

A Conditional Splitting Framework for Efficient Constituency Parsing

Thanh-Tung Nguyen^{†¶}, Xuan-Phi Nguyen^{†¶}, Shafiq Joty^{¶§}, Xiaoli Li^{†¶}

[¶]Nanyang Technological University

[§]Salesforce Research Asia

[†]Institute for Infocomm Research, A-STAR

Singapore

{ng0155ng@e.;nguyenxu002@e.;srjoty@}ntu.edu.sg

xlli@i2r.a-star.edu.sg

Abstract

We introduce a generic seq2seq parsing framework that casts constituency parsing problems (syntactic and discourse parsing) into a series of conditional splitting decisions. Our parsing model estimates the conditional probability distribution of possible splitting points in a given text span and supports efficient top-down decoding, which is linear in number of nodes. The conditional splitting formulation together with efficient beam search inference facilitate structural consistency without relying on expensive structured inference. Crucially, for discourse analysis we show that in our formulation, discourse segmentation can be framed as a special case of parsing which allows us to perform discourse parsing without requiring segmentation as a pre-requisite. Experiments show that our model achieves good results on the standard syntactic parsing tasks under settings with/without pre-trained representations and rivals state-of-the-art (SoTA) methods that are more computationally expensive than ours. In discourse parsing, our method outperforms SoTA by a good margin. Our source code will be publicly available.

1 Introduction

A number of formalisms have been introduced to analyze natural language at different linguistic levels. This includes syntactic structures in the form of phrasal and dependency trees, semantic structures in the form of meaning representations (Banarescu et al., 2013; Artzi et al., 2013), and discourse structures with Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) or Discourse-LTAG (Webber, 2004). Many of these formalisms have a *constituency* structure, where textual units (e.g., phrases, sentences) are organized into nested constituents. For example, Figure 1 shows examples of a phrase structure tree and a sentence-level discourse tree (RST) that respectively represent

how the phrases and clauses are hierarchically organized into a constituency structure. Developing efficient and effective parsing solutions has always been a key focus in NLP. In this work, we consider both phrasal (syntactic) and discourse parsing.

In recent years, neural end-to-end parsing methods have outperformed traditional methods that use grammar, lexicon and hand-crafted features. These methods can be broadly categorized based on whether they employ a greedy transition-based, a globally optimized chart parsing or a greedy top-down algorithm. Transition-based parsers (Dyer et al., 2016; Cross and Huang, 2016; Liu and Zhang, 2017; Wang et al., 2017) generate trees auto-regressively as a form of shift-reduce decisions. Though computationally attractive, the local decisions made at each step may propagate errors to subsequent steps due to exposure bias (Bengio et al., 2015). Moreover, there may be mismatches in shift and reduce steps, resulting in invalid trees.

Chart based methods, on the other hand, train neural scoring functions to model the tree structure globally (Durrett and Klein, 2015; Gaddy et al., 2018; Kitaev and Klein, 2018; Zhang et al., 2020b; Joty et al., 2012). By utilizing dynamic programming, these methods can perform exact inference to combine these constituent scores into finding the highest probable tree. However, they are generally slow with at least $\mathcal{O}(n^3)$ time complexity. Greedy top-down parsers find the split points recursively and have received much attention lately due to their efficiency, which is usually $\mathcal{O}(n^2)$ (Stern et al., 2017a; Shen et al., 2018; Lin et al., 2019; Nguyen et al., 2020). However, they still suffer from exposure bias, where one incorrect splitting step may affect subsequent steps.

Discourse parsing in RST requires an additional step – *discourse segmentation* which involves breaking the text into contiguous clause-like units called Elementary Discourse Units or EDUs

(Figure 1). Traditionally, segmentation has been considered separately and as a prerequisite step for the parsing task which links the EDUs (and larger spans) into a discourse tree (Soricut and Marcu, 2003; Joty et al., 2012; Wang et al., 2017). In this way, the errors in discourse segmentation can propagate to discourse parsing (Lin et al., 2019).

In this paper, we propose a generic top-down neural framework for constituency parsing that we validate on both syntactic and sentence-level discourse parsing. Our main contributions are:

- We cast the constituency parsing task into a series of conditional splitting decisions and use a seq2seq architecture to model the splitting decision at each decoding step. Our parsing model, which is an instance of a Pointer Network (Vinyals et al., 2015a), estimates the pointing score from a span to a splitting boundary point, representing the likelihood that the span will be split at that point and create two child spans.
- The conditional probabilities of the splitting decisions are optimized using a cross entropy loss and structural consistency is maintained through a global pointing mechanism. The training process can be fully parallelized without requiring structured inference as in (Shen et al., 2018; Gómez and Vilares, 2018; Nguyen et al., 2020).
- Our model enables efficient top-down decoding with $\mathcal{O}(n)$ running time like transition-based parsers, while also supporting a customized beam search to get the best tree by searching through a reasonable search space of high scoring trees. The beam-search inference along with the structural consistency from the modeling makes our approach competitive with existing structured chart methods for syntactic (Kitaev and Klein, 2018) and discourse parsing (Zhang et al., 2020b). Moreover, our parser does not rely on any hand-crafted features (not even part-of-speech tags), which makes it more efficient and be flexible to different domains or languages.
- For discourse analysis, we demonstrate that our method can effectively find the segments (EDUs) by simply performing one additional step in the top-down parsing process. In other words, our method can parse a text into the discourse tree without needing discourse segmentation as a prerequisite; instead, it produces the segments as a by-product. To the best of our knowledge, this is the first model that can perform segmentation and parsing in a single embedded framework.

In the experiments with English Penn Treebank, our model without pre-trained representations achieves 93.8 F1, outperforming all existing methods with similar time complexity. With pre-training, our model pushes the F1 score to 95.7, which is on par with the SoTA while supporting faster decoding with a speed of over 1,100 sentences per second (fastest so far). Our model also performs competitively with SoTA methods on the multilingual parsing tasks in the SPMRL 2013/2014 shared tasks. In discourse parsing, our method establishes a new SoTA in end-to-end sentence-level parsing performance on the RST Discourse Treebank with an F1 score of 78.82.

2 Parsing as a Splitting Problem

Constituency parsing (both syntactic and discourse) can be considered as the problem of finding a set of labeled spans over the input text (Stern et al., 2017a). Let $\mathbb{S}(T)$ denote the set of *labeled spans* for a parse tree T , which can formally be expressed as (excluding the trivial singleton span layer):

$$\mathbb{S}(T) := \{((i_t, j_t), l_t)\}_{t=1}^{|\mathbb{S}(T)|} \text{ for } i_t < j_t \quad (1)$$

where l_t is the label of the text span (i_t, j_t) encompassing tokens from index i_t to index j_t .

Previous approaches to syntactic parsing (Stern et al., 2017a; Kitaev and Klein, 2018; Nguyen et al., 2020) train a neural model to score each possible span and then apply a greedy or dynamic programming algorithm to find the parse tree. In other words, these methods are span-based formulation.

In contrary, we formulate constituency parsing as the problem of finding the splitting points in a recursive, top-down manner. For each parent node in a tree that spans over (i, j) , our parsing model is trained to point to the *boundary* between the tokens at k and $k + 1$ positions to split the parent span into two child spans (i, k) and $(k + 1, j)$. This is done through the Pointing mechanism (Vinyals et al., 2015a), where each splitting decision is modeled as a multinomial distribution over the input elements, which in our case are the token boundaries.

The correspondence between token- and boundary-based representations of a tree is straightforward. After including the start (`<eos>`) and end (`<eos>`) tokens, the token-based span (i, j) is equivalent to the boundary-based span $(i - 1, j)$ and the boundary between i -th and $(i + 1)$ -th tokens is indexed as i . For example, the (boundary-based) span “enjoys playing tennis” in Figure 1 is defined

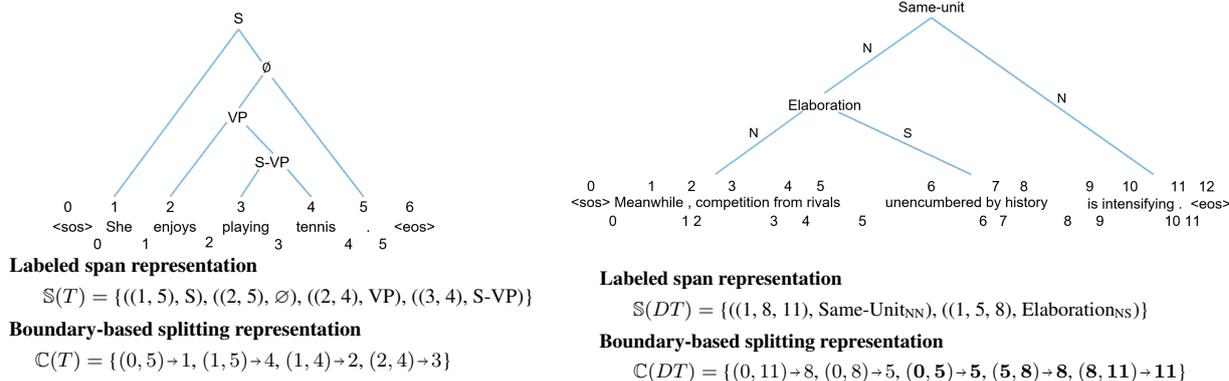


Figure 1: A syntactic tree at the left and a discourse tree (DT) at the right; both have a constituency structure. The internal nodes in the discourse tree (*Elaboration*, *Same-Unit*) represent coherence relations and the edge labels indicate the nuclearity statuses (‘N’ for Nucleus and ‘S’ for Satellite) of the child spans. Below the tree, we show the labeled span and splitting representations. The bold splits in the DT representation ($\mathbb{C}(DT)$) indicate the end of further splitting into smaller spans (*i.e.*, they are EDUs).

as (1, 4). Similarly, the boundary between the tokens “enjoys” and “playing” is indexed with 2.¹

Following the common practice in syntactic parsing, we binarize the n -ary tree by introducing a dummy label \emptyset . We also collapsed the nested labeled spans in the unary chains into unique atomic labels, such as S-VP in Figure 1. Every span represents an internal node in the tree, which has a left and a right child. Therefore, we can represent each internal node by its split into left and right children. Based on this, we define the set of splitting decisions $\mathbb{C}(T)$ for a syntactic tree T as follows.

Proposition 1 *A binary syntactic tree T of a sentence containing n tokens can be transformed into a set of splitting decisions $\mathbb{C}(T) = \{(i, j) \rightarrow k : i < k \leq j\}$ such that the parent span (i, j) is split into two child spans (i, k) and (k, j) .*

An example of the splitting representation of a tree is shown in Figure 1 (without the node labels). Note that our transformed representation has a one-to-one mapping with the tree since each splitting decision corresponds to one and only one internal node in the tree. We follow a depth-first order of the decision sequence, which in our preliminary experiments showed more consistent performance than other alternatives like breadth-first order.

Extension to End-to-End Discourse Parsing

Note that in syntactic parsing, the split position must be within the span but not at its edge, that is, k must satisfy $i < k < j$ for each boundary span (i, j) . Otherwise, it will not produce valid

¹We use the same example from (Stern et al., 2017a; Shen et al., 2018; Nguyen et al., 2020) to distinguish the differences between the methods.

sub-trees. In this case, we keep splitting until each span contains a single leaf token. However, for discourse trees, each leaf is an EDU – a clause-like unit that can contain one or multiple tokens.

Unlike previous studies which assume discourse segmentation as a pre-processing step, we propose a unified formulation that treats segmentation as one additional step in the top-down parsing process. To accommodate this, we relax Proposition 1 as:

Proposition 2 *A binary discourse tree DT of a text containing n tokens can be transformed into a set of splitting decisions $\mathbb{C}(DT) = \{(i, j) \rightarrow k : i < k \leq j\}$ such that the parent span (i, j) gets split into two child spans (i, k) and (k, j) for $k < j$ or a terminal span or EDU for $k = j$ (end of splitting the span further).*

We illustrate it with the DT example in Figure 1. Each splitting decision in $\mathbb{C}(DT)$ represents either the splitting of the parent span into two child spans (when the splitting point is strictly within the span) or the end of any further splitting (when the splitting point is the right endpoint of the span). By making this simple relaxation, our formulation can not only generate the discourse tree (in the former case) but can also find the discourse segments (EDUs) as a by-product (in the latter case).

3 Seq2Seq Parsing Framework

Let $\mathbb{C}(T)$ and $L(T)$ respectively denote the structure (in split representation) and labels of a tree T (syntactic or discourse) for a given text \mathbf{x} . We can express the probability of the tree as:

$$\begin{aligned}
 P_\theta(T|\mathbf{x}) &= P_\theta(L(T), \mathbb{C}(T)|\mathbf{x}) \\
 &= P_\theta(L(T)|\mathbb{C}(T), \mathbf{x})P_\theta(\mathbb{C}(T)|\mathbf{x})
 \end{aligned}
 \tag{2}$$

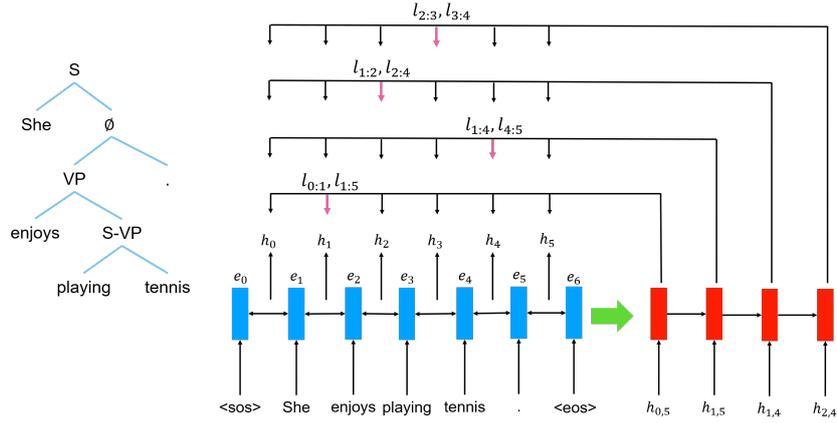


Figure 2: Our syntactic parser along with the decoding process for a given sentence. The input to the decoder at each step is the representation of the span to be split. We predict the splitting point using a biaffine function between the corresponding decoder state and the boundary-based encoder representations. A label classifier is used to assign labels to the left and right spans.

This factorization allows us to first infer the tree structure from the input text, and then find the corresponding labels. As discussed in the previous section, we consider the structure prediction as a sequence of splitting decisions to generate the tree in a top-down manner. Specifically, at each decoding step t , the output y_t represents the splitting decision $(i_t, j_t) \rightarrow k_t$. Thus, we can express the probability of the tree structure as follows:

$$\begin{aligned}
 P_\theta(\mathbb{C}(T)|\mathbf{x}) &= \prod_{y_t \in \mathbb{C}(T)} P_\theta(y_t | y_{<t}, \mathbf{x}) \\
 &= \prod_{t=1}^{|\mathbb{C}(T)|} P_\theta((i_t, j_t) \rightarrow k_t | ((i, j) \rightarrow k)_{<t}, \mathbf{x})
 \end{aligned} \tag{3}$$

This can effectively be modeled within a Seq2Seq pointing framework as shown in Figure 2. At each step t , the decoder autoregressively predicts the split point k_t in the input by conditioning on the current input span (i_t, j_t) and previous splitting decisions $((i, j) \rightarrow k)_{<t}$. This conditional splitting formulation (decision at step t depends on previous steps) can help our model to find better trees compared to *non-conditional* top-down parsers (Stern et al., 2017a; Shen et al., 2018; Nguyen et al., 2020), thus bridging the gap between the global (but expensive) and the local (but efficient) models. The labels $L(T)$ can be modeled by using a label classifier, as described later in the next section.

3.1 Model Architecture

We now describe the components of our parsing model: the sentence encoder, the span representation, the pointing model and the labeling model.

Sentence Encoder Given an input sequence of n tokens $\mathbf{x} = (x_1, \dots, x_n)$, we first add $\langle \text{sos} \rangle$ and

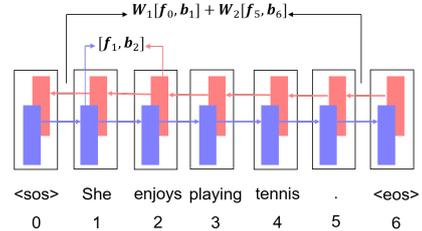


Figure 3: Illustration of our boundary-based span encoder. Here we have shown the representation for the boundary at 1 and the representation of the boundary-based span $(0, 5)$ that corresponds to the sentence “*She enjoys playing tennis .*”.

$\langle \text{eos} \rangle$ markers to the sequence. After that, each token t in the sequence is mapped into its dense vector representation e_t as

$$e_t = [e_t^{\text{char}}, e_t^{\text{word}}] \tag{4}$$

where e_t^{char} , e_t^{word} are respectively the character and word embeddings of token t . Similar to (Kiatav and Klein, 2018; Nguyen et al., 2020), we use a character LSTM to compute the character embedding of a token. We experiment with both randomly initialized and pretrained token embeddings. When pretrained embedding is used, the character embedding is replaced by the pretrained token embedding. The token representations are then passed to a 3-layer Bi-LSTM encoder to obtain their contextual representations. In the experiments, we find that even without the POS-tags, our model performs competitively with other baselines that use them.

Boundary and Span Representations To represent each boundary between positions k and $k + 1$, we use the fencepost representation (Cross and Huang, 2016; Stern et al., 2017a):

$$h_k = [f_k, b_{k+1}] \tag{5}$$

where \mathbf{f}_k and \mathbf{b}_{k+1} are the forward and backward LSTM hidden vectors at positions k and $k + 1$, respectively. To represent the span (i, j) , we compute a linear combination of the two endpoints

$$\mathbf{h}_{i,j} = \mathbf{W}_1 \mathbf{h}_i + \mathbf{W}_2 \mathbf{h}_j \quad (6)$$

This span representation will be used as input to the decoder. Figure 3 shows the boundary-based span representations for our example.

The Decoder Our model uses a unidirectional LSTM as the decoder. At each decoding step t , the decoder takes as input the corresponding span (i, j) (specifically, $\mathbf{h}_{i,j}$) and its previous state \mathbf{d}_{t-1} to generate the current state \mathbf{d}_t and then apply a biaffine function (Dozat and Manning, 2017) between \mathbf{d}_t and **all** of the encoded boundary representations $(\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_n)$ as follows:

$$\mathbf{d}'_t = \text{MLP}_d(\mathbf{d}_t) \quad \mathbf{h}'_i = \text{MLP}_h(\mathbf{h}_i) \quad (7)$$

$$s_{t,i} = \mathbf{d}'_t{}^T \mathbf{W}_{dh} \mathbf{h}'_i + \mathbf{h}'_i{}^T \mathbf{w}_h \quad (8)$$

$$a_t = \frac{\exp(s_{t,i})}{\sum_{i=1}^n \exp(s_{t,i})} \quad (9)$$

where each MLP operation includes a linear transformation with LeakyReLU activation to transform \mathbf{d} and \mathbf{h} into equal-sized vectors, and $\mathbf{W}_{dh} \in \mathbb{R}^{d \times d}$ and $\mathbf{w}_h \in \mathbb{R}^d$ are respectively the weight matrix and weight vector for the biaffine function. The biaffine scores are then passed through a softmax layer to acquire the pointing distribution $\mathbf{a}_t \in [0, 1]^n$ for the splitting decision.

When decoding the tree during inference, at each step we only examine the ‘valid’ splitting points between i and j – for syntactic parsing, it is $i < k < j$ and for discourse parsing, it is $i < k \leq j$.

Label Classifier For syntactic parsing, we perform the label assignments for a span (i, j) as:

$$\mathbf{h}_i^l = \text{MLP}_l(\mathbf{h}_i); \quad \mathbf{h}_j^r = \text{MLP}_r(\mathbf{h}_j) \quad (10)$$

$$P_\theta(l|i, j) = \text{softmax}((\mathbf{h}_i^l)^T \mathbf{W}_{lr} \mathbf{h}_j^r + (\mathbf{h}_i^l)^T \mathbf{W}_l + (\mathbf{h}_j^r)^T \mathbf{W}_r + \mathbf{b}) \quad (11)$$

$$l_{i,j} = \arg \max_{l \in L} P_\theta(l|i, j) \quad (12)$$

where each of MLP_l and MLP_r includes a linear transformation with LeakyReLU activations to transform the left and right spans into equal-sized vectors, and $\mathbf{W}_{lr} \in \mathbb{R}^{d \times L \times d}$, $\mathbf{W}_l \in \mathbb{R}^{d \times L}$, $\mathbf{W}_r \in \mathbb{R}^{d \times L}$ are the weights and \mathbf{b} is a bias vector with L being the number of phrasal labels.

For discourse parsing, we perform label assignment *after* every split decision since the label here represents the relation between the child spans.

Specifically, as we split a span (i, j) into two child spans (i, k) and (k, j) , we determine the relation label as the following.

$$\mathbf{h}_{ik}^l = \text{MLP}_l([\mathbf{h}_i, \mathbf{h}_k]); \quad \mathbf{h}_{kj}^r = \text{MLP}_r([\mathbf{h}_k, \mathbf{h}_j]) \quad (13)$$

$$P_\theta(l|(i, k), (k, j)) = \text{softmax}((\mathbf{h}_{ik}^l)^T \mathbf{W}_{lr} \mathbf{h}_{kj}^r + (\mathbf{h}_{ik}^l)^T \mathbf{W}_l + (\mathbf{h}_{kj}^r)^T \mathbf{W}_r + \mathbf{b}) \quad (14)$$

$$l_{(i,k),(k,j)} = \arg \max_{l \in L} P_\theta(l|(i, k), (k, j)) \quad (15)$$

where MLP_l , MLP_r , \mathbf{W}_{lr} , \mathbf{W}_l , \mathbf{W}_r , \mathbf{b} are similarly defined.

Training Objective The total loss is simply the sum of the cross entropy losses for predicting the structure (split decisions) and the labels:

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{L}_{\text{split}}(\theta_e, \theta_d) + \mathcal{L}_{\text{label}}(\theta_e, \theta_{\text{label}}) \quad (16)$$

where $\theta = \{\theta_e, \theta_d, \theta_{\text{label}}\}$ denotes the overall model parameters, which includes the encoder parameters θ_e shared by all components, parameters for splitting θ_d and parameters for labeling θ_{label} .

3.2 Top-Down Beam-Search Inference

As mentioned, existing top-down syntactic parsers do not consider the decoding history. They also perform greedy inference. With our conditional splitting formulation, our method can not only model the splitting history but also enhance the search space of high scoring trees through beam search.

At each step, our decoder points to *all* the encoded *boundary* representations which ensures that the pointing scores are in the same scale, allowing a fair comparison between the total scores of all candidate subtrees. With these uniform scores, we could apply a beam search to infer the most probable tree using our model. Specifically, the method generates the tree in depth-first order while maintaining top- B (beam size) partial trees at each step. It terminates exactly after $n - 1$ steps, which matches the number of internal nodes in the tree. Each decoding step can be parallelized ($\mathcal{O}(1)$ complexity) using GPUs. This makes our algorithm run at $\mathcal{O}(n)$ time complexity, which is faster than most top-down methods. We leave further details about the inference algorithm to the Appendix.

By enabling beam search, our method can find the best tree by comparing high scoring trees within a reasonable search space, making our model competitive with existing structured (globally) inference methods that use more expensive algorithms like CKY and/or larger models (Kitaev and Klein, 2018; Zhang et al., 2020b).

4 Experiment

Datasets and Metrics To show the effectiveness of our approach, we conduct experiments on both syntactic and sentence-level RST parsing tasks.² We use the standard Wall Street Journal (WSJ) part of the Penn Treebank (PTB) (Marcus et al., 1993) for syntactic parsing and RST Discourse Treebank (RST-DT) (Lynn et al., 2002) for discourse parsing. For syntactic parsing, we also experiment with the multilingual parsing tasks on seven different languages from the SPMRL 2013-2014 shared task (Seddah et al., 2013): Basque, French, German, Hungarian, Korean, Polish and Swedish.

For evaluation on syntactic parsing, we report the standard labeled precision (LP), labeled recall (LR), and labelled F1 computed by `evalb`³. For evaluation on RST-DT, we report the standard span, nuclearity label, relation label F1 scores, computed using the implementation of (Lin et al., 2019).⁴

4.1 English (PTB) Syntactic Parsing

Setup We follow the standard train/valid/test split, which uses Sections 2-21 for training, Section 22 for development and Section 23 for evaluation. This results in 39,832 sentences for training, 1,700 for development, and 2,416 for testing. For our model, we use an LSTM encoder-decoder framework with a 3-layer bidirectional encoder and 3-layer unidirectional decoder. The word embedding size is 100 while the character embedding size is 50; the LSTM hidden size is 400. The hidden dimension in MLP modules and biaffine function for split point prediction is 500. The beam width B is set to 20. We use the Adam optimizer (Kingma and Ba, 2015) with a batch size of 5000 tokens, and an initial learning rate of 0.002 which decays at the rate 0.75 exponentially at every 5k steps. Model selection for final evaluation is performed based on the labeled F1 score on the development set.

Results without Pre-training From the results shown in Table 1, we see that our model achieves an F1 of 93.77, the highest among models that use *top-down* methods. Specifically, our parser outperforms Stern et al. (2017a); Shen et al. (2018) by about 2 points in F1-score and Nguyen et al. (2020) by ~ 1 point. Notably, without beam search (beam

²Extending the discourse parser to the document level may require handling of intra- and multi-sentential constituents differently, which we leave for future work.

³<http://nlp.cs.nyu.edu/evalb/>

⁴https://github.com/ntunlp/sg/UnifiedParser_RST

Model	LR	LP	F1
Top-Down Inference			
Stern et al. (2017a)	93.20	90.30	91.80
Shen et al. (2018)	92.00	91.70	91.80
Nguyen et al. (2020)	92.91	92.75	92.78
Our Model	93.90	93.63	93.77
CKY/Chart Inference			
Gaddy et al. (2018)	91.76	92.41	92.08
Kitaev and Klein (2018)	93.20	93.90	93.55
Wei et al. (2020)	93.3	94.1	93.7
Zhang et al. (2020b)	93.84	93.58	93.71
Other Approaches			
Gómez and Vilares (2018)	-	-	90.7
Liu and Zhang (2017)	-	-	91.8
Stern et al. (2017b)	92.57	92.56	92.56
Zhou and Zhao (2019)	93.64	93.92	93.78

Table 1: Results for single models (no pre-training) on the PTB WSJ test set, Section 23.

Model	F1
Nguyen et al. (2020)	95.5
Our model	95.7
Kitaev et al. (2019)	95.6
Zhang et al. (2020b)	95.7
Wei et al. (2020)	95.8
Zhou and Zhao (2019)	95.8

Table 2: Results on PTB WSJ test set with pretraining.

width 1 or greedy decoding), our model achieves an F1 of 93.40, which is still better than other top-down methods. Our model also performs competitively with CKY-based methods like (Kitaev and Klein, 2018; Zhang et al., 2020b; Wei et al., 2020; Zhou and Zhao, 2019), while these methods run slower than ours. Besides, Zhou and Zhao (2019) also use external supervision from the dependency parsing task which our method does not use.

Results with Pre-training Similar to (Kitaev and Klein, 2018; Kitaev et al., 2019), we also evaluate our parser with BERT embeddings (Devlin et al., 2019). They fine-tuned *Bert-large-cased* on the task, while in our work keeping it frozen was already good enough (gives training efficiency). As shown in Table 2, our model achieves an F1 of 95.7, which is on par with SoTA models. However, our parser runs faster than other methods. Specifically, our model runs at $\mathcal{O}(n)$ time complexity, while CKY needs $\mathcal{O}(n^3)$. Comprehensive comparisons on parsing speed are presented later.

Model	Basque	French	German	Hungarian	Korean	Polish	Swedish
(Bjorkelund et al., 2014)	88.24	82.53	81.66	91.72	83.81	90.50	85.50
(Coavoux and Crabbé, 2017) ⁺	88.81	82.49	85.34	92.34	86.04	93.64	84.0
(Kitaev and Klein, 2018)	89.71	84.06	87.69	92.69	86.59	93.69	84.45
(Nguyen et al., 2020) ⁺	90.23	82.20	84.91	91.07	85.36	93.99	86.87
Our Model	89.74	84.12	85.21	92.84	86.72	92.10	85.81

Table 3: Results on SPMRL test sets without pre-training. The sign ⁺ denotes that systems use predicted POS tags.

4.2 SPMRL Multilingual Syntactic Parsing

We use the identical hyper-parameters and optimizer setups as in English PTB. We follow the standard train/valid/test split provided in the SPMRL datasets; details are reported in the Appendix.

From the results in Table 3, we see that our model achieves the highest F1 in French, Hungarian and Korean and higher than the best baseline by 0.06, 0.15 and 0.13, respectively. Our method also rivals existing SoTA methods on other languages even though some of them use predicted POS tags.

4.3 Discourse Parsing

Setup For discourse parsing, we follow the standard split from (Lin et al., 2019), which has 7321 sentence-level discourse trees for training and 951 for testing. We also randomly select 10% of the training for validation. Model selection for testing is performed based on the F1 of relation labels on the validation set. We use the same model settings as the constituency parsing experiments, with BERT as pretrained embeddings.⁵

Results Table 4 compares the results on the discourse parsing tasks in two settings: (i) when the EDUs are given (gold segmentation) and (ii) end-to-end parsing. We see that our model outperforms the baselines in both parsing conditions achieving SoTA. When gold segmentation is provided, our model outperforms the single-task training model of (Lin et al., 2019) by 0.43%, 1.06% and 0.82% absolute in Span, Nuclearity and Relation, respectively. Our parser also surpasses their joint training model, which uses multi-task training (segmentation and parsing), with 0.61% and 0.4% absolute improvements in Nuclearity and Relation, respectively. For end-to-end parsing, compared to the best baseline (Lin et al., 2019), our model yields 0.27%, 0.67%, and 1.30% absolute improvements in Span, Nuclearity, Relation, respectively. This demonstrates the effectiveness of our conditional

⁵Lin et al. (2019) used ELMo (Peters et al., 2018) as pretrained embeddings. With BERT, their model performs worse which we have confirmed with the authors.

Approach	Span	Nuclearity	Relation
Parsing with gold EDU segmentation			
Human Agreement	95.7	90.4	83.0
Baselines			
Wang et al. (2017)	95.6	87.8	77.6
Lin et al. (2019) (single)	96.94	90.89	81.28
Lin et al. (2019) (joint)	97.44	91.34	81.70
Our Model	97.37	91.95	82.10
End-to-End parsing			
Baselines			
Soricut and Marcu (2003)	76.7	70.2	58.0
Joty et al. (2012)	82.4	76.6	67.5
Lin et al. (2019) (pipeline)	91.14	85.80	76.94
Lin et al. (2019) (joint)	91.75	86.38	77.52
Our Model	92.02	87.05	78.82

Table 4: Results on discourse parsing tasks on the RST-DT test set with and without gold segmentation.

splitting approach and end-to-end formulation of the discourse analysis task. The fact that our model improves on *span* identification indicates that our method also yields better EDU segmentation.

4.4 Parsing Speed Comparison

We compare parsing speed of different models in Table 5. We ran our models on both CPU (Intel Xeon W-2133) and GPU (Nvidia GTX 1080 Ti).

Syntactic Parsing The Berkeley Parser and ZPar are two representative non-neural parsers without access to GPUs. Stern et al. (2017a) employ max-margin training and perform top-down greedy decoding on CPUs. Meanwhile, Kitaev and Klein (2018); Zhou and Zhao (2019); Wei et al. (2020) use a self-attention encoder and perform decoding using Cython for acceleration. Zhang et al. (2020b) perform CKY decoding on GPU. The parser proposed by Gómez and Vilares (2018) is also efficient as it treats parsing as a sequence labeling task. However, its parsing accuracy is much lower compared to others (90.7 F1 in Table 1).

We see that our parser is much more efficient than existing ones. It utilizes neural modules to perform splitting, which is optimized and parallelized with efficient GPU implementation. It can parse 1, 127 sentences/second, which is faster than

System	Speed (Sents/s)	Speedup
Syntactic Parser		
Petrov and Klein (2007) (Berkeley)	6	1.0x
Zhu et al. (2013)(ZPar)	90	15.0x
Stern et al. (2017a)	76	12.7x
Shen et al. (2018)	111	18.5x
Nguyen et al. (2020)	130	21.7x
Zhou and Zhao (2019)	159	26.5x
Wei et al. (2020)	220	36.7x
Kitaev and Klein (2018)	332	55.3x
Gómez and Vilares (2018)	780	130x
Zhang et al. (2020b)	924	154x
Our model (GPU)	1127	187.3x
End-to-End Discourse parsing (Segmenter + Parser)		
CODRA (Joty et al., 2015)	3.05	1.0x
SPADE (Soricut and Marcu, 2003)	4.90	1.6x
(Lin et al., 2019)	28.96	9.5x
Our end-to-end parser (CPU)	59.03	19.4x
Our end-to-end parser (GPU)	135.85	44.5x

Table 5: Speed comparison of our parser with existing syntactic and discourse parsers.

existing parsers. In fact, we could still improve our speed by choosing faster architecture like Transformer and allowing parallel same tree level generation at each decoding step. We leave these extensions to future work.

Discourse Parsing For measuring discourse parsing speed, we follow the same set up as Lin et al. (2019). We include the model loading time for all systems. Because of using a number of features, SPADE and CODRA are typically slower than the neural models which use pretrained embeddings. As shown, our parser is 4.7x faster than the fastest end-to-end parser of Lin et al. (2019), making it not only effective but also highly efficient. Running on CPU or GPU, our model is faster than all the other models, thanks to the end-to-end formulation that does not need EDU segmentation beforehand.

5 Related Work

With the recent popularity of neural architectures, such as LSTMs (Hochreiter and Schmidhuber, 1997) and Transformers (Vaswani et al., 2017), various neural models have been proposed to encode the input sentences and infer their constituency trees. To enforce structural consistency, such methods employ either a greedy transition-based (Dyer et al., 2016; Liu and Zhang, 2017), a globally optimized chart parsing (Gaddy et al., 2018; Kitaev and Klein, 2018), or a greedy top-down algorithm (Stern et al., 2017a; Shen et al., 2018). Meanwhile, researchers also tried to cast the parsing problem into tasks that can be solved differently. For example, Gómez and Vilares (2018); Shen et al. (2018) proposed to map the syntactic tree of a sentence

containing n tokens into a sequence of $n - 1$ labels or scalars. However, parsers of this type suffer from the exposure bias during inference. Beside these methods, Seq2Seq models have been used to generate a linearized form of the tree (Vinyals et al., 2015b; Kamigaito et al., 2017; Suzuki et al., 2018; Fernández-González and Gómez-Rodríguez, 2020). However, these methods may generate invalid trees when the open and end brackets do not match.

In discourse parsing, existing parsers receive the EDUs from a segmenter to build the discourse tree, which makes them susceptible to errors when the segmenter produces incorrect EDUs (Joty et al., 2015; Lin et al., 2019; Zhang et al., 2020a).

Our approach differs from previous methods in that it represents the constituency structure as a series of splitting representations, and uses a Seq2Seq framework to model the splitting decision at each step. By enabling beam search, our model can find the best trees without the need to perform an expensive global search. We also unify discourse segmentation and parsing into one system by generalizing our model, which has been done for the first time to the best of our knowledge.

Our splitting mechanism shares some similarities with Pointer Network (Vinyals et al., 2015a; Ma et al., 2018; Fernández-González and Gómez-Rodríguez, 2019), but is distinct from it that in each decoding step, our method identifies the splitting point of a span and generates a new input for future steps instead of pointing to generate the next decoder input.

6 Conclusion

We have presented a novel, generic parsing method for constituency parsing based on a Seq2Seq framework. Our method supports an efficient top-down decoding algorithm that uses a pointing function for scoring possible splitting points. The pointing mechanism captures global structural properties of a tree and allows efficient training with a cross entropy loss. Our formulation, when applied to discourse parsing, can bypass discourse segmentation as a pre-requisite step. Through experiments we have shown that our method outperforms all existing top-down methods on English Penn Treebank and RST Discourse Treebank sentence-level parsing tasks. With pre-trained representations, our method rivals state-of-the-art methods, while being faster. Our model also establishes a new state-of-the-art for sentence-level RST parsing.

References

- Yoav Artzi, Nicholas FitzGerald, and Luke Zettlemoyer. 2013. Semantic parsing with Combinatory Categorical Grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Tutorials)*, page 2, Sofia, Bulgaria. Association for Computational Linguistics.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.
- Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. 2015. Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, volume 28, pages 1171–1179. Curran Associates, Inc.
- Anders Bjorkelund, Ozlem Cetinoglu, Agnieszka Falenska, Richard Farkas, Thomas Mueller, Wolfgang Seeker, and Zsolt Szanto. 2014. The ims-wrocław-szeged-cis entry at the spmrl 2014 shared task: Reranking and morphosyntax meet unlabeled data. In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of NonCanonical Languages*, pages 97–102.
- Maximin Coavoux and Benoît Crabbé. 2017. Multilingual lexicalized constituency parsing with word-level auxiliary tasks. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 331–336, Valencia, Spain. 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. 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.
- Timothy Dozat and Christopher D. Manning. 2017. Deep biaffine attention for neural dependency parsing. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Greg Durrett and Dan Klein. 2015. Neural CRF 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 302–312, Beijing, China. 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.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2019. Left-to-right dependency parsing with pointer networks. 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 710–716, Minneapolis, Minnesota. Association for Computational Linguistics.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2020. Enriched in-order linearization for faster sequence-to-sequence constituent parsing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4092–4099, Online. Association for Computational Linguistics.
- 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, Volume 1 (Long Papers)*, pages 999–1010, New Orleans, Louisiana. Association for Computational Linguistics.
- Carlos Gómez, Rodríguez and David Vilares. 2018. Constituent parsing as sequence labeling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1314–1324, Brussels, Belgium. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Shafiq Joty, Giuseppe Carenini, and Raymond Ng. 2012. A novel discriminative framework for sentence-level discourse analysis. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 904–915, Jeju Island, Korea. Association for Computational Linguistics.
- Shafiq Joty, Giuseppe Carenini, and Raymond T. Ng. 2015. CODRA: A novel discriminative framework for rhetorical analysis. *Computational Linguistics*, 41(3):385–435.

- 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. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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.
- 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. Association for Computational Linguistics.
- Xiang Lin, Shafiq Joty, Prathyusha Jwalapuram, and M Saiful Bari. 2019. [A unified linear-time framework for sentence-level discourse parsing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4200, Florence, Italy. Association for Computational Linguistics.
- Jiangming Liu and Yue Zhang. 2017. Shift-reduce constituent parsing with neural lookahead features. *Transactions of the Association for Computational Linguistics*, 5:45–58.
- Carlson Lynn, Daniel Marcu, and Mary Ellen Okunowski. 2002. Rst discourse treebank (rst-dt) ldc2002t07. *Linguistic Data Consortium*.
- Xuezhe Ma, Zecong Hu, Jingzhou Liu, Nanyun Peng, Graham Neubig, and Eduard Hovy. 2018. [Stack-pointer networks for dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1403–1414, Melbourne, Australia. Association for Computational Linguistics.
- William Mann and Sandra Thompson. 1988. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8(3):243–281.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.*, 19(2):313–330.
- Thanh-Tung Nguyen, Xuan-Phi Nguyen, Shafiq Joty, and Xiaoli Li. 2020. Efficient constituency parsing by pointing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3284–3294, Online. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proc. of NAACL*.
- Slav Petrov and Dan Klein. 2007. Improved inference for unlexicalized parsing. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 404–411, Rochester, New York. Association for Computational Linguistics.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho D. Choi, Richárd Farkas, Jennifer Foster, Iakes Goenaga, Koldo Gojenola Galletebeitia, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiórkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villemonte de la Clergerie. 2013. Overview of the SPMRL 2013 shared task: A cross-framework evaluation of parsing morphologically rich languages. In *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 146–182, Seattle, Washington, USA. Association for Computational Linguistics.
- Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordani, Aaron Courville, and Yoshua Bengio. 2018. Straight to the tree: Constituency parsing with neural syntactic distance. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1171–1180, Melbourne, Australia. Association for Computational Linguistics.
- Radu Soricut and Daniel Marcu. 2003. Sentence level discourse parsing using syntactic and lexical information. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 228–235.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017a. 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.
- Mitchell Stern, Daniel Fried, and Dan Klein. 2017b. [Effective inference for generative neural parsing](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1695–1700, Copenhagen, Denmark. Association for Computational Linguistics.
- 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.
- 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015a. Pointer networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 2692–2700. Curran Associates, Inc.
- Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015b. [Grammar as a foreign language](#). In *Advances in Neural Information Processing Systems*, volume 28, pages 2773–2781. Curran Associates, Inc.
- Yizhong Wang, Sujian Li, and Houfeng Wang. 2017. A two-stage parsing method for text-level discourse analysis. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 184–188. Association for Computational Linguistics.
- B. Webber. 2004. D-LTAG: Extending Lexicalized TAG to Discourse. *Cognitive Science*, 28(5):751–779.
- Yang Wei, Yuanbin Wu, and Man Lan. 2020. A span-based linearization for constituent trees. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3267–3277, Online. Association for Computational Linguistics.
- Longyin Zhang, Yuqing Xing, Fang Kong, Peifeng Li, and Guodong Zhou. 2020a. [A top-down neural architecture towards text-level parsing of discourse rhetorical structure](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6386–6395, Online. Association for Computational Linguistics.
- Yu Zhang, Houquan Zhou, and Zhenghua Li. 2020b. [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.
- Junru Zhou and Hai Zhao. 2019. Head-driven phrase structure grammar parsing on penn treebank. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2396–2408.
- 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 (Volume 1: Long Papers)*, pages 434–443, Sofia, Bulgaria. Association for Computational Linguistics.