. 1987. Discontinuous constituents in trees, rules, and parsing. In *Third Conference of the European Chapter of the Association for Computational Linguistics*, Copenhagen, Denmark, April. Association for Computational Linguistics.
- Joan Chen-Main and Aravind K. Joshi. 2010. Unavoidable ill-nestedness in natural language and the adequacy of tree local-MCTAG induced dependency structures. In *Proceedings of the 10th International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+10)*, pages 53–60, Yale University, June. Linguistic Department, Yale University.
- Maximin Coavoux and Shay B. Cohen. 2019. Discontinuous constituency parsing with a stack-free transition system and a dynamic oracle. 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 204–217, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Maximin Coavoux and Benoît Crabbé. 2017. Incremental discontinuous phrase structure parsing with the GAP transition. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1259–1270, Valencia, Spain, April. Association for Computational Linguistics.
- Maximin Coavoux, Benoît Crabbé, and Shay B. Cohen. 2019. Unlexicalized transition-based discontinuous constituency parsing. *Transactions of the Association for Computational Linguistics*, 7:73–89.
- John Cocke. 1969. *Programming Languages and Their Compilers: Preliminary Notes*. New York University, New York, NY, USA.
- Caio Corro, Joseph Le Roux, Mathieu Lacroix, Antoine Rozenknop, and Roberto Wolfler Calvo. 2016. Dependency parsing with bounded block degree and well-nestedness via Lagrangian relaxation and branch-and-bound. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 355–366, Berlin, Germany, August. Association for Computational Linguistics.

- Caio Corro, Joseph Le Roux, and Mathieu Lacroix. 2017. Efficient discontinuous phrase-structure parsing via the generalized maximum spanning arborescence. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1644–1654, Copenhagen, Denmark, September. Association for Computational Linguistics.
- James Cross and Liang Huang. 2016. Span-based constituency parsing with a structure-label system and provably optimal dynamic oracles. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1–11, Austin, Texas, November. Association for Computational Linguistics.
- Timothy Dozat and Christopher D Manning. 2016. Deep biaffine attention for neural dependency parsing. *arXiv preprint arXiv:1611.01734*.
- Amit Dubey and Frank Keller. 2003. Probabilistic parsing for German using sister-head dependencies. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, pages 96–103. Association for Computational Linguistics.
- 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, June. Association for Computational Linguistics.
- Jason Eisner and Giorgio Satta. 2000. A faster parsing algorithm for lexicalized tree-adjoining grammars. In *Proceedings of the Fifth International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+5)*, pages 79–84, Université Paris 7, May.
- Jason Eisner. 1997. Bilexical grammars and a cubic-time probabilistic parser. In *Advances in probabilistic and other parsing technologies*, pages 54–65.
- Kilian Evang and Laura Kallmeyer. 2011. PLCFRS parsing of English discontinuous constituents. In *Proceedings of the 12th International Conference on Parsing Technologies*, pages 104–116, Dublin, Ireland, October. Association for Computational Linguistics.
- Daniel Fernández-González and André F. T. Martins. 2015. Parsing as reduction. 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 1523–1533, Beijing, China, July. Association for Computational Linguistics.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2020. Discontinuous constituent parsing with pointer networks.
- Kilian Gebhardt. 2018. Generic refinement of expressive grammar formalisms with an application to discontinuous constituent parsing. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3049–3063, Santa Fe, New Mexico, USA, August. Association for Computational Linguistics.
- Carlos Gómez-Rodríguez, David Weir, and John Carroll. 2009. Parsing mildly non-projective dependency structures. In *Proceedings of the 12th Conference of the European Chapter of the ACL (EACL 2009)*, pages 291–299, Athens, Greece, March. Association for Computational Linguistics.
- Carlos Gómez-Rodríguez, John Carroll, and David Weir. 2011. Dependency parsing schemata and mildly non-projective dependency parsing. *Computational Linguistics*, 37(3):541–586.
- Joshua Goodman. 1999. Semiring parsing. *Computational Linguistics*, 25(4):573–606.
- David Hall, Greg Durrett, and Dan Klein. 2014. Less grammar, more features. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 228–237, Baltimore, Maryland, June. Association for Computational Linguistics.
- Laura Kallmeyer and Wolfgang Maier. 2010. Data-driven parsing with probabilistic linear context-free rewriting systems. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 537–545, Beijing, China, August. Coling 2010 Organizing Committee.
- Tadao Kasami. 1966. An efficient recognition and syntax-analysis algorithm for context-free languages. *Coordinated Science Laboratory Report no. R-257*.
- 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. 2018. Constituency parsing with a self-attentive encoder. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2676–2686, Melbourne, Australia, July. Association for Computational Linguistics.
- 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, July. Association for Computational Linguistics.
- Marco Kuhlmann and Joakim Nivre. 2006. Mildly non-projective dependency structures. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 507–514, Sydney, Australia, July. Association for Computational Linguistics.
- Wolfgang Maier and Laura Kallmeyer. 2010. Discontinuity and non-projectivity: Using mildly context-sensitive formalisms for data-driven parsing. In *Proceedings of the 10th International Workshop on Tree Adjoining Grammar and Related Frameworks (TAG+10)*, pages 119–126, Yale University, June. Linguistic Department, Yale University.
- Wolfgang Maier and Timm Lichte. 2016. Discontinuous parsing with continuous trees. In *Proceedings of the Workshop on Discontinuous Structures in Natural Language Processing*, pages 47–57, San Diego, California, June. Association for Computational Linguistics.
- Wolfgang Maier, Miriam Kaeshammer, and Laura Kallmeyer. 2012. PLCFRS parsing revisited: Restricting the fan-out to two. In *Proceedings of the 11th International Workshop on Tree Adjoining Grammars and Related Formalisms (TAG+11)*, pages 126–134, Paris, France, September.
- Wolfgang Maier. 2015. Discontinuous incremental shift-reduce parsing. 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 1202–1212, Beijing, China, July. Association for Computational Linguistics.
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank.
- James D. McCawley. 1982. Parentheticals and discontinuous constituent structure. *Linguistic Inquiry*, 13(1):91–106.
- 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, October. Association for Computational Linguistics.
- Stefan Müller. 2004. Continuous or discontinuous constituents? a comparison between syntactic analyses for constituent order and their processing systems. *Research on Language and Computation*, 2(2):209–257.
- Young-Soo Myung, Chang-Ho Lee, and Dong-Wan Tcha. 1995. On the generalized minimum spanning tree problem. *Networks*, 26(4):231–241.
- Joakim Nivre. 2009. Non-projective dependency parsing in expected linear time. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 351–359, Suntec, Singapore, August. Association for Computational Linguistics.
- Fernando C. N. Pereira and David H. D. Warren. 1983. Parsing as deduction. In *21st Annual Meeting of the Association for Computational Linguistics*, pages 137–144, Cambridge, Massachusetts, USA, June. Association for Computational Linguistics.
- Emily Pitler, Sampath Kannan, and Mitchell Marcus. 2012. Dynamic programming for higher order parsing of gap-minding trees. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 478–488, Jeju Island, Korea, July. Association for Computational Linguistics.
- Owen Rambow. 2010. The simple truth about dependency and phrase structure representations: An opinion piece. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 337–340, Los Angeles, California, June. Association for Computational Linguistics.

- Giorgio Satta and Marco Kuhlmann. 2013. Efficient parsing for head-split dependency trees. *Transactions of the Association for Computational Linguistics*, 1:267–278.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. 2013. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks.
- Djamé Seddah, Sandra Kübler, and Reut Tsarfaty. 2014. Introducing the spmrl 2014 shared task on parsing morphologically-rich languages. In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 103–109.
- Hiroyuki Seki, Takashi Matsumura, Mamoru Fujii, and Tadao Kasami. 1991. On multiple context-free grammars. *Theoretical Computer Science*, 88(2):191–229.
- Wojciech Skut, Brigitte Krenn, Thorsten Brants, and Hans Uszkoreit. 1997. An annotation scheme for free word order languages. *arXiv preprint cmp-lg/9702004*.
- Miloš Stanojević and Raquel G. Alhama. 2017. Neural discontinuous constituency parsing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1666–1676, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Miloš Stanojević and Raquel G. Alhama. 2017. Neural discontinuous constituency parsing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1666–1676, Copenhagen, Denmark, September. Association for Computational Linguistics.
- 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 (Volume 1: Long Papers)*, pages 818–827, Vancouver, Canada, July. Association for Computational Linguistics.
- Ben Taskar, Vassil Chatalbashev, Daphne Koller, and Carlos Guestrin. 2005. Learning structured prediction models: A large margin approach. In *Proceedings of the 22nd international conference on Machine learning*, pages 896–903. ACM.
- Yannick Versley. 2014. Experiments with easy-first nonprojective constituent parsing. In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 39–53, Dublin, Ireland, August. Dublin City University.
- Yannick Versley. 2016. Discontinuity (re)²-visited: A minimalist approach to pseudoprojective constituent parsing. In *Proceedings of the Workshop on Discontinuous Structures in Natural Language Processing*, pages 58–69, San Diego, California, June. Association for Computational Linguistics.
- K. Vijay-Shanker, David J. Weir, and Aravind K. Joshi. 1987. Characterizing structural descriptions produced by various grammatical formalisms. In *25th Annual Meeting of the Association for Computational Linguistics*, pages 104–111, Stanford, California, USA, July. Association for Computational Linguistics.
- Daniel H Younger. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and control*, 10(2):189–208.
- Xingxing Zhang, Jianpeng Cheng, and Mirella Lapata. 2017. Dependency parsing as head selection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 665–676, Valencia, Spain, April. Association for Computational Linguistics.

6 Neural parameterization

In this appendix, we describe the different components of our neural network. If unspecified, parameters are initialized with Pytorch default initialization.

6.1 Word-level features

We use three kind of word-level features: word embeddings, character embeddings and, for a few experiments, part-of-speech embeddings. All embeddings are concatenated to form word-level embeddings.

Word embeddings can either be pre-trained or trained end-to-end. In the case of pre-trained word embeddings, we fix them and sum them with end-to-end learned word embeddings initialized at 0.

Character embeddings are fed to a BiLSTM. The hidden states of the two endpoints are then concatenated together. Words are truncated to 20 characters for this feature.

Name	Dimension
Word embeddings	300
Character embeddings	64
Character BiLSTM	100
Character BiLSTMs layer	1
Sentence BiLSTMs	800
Sentence BiLSTMs layer	1
Sentence BiLSTMs stack	2
Span projection	500
Label projection	100

Table 5: Hyperparameters

6.2 Sentence-level features

We follow (Kiperwasser and Goldberg, 2016) by using two stacked BiLSTM, i.e. the input of the second BiLSTM is the concatenation of the forward and backward hidden states of the first one. All LSTMs have a single layer. Projection matrices are initialized with the orthogonal approach proposed by (Saxe et al., 2013) and bias vectors are initialized to 0.

For models using Bert, we learn a convex combination of the last 4 layers, in a similar spirit to ELMO. When word are tokenized in subwords by the Bert tokenizer, we use the embedding of the first sub-token.

6.3 Output weights

We have two different output layers. First, we predict part-of-speech tags with a linear projection on top of the hidden states of the first BiLSTM. During training, we use an auxiliary negative log-likelihood loss. Second, after the second BiLSTM we add the biaffine layers to compute span scores.

6.4 Hyperparameters

We report the dimensions of the building blocks of the network in Table 5.

We optimize the parameters with the Adam variant of stochastic gradient descent descent with mini-batches containing at most 5000 words for 200 epochs. We apply dropout with ratio 0.3 before the input of the character BiLSTM, before the first stack of sentence-level BiLSTM and after the second one by following the methodology of Dozat and Manning (2016).