
Strongly Incremental Constituency Parsing with Graph Neural Networks

Kaiyu Yang

Princeton University
kaiyuy@cs.princeton.edu

Jia Deng

Princeton University
jiadeng@cs.princeton.edu

Abstract

Parsing sentences into syntax trees can benefit downstream applications in NLP. Transition-based parsers build trees by executing actions in a state transition system. They are computationally efficient, and can leverage machine learning to predict actions based on partial trees. However, existing transition-based parsers are predominantly based on the shift-reduce transition system, which does not align with how humans are known to parse sentences. Psycholinguistic research suggests that human parsing is *strongly incremental*—humans grow a single parse tree by adding exactly one token at each step. In this paper, we propose a novel transition system called *attach-juxtapose*. It is strongly incremental; it represents a partial sentence using a single tree; each action adds exactly one token into the partial tree. Based on our transition system, we develop a strongly incremental parser. At each step, it encodes the partial tree using a graph neural network and predicts an action. We evaluate our parser on Penn Treebank (PTB) and Chinese Treebank (CTB). On PTB, it outperforms existing parsers trained with only constituency trees; and it performs on par with state-of-the-art parsers that use dependency trees as additional training data. On CTB, our parser establishes a new state of the art. Code is available at <https://github.com/princeton-vl/attach-juxtapose-parser>.

1 Introduction

Constituency parsing is a core task in natural language processing. It recovers the syntactic structures of sentences as trees (Fig. 1). State-of-the-art parsers are based on deep neural networks and typically consist of an encoder and a decoder. The encoder embeds input tokens into vectors, from which the decoder generates a parse tree. A main class of decoders builds trees by executing a sequence of actions in a state transition system [33, 12, 23]. These transition-based parsers achieve linear runtime in sentence length. More importantly, they construct partial trees during decoding, enabling the parser to leverage explicit structural information for predicting the next action.

Most existing transition-based parsers adopt the shift-reduce transition system [33, 12, 23]. They represent the partial sentence as a stack of subtrees. At each step, the parser either pushes a new token onto the stack (*shift*) or combines two existing subtrees in the stack (*reduce*).

Despite their empirical success, shift-reduce parsers appear to differ from how humans are known to perform parsing. Psycholinguistic research [26, 38, 36] has suggested that human parsing is *strongly incremental*: at each step, humans process exactly one token—no more, no less—and integrate it into a single parse tree for the partial sentence. In a shift-reduce system, however, only *shift* actions process new tokens, and the partial sentence is represented as a stack of disconnected subtrees rather than a single connected tree.




Figure 1: The constituency tree for “Arthur is King of the Britons.” Leaves are labeled with tokens, and internal nodes are labeled with syntactic categories, e.g., *S* for *sentence*, *NP* for *noun phrase*, *VP* for *verb phrase*, and *PP* for *prepositional phrase*.

This observation puts forward an intriguing question: can a strongly incremental transition system lead to a better parser? Intuitively, a strongly incremental system is more aligned with human processing; as a result, the sequence of actions may be easier to learn.

Attach-juxtapose transition system We propose a novel transition system named *attach-juxtapose*, which enables strongly incremental constituency parsing. For a sentence of length n , we start with an empty tree and execute n actions; each action integrates exactly one token into the current partial tree, deciding on *where* and *how* to integrate the new token. There are two types of actions: *attach*, which attaches the new token as a child to an existing node, and *juxtapose*, which juxtaposes the new token as a sibling to an existing node while also creating a shared parent node (Fig. 2). We can prove that any parse tree without unary chains can be constructed by a unique sequence of actions in this attach-juxtapose system.

Being strongly incremental, our system represents the state as a single tree rather than a stack of multiple subtrees. Not only is the single-tree representation more aligned with humans, but it also allows us to tap into a large inventory of model architectures for learning from graph data, such as TreeLSTMs [40] and graph neural networks (GNNs) [19]. Further, the single-tree representation provides valid syntax trees for partial sentences, which is impossible in bottom-up shift-reduce systems [33]. Taking “Arthur is King of the Britons” as an example, we can produce a valid tree for the prefix “Arthur is King” (Fig. 2 *Bottom*). Whereas in bottom-up shift-reduce systems, you must complete the subtree for “is King of the Britons” before connecting it to “Arthur.” Therefore, our representation captures the complete syntactic structure of the partial sentence, and thus provides stronger guidance for action generation.

Our transition system can be understood as a refactorization of In-order Shift-reduce System (ISR) proposed by Liu and Zhang [23]. We prove that a sequence of actions in our system can be translated into a sequence of actions in ISR, but our sequence is shorter (Theorem 4). Specifically, to generate a parse tree with n leaves and m internal nodes (assuming no unary chains), our sequence has length n , whereas ISR has length $n + 2m$. On the other hand, each of our actions has a larger number of choices, resulting in a different trade-off between the sequence length and the number of choices per action. We hypothesize that our system achieves a trade-off more amenable to machine learning due to closer alignment with human processing.

Action generation with graph neural networks Based on the attach-juxtapose system, we develop a strongly incremental parser by training a deep neural network to generate actions. Specifically, we adopt the encoder in prior work [21, 50] and propose a novel graph-based decoder. It uses GNNs to learn node features in the partial tree, and uses attention to predict where and how to integrate the new token. To our knowledge, this is the first time GNNs are applied to constituency parsing.

We evaluate our method on two standard benchmarks for constituency parsing: Penn Treebank (PTB) [25] and Chinese Treebank (CTB) [47]. On PTB, our method outperforms existing parsers trained with only constituency trees. And it performs competitively with state-of-the-art parsers that use dependency trees as additional training data. On CTB, we achieve an F1 score of 93.59—a significant improvement of 0.95 upon previous best results. These results demonstrate the effectiveness of our strongly incremental parser.

Contributions Our contributions are threefold. First, we propose attach-juxtapose, a novel transition system for constituency parsing. It is strongly incremental and motivated by psycholinguistics. Second, we provide theoretical results characterizing its capability and its connections with an existing shift-reduce system [23]. Third, we develop a parser by generating actions in the attach-juxtapose system. Our parser achieves state-of-the-art performance on two standard benchmarks.

2 Related Work

Constituency parsing Significant progress in constituency parsing has been made by powerful token representations. Stern et al. [37] fed tokens into an LSTM [17] to obtain contextualized embeddings. Gaddy et al. [14] demonstrated the value of character-level features. Kitaev and Klein [21] replaced LSTMs with self-attention layers [42]. They also show that pre-trained contextualized embeddings such as ELMo [31] significantly improve parsing performance. Further improvements [20, 50] came with more powerful pre-trained embeddings, including BERT [10] and XLNet [48]. Mrini et al. [28] proposed label attention layers that improved upon self-attention layers. All these works use an existing decoder (chart-based) and focus on designing a new encoder. In contrast, we use an existing encoder by Kitaev and Klein [21] and focus on designing a new decoder.

There are several types of decoders in prior work: chart-based decoders search for a tree maximizing the sum of span scores via dynamic programming (e.g., the CKY algorithm) [37, 14, 21, 50, 28]; transition-based decoders build trees through a sequence of actions [33, 51, 12, 8, 23]; sequence-based decoders generate a linearized sequence of the tree using seq2seq models [43, 5, 39, 15, 24]. Our method is transition-based; but unlike the conventional shift-reduce methods, we propose a novel transition system, which is strongly incremental.

State-of-the-art parsers perform joint constituency parsing and dependency parsing, e.g., through head-driven phrase structure grammar (HPSG) [50, 28]. They use dependency trees as additional training data, which are converted from constituency trees by a set of hand-crafted rules [9]. In contrast, our parser is trained with only constituency trees and achieves competitive performance.

Transition-based constituency parsers Most transition-based parsers adopt the shift-reduce transition system: a state consists of a buffer B holding unprocessed tokens and a stack S holding processed subtrees. Initially, S is empty, and B contains the entire input sentence. In a successful final state, B is empty, and S contains a single subtree—the complete parse tree. Below are actions for a standard bottom-up shift-reduce system such as Sagae and Lavie [33], which corresponds to post-order traversal of the complete parse tree.

- **shift**: Remove the first element from B and push it onto S .
- **unary_reduce(X)**: Pop a subtree from S ; add a label X as its parent; and push it back onto S .
- **binary_reduce(X)**: Pop two subtrees; add a label X as their shared parent; and push back.

Dyer et al. [12] proposed a top-down (pre-order) variant. Liu and Zhang [23] proposed an in-order shift-reduce system, outperforming bottom-up [33] and top-down [12] baselines. Compared to our transition system, none of these shift-reduce systems is strongly incremental, because they represent the partial sentence as a stack of disconnected subtrees and they do not process exactly one token per action—only the **shift** action consumes a token.

Among the shift-reduce systems, our approach is most related to the in-order system by Liu and Zhang [23] in that each action in our system can be mapped to a combination of actions in their system (see Sec. 3 for details). An analogy is that their actions resemble a set of microinstructions for CPUs, where each instruction is simple but it takes many instructions to complete a task; our actions resemble a set of complex instructions, where each instruction is more complex but it takes fewer instructions to complete the same task.

Transition-based parsers use machine learning to make local decisions—determining the action to take at each step. This poses the question of how to represent a stack of subtrees in shift-reduce systems. Earlier works such as Sagae and Lavie [33] and Zhu et al. [51] use hand-crafted features. More recent works [44, 12, 8, 23] have switched to recurrent neural networks and LSTMs. In particular, Dyer et al. [11] propose an LSTM-based model named Stack LSTM for representing the stack. It is designed for dependency parsing but applies to constituency parsing as well [11, 23]. However, we

do not have to represent stacks thanks to the single-tree state representation. Instead, we use graph neural networks (GNNs) [19] to represent partial trees.

Incremental parsing Prior work has built parsers inspired by the incremental syntax processing of humans. Earlier works focused on psycholinguistic modeling of humans and evaluated on a handful of carefully curated sentences [27, 1]. More recent methods switched to developing efficient parsers with a wide coverage of real-world texts. Roark [32] proposed a top-down incremental parser that expands nodes in the partial tree using a probabilistic context-free grammar (PCFG). In contrast, our system is more flexible by not restricted to any predefined grammar. Also, we predict actions leveraging not only top-down information but also bottom-up information.

Costa et al. [7] proposed a transition system for incremental parsing. Similar to ours, it integrates exactly one token per step into the partial tree. However, at each step, they have to predict an unbounded number of labels, whereas we have to predict no more than two. Therefore, our action space is smaller than theirs and thus easier to navigate by learning-based parsers. In fact, this limitation may have prevented Costa et al. [7] from building a fully functional parser, and they only evaluated on action generation. Collins and Roark [6] developed a parser based on Costa et al. [7] by using grammar rules and heuristics to prune the action space. In contrast, our action space is more flexible without grammar rules but still tractable for machine learning models.

Graph neural networks for syntactic processing GNNs have been used to process syntactic information. These methods obtain syntax trees using external parsers and apply GNNs to the trees for downstream tasks such as pronoun resolution [46], relation extraction [16, 49, 34], and machine translation [3, 4]. In contrast, we apply GNNs to partial trees for the task of parsing. Ji et al. [18] used GNNs for graph-based dependency parsing. However, their method is not transition-based. They apply GNNs to complete graphs formed by all tokens, whereas we apply GNNs to partial trees.

3 Attach-juxtapose Transition System

Overview We introduce a novel transition system named *attach-juxtapose* for strongly incremental constituency parsing. Our system is inspired by psycholinguistic research [26, 38, 36]; it maintains a single parse tree and adds one token to it at each step. Our system can produce valid syntax trees for partial sentences and can handle trees with arbitrary branching factors.

For parsing a sentence of length n , we start with an empty tree and sequentially execute n actions; each action integrates the next token into the current partial tree. Formally speaking, for the sentence $[w_0, w_1, \dots, w_{n-1}]$, the state at i th step is $s_i = (T_i, w_i)$, where T_i is the partial tree for the prefix $[w_0, \dots, w_{i-1}]$. The state transition rules are $T_0 = \text{empty_tree}$ and $T_{i+1} = T_i(a_i)$, where $T_i(a_i)$ denotes the result of executing action a_i on tree T_i . After n steps, we end up with a complete tree T_n . The actions are designed to capture *where* and *how* to integrate a new token into the partial tree.

Where to integrate the new token Since the new token is to the right of existing tokens, it must appear on the *rightmost chain*—the chain of nodes starting from the root and iteratively descending to the rightmost child (A similar observation was also made by Costa et al. [7]). Formally speaking, at the i th step, we have a partial tree T_i and a new token w_i . Let $\text{rightmost_chain}(T_i)$ denote the set of *internal nodes* on the rightmost chain of T_i . We pick $\text{target_node} \in \text{rightmost_chain}(T_i)$ as where the new token should be integrated.

How to integrate the new token Fig. 2 *Top* shows the rightmost chain and target_node (*orange*), we design two types of actions specifying how to integrate the new token (*blue*):

- `attach(target_node, parent_label)`: Attach the token as a descendant of target_node . The parameter `parent_label` is optional; when provided, we create an internal node labeled `parent_label` (*green*) as the parent of the new token. `parent_label` then becomes the rightmost child of target_node (as in Fig. 2 *Top*). When `parent_label` is not provided, the new token itself becomes the rightmost child of target_node .
- `juxtapose(target_node, parent_label, new_label)`: Create an internal node labeled `new_label` (*gray*) as the shared parent of target_node and the new token. It then replaces

target_node in the tree. Similar to attach, we can optionally create a parent for the new token via the parent_label parameter.

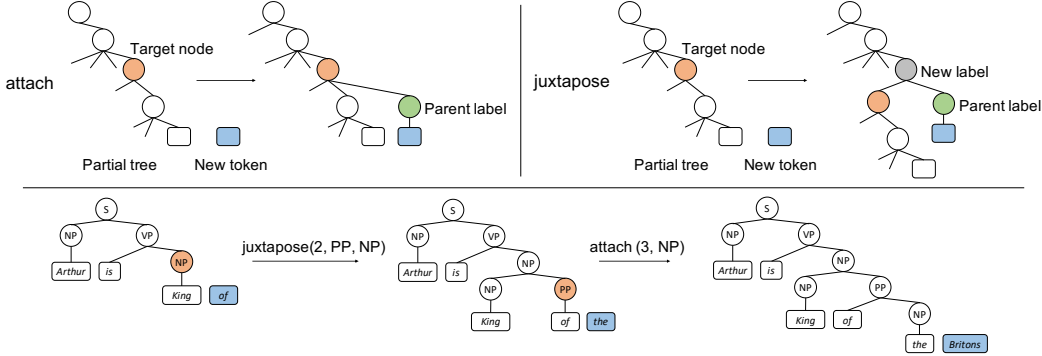


Figure 2: Actions in the attach-juxtapose transition system. *Top-left*: Given a target node (orange) on the rightmost chain, the attach action attaches the new token (blue) as its descendant. *Top-right*: The juxtapose action juxtaposes the new token and the target node in different branches of a shared ancestor (gray). Both actions can optionally create a parent (green) for the new token. *Bottom*: Two example actions when parsing the sentence “Arthur is King of the Britons.”

Fig. 2 *Bottom* shows two example actions when parsing “Arthur is King of the Britons.” We represent target_node using its index on the rightmost chain (starting from 0). The complete action sequence to parse the sentence correctly (Fig. 1) would be: attach(0, NP), juxtapose(0, VP, S), attach(1, NP), juxtapose(2, PP, NP), attach(3, NP), attach(4, None). Note that the first action attach(0, NP) is a degenerated case. Since $T_0 = \text{empty_tree}$, it is impossible to pick target_node $\in \text{rightmost_chain}(T_0)$. In this case, imagine a dummy root node for T_0 ; then we can execute attach(0, parent_label), making parent_label the new root.

Oracle actions Having defined the attach-juxtapose transition system, we are yet to show its capability for constituency parsing: *Given a constituency tree, is it always possible to find a sequence of oracle actions to produce the tree?* This question is important because if the oracle actions did not exist, it would not be possible to parse the sentence correctly. If the oracle actions do exist, a further question is: *For a given tree, is the sequence of oracle actions unique?* Uniqueness is desirable because it guarantees an unambiguous supervision signal at each step when training the parser. We prove that the answers to both questions are positive under mild conditions:

Theorem 1 (Existence of oracle actions). *Let T be a constituency tree for a sentence of length n . If T does not contain unary chains, there exists a sequence of actions a_0, a_1, \dots, a_{n-1} such that $\text{empty_tree}(a_0)(a_1) \dots (a_{n-1}) = T$.*

Theorem 2 (Uniqueness of oracle actions). *Let T be a constituency tree for a sentence of length n , and T does not contain unary chains. If a_0, a_1, \dots, a_{n-1} is a sequence of oracle actions, it is the only action sequence that satisfies $\text{empty_tree}(a_0)(a_1) \dots (a_{n-1}) = T$.*

The condition regarding unary chains is not a restriction in practice, as we can remove unary chains using the preprocessing technique in prior work [21, 50, 28]. The theorems above can be proved by constructing an algorithm to compute the oracle actions. We present the algorithm and detailed proofs in Appendix A. Intuitively, given a tree T , we recursively find and undo the last action until T becomes empty_tree.

Connections with In-order Shift-reduce System Our attach-juxtapose transition system is closely related to In-order Shift-reduce System (ISR) proposed by Liu and Zhang [23]. ISR’s state space is strictly larger than ours; we prove it to be equivalent to an augmented version of our state space. Given a sentence $[w_0, w_1, \dots, w_{n-1}]$, the space of partial trees is $\mathcal{U} = \{t \mid \exists 0 \leq m \leq n, \text{ s.t. } t \text{ is a constituency tree for } [w_0, w_1, \dots, w_{m-1}]\}$. By Theorem 1, our state space (for the given sentence) is a subset of \mathcal{U} , i.e., $\mathcal{U}_{AJ} = \{t \mid t \in \mathcal{U}, t \text{ does not contain unary chains}\}$. To bridge \mathcal{U}_{AJ} and \mathcal{U}_{ISR} (the state space of IRS), we define the augmented space of partial trees to be $\mathcal{U}' = \{(t, i) \mid t \in \mathcal{U}, i \in \mathbb{Z}, -1 \leq i < L(t)\}$, where $L(t)$ denotes the number of internal nodes on the rightmost chain of t . We assert that \mathcal{U}' is equivalent to \mathcal{U}_{ISR} .

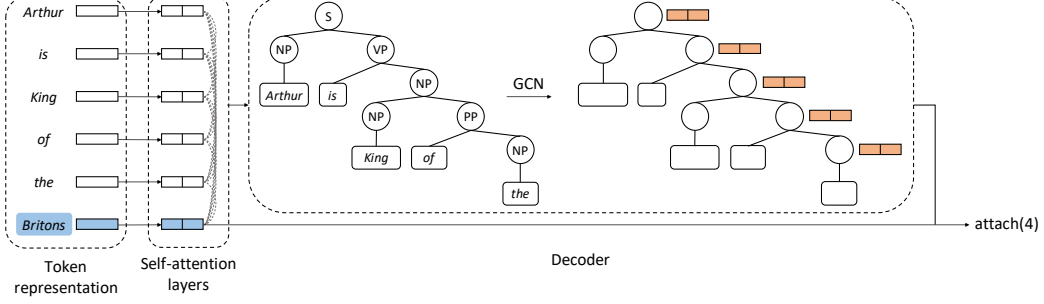


Figure 3: The architecture of our model for action generation. We use the self-attention encoder in prior work [21, 50] and generate actions using a GCN-based [19] decoder.

Theorem 3 (Connection in state spaces). *There is a bijective mapping $\varphi : \mathcal{U}_{ISR} \rightarrow \mathcal{U}'$ between the set of legal states in In-order Shift-reduce System and the augmented space of partial trees.*

We prove the theorem in Appendix B. Intuitively, a state in ISR is a stack configuration; it corresponds to an element of \mathcal{U}' , which can be understood as a partial tree $t \in \mathcal{U}$ with a special node on the rightmost chain marked by an integer i .

Not only is φ bijective, but it also preserves actions. In other words, each action in our system can be mapped to a combination of actions in ISR. To see this, we define an injective mapping $\xi : \mathcal{U}_{AJ} \rightarrow \mathcal{U}'$ such that $\xi(t) = (t, L(t) - 1)$. Then by Theorem 3, $\varphi^{-1} \circ \xi : \mathcal{U}_{AJ} \rightarrow \mathcal{U}_{ISR}$ is an injective mapping from our state space to ISR’s state space. And we have the following connection between actions:

Theorem 4 (Connection in actions). *Let t_1 and t_2 be two partial trees without unary chains, i.e., $t_1, t_2 \in \mathcal{U}_{AJ}$. If a is an attach-juxtapose action that brings t_1 to t_2 , there must exist a sequence of actions in In-order Shift-reduce System that brings $\varphi^{-1} \circ \xi(t_1)$ to $\varphi^{-1} \circ \xi(t_2)$.*

We present a constructive proof in Appendix B, making it possible to translate any action sequence on our system to a longer sequence in ISR. Theorem 4 also implies any reachable parse tree in our system can also be reached in ISR. And by Theorem 1, both our system and ISR can generate any tree without unary chains.

4 Action Generation with Graph Neural Networks

Given the attach-juxtapose system, we develop a model for constituency parsing by generating actions based on the partial tree and the new token (Fig. 3). First, it encodes input tokens as vectors using the self-attention encoder in prior work [21, 50]. These vectors are used to initialize node features in the partial tree, which is then fed into a graph convolutional network (GCN) [19]. The GCN produces features for each node, and in particular, the features on the rightmost chain. Finally, an attention-based action decoder generates the action based on the new token and the rightmost chain.

Token encoder We use the same encoder as Kitaev and Klein [21] and Zhou and Zhao [50]. It consists of a pre-trained contextualized embedding such as BERT [10] or XLNet [48], followed by a few additional self-attention layers. Like in prior work, we separate content and position information; the resulting token features are the concatenation of content features and position features. The encoder is not incremental due to how self-attention works; it is applied once to the entire sentence. Then we apply the decoder to generate actions in a strongly incremental manner.

Graph convolutional neural network We use GCN on the partial tree to produce features for nodes on the rightmost chain. Initially, leaf features are provided by the encoder, whereas for an internal node labeled l that spans from position i to j (endpoints included), the initial feature $x = [W_l, x_p] \in \mathbb{R}^D$ is a concatenation of label and position embeddings: W_l is a $\frac{D}{2}$ -dimensional learned embedding for label l . And $x_p = (P_i + P_j)/2$ is the position embedding averaging two endpoints, where P is the same position embedding matrix in the self-attention encoder.

The initial node features go through several GCN layers with residual connections. We use a variant of the original GCN layer [19] to separate content features and position features (details in Appendix C). The GCN produces features for all nodes. However, we are only interested in nodes on the rightmost chain, as they are candidates for `target_node` in actions.

Action decoder Given the structure of our actions (Sec. 3), the action decoder has to (1) choose a `target_node` on the rightmost chain; (2) decide between `attach` and `juxtapose`; and (3) generate the parameters `parent_label` and `new_label`.

We choose `target_node` using attention on the rightmost chain. Let L be the size of the chain, $Y = [Y_c, Y_p] \in \mathbb{R}^{L \times D}$ be the features on the chain produced by the GCN, $z = [z_c, z_p] \in \mathbb{R}^{1 \times D}$ be the feature of the new token given by the encoder. They are both concatenation of content and position features. We generate attention weights for nodes on the rightmost chain as $w = f_c([Y_c, \mathbf{1}^{L \times 1} z_c]) + f_p([Y_p, \mathbf{1}^{L \times 1} z_p])$, where $\mathbf{1}^{L \times 1}$ is a $L \times 1$ matrix of ones. $[\cdot, \cdot]$ concatenates two matrices horizontally, and f_c, f_p are two-layer fully-connected networks with ReLU [29] and layer normalization [2]. $w \in \mathbb{R}^{L \times 1}$ and we pick the node with maximum attention as `target_node`.

Since the parameter `new_label` is only for `juxtapose`, we can interpret `new_label = None` as `attach`. So we only need to generate `parent_label, new_label` $\in V \cup \{\text{None}\}$, where V is the vocabulary of labels. We generate them using the new token and a weighted average of the rightmost chain: $[u, v] = g([z, \sigma(w)^T Y])$, where σ is the sigmoid function, $u, v \in \mathbb{R}^{|V|+1}$ are predicted log probabilities of `parent_label` and `new_label`, and g is a two-layer layer-normalized network.

Training We train the model to predict oracle actions. At any step, let $a = (\text{target_node}, \text{parent_label}, \text{new_label})$ be the oracle action; recall that `new_label = None` implies `attach`, and `new_label \neq None` implies `juxtapose`. The loss function is a sum of cross-entropy losses for each component: $\mathcal{L} = \mathcal{CE}(w, \text{target_node}) + \mathcal{CE}(u, \text{parent_label}) + \mathcal{CE}(v, \text{new_label})$. For a batch of training examples, the losses are summed across steps and averaged across different examples.

5 Experiments

Setup We evaluate our model for constituency parsing on two standard benchmarks: Chinese Treebank 5.1 (CTB) [47] and the Wall Street Journal part of Penn Treebank (PTB) [25]. PTB consists of 39,832 training examples; 1,700 validation examples; and 2,416 testing examples. Whereas CTB consists of 17,544/352/348 examples for training/validation/testing respectively. Each example is a constituency tree with words and POS tags.

For both datasets, we follow the standard data splits and preprocessing in prior work [23, 35, 21, 50]. In evaluation, we report four metrics—exact match (EM), F1 score, labeled precision (LP), and labeled recall (LR)—all computed by the standard Evalb¹ tool. The testing numbers are produced by models trained on training data alone (not including validation data).

We use the same technique in prior work [21, 50, 28] to remove unary chains by collapsing multiple labels in a unary chain into a single label. It does not affect evaluation, as we revert this process before computing evaluation metrics.

Training details We train the model to predict oracle actions through teacher forcing [45]—the model takes actions according to the oracle rather than the predictions. Model parameters are optimized using RMSProp [41] with a batch size of 32. We decrease the learning rate by a factor of 2 when the best validation F1 score plateaus. The model is implemented in PyTorch [30] and takes 2 \sim 3 days to train on a single Nvidia GeForce GTX 2080 Ti GPU. For fair comparisons with prior work, we use the same pre-trained BERT and XLNet models²: `xlnet-large-cased` and `bert-large-uncased` for PTB; `bert-base-chinese` for CTB.

Parsing performance Table 1 summarizes our PTB results compared to state-of-the-art parsers, including both chart-based parsers [21, 20, 50, 28] and transition-based parsers [23]. Methods with \star

¹<https://nlp.cs.nyu.edu/evalb/>

²https://huggingface.co/transformers/pretrained_models.html

Table 1: Constituency parsing performance on Penn Treebank (PTB). Methods with \star are trained with extra supervision from dependency parsing data. Methods with \dagger are reported in the re-implementation by Fried et al. [13]. Liu and Zhang [23] is transition-based, whereas other baselines are chart-based. We run each experiment 5 times to report the mean and standard error (SEM) of four metrics—exact match (EM), F1 score, labeled precision (LP), and labeled recall (LR). Our method performs competitively with state of the art and achieves the highest EM.

Model	EM	F1	LP	LR	#Params
Liu and Zhang [23]	-	91.8	-	-	-
Liu and Zhang [23] (BERT) \dagger	57.05	95.71	-	-	-
Kitaev and Klein [21]	47.31	93.55	93.90	93.20	26M
Kitaev and Klein [21] (ELMo)	53.06	95.13	95.40	94.85	107M
Kitaev et al. [20] (BERT)	-	95.59	95.46	95.73	342M
Zhou and Zhao [50] (GloVe) \star	47.72	93.78	93.92	93.64	51M
Zhou and Zhao [50] (BERT) \star	55.84	95.84	95.98	95.70	349M
Zhou and Zhao [50] (XLNet) \star	<u>58.73</u>	96.33	96.46	<u>96.21</u>	374M
Mrini et al. [28] (XLNet) \star	58.65	96.38	<u>96.53</u>	96.24	459M
Ours (BERT)	57.29 \pm 0.57	95.79 \pm 0.05	96.04 \pm 0.05	95.55 \pm 0.06	377M
Ours (XLNet)	59.17 \pm 0.33	<u>96.34</u> \pm 0.03	96.55 \pm 0.02	96.13 \pm 0.04	391M

are trained with extra supervision from dependency parsing data. Methods with \dagger are reported in not their original papers but the re-implementation by Fried et al. [13], since the original versions did not use BERT. Liu and Zhang [23] (BERT) \dagger performs beam search during testing with a beam size of 10. We do the same for fair comparisons, which improves the performance marginally (0.05 in EM and 0.02 in F1 for our model with XLNet). Some metrics for prior work are missing because they are neither reported in the original papers nor available using the released model and code. We run each experiment 5 times with different random seeds to report the mean and its standard error (SEM). Overall, our method performs competitively with state-of-the-art parsers on PTB. It achieves higher EM using the same pre-trained embedding (BERT or XLNet). Also, our method has a comparable number of parameters with existing methods.

Table 2: Constituency parsing performance on Chinese Treebank (CTB). \star and \dagger bear the same meaning as in Table 1. Our method outperforms state-of-the-art parsers by a large margin.

Model	EM	F1	LP	LR
Kitaev et al. [20]	-	91.75	91.96	91.55
Kitaev et al. [20] (BERT) \dagger	44.42	92.14	-	-
Zhou and Zhao [50] \star	-	92.18	92.33	<u>92.03</u>
Mrini et al. [28] (BERT) \star	-	<u>92.64</u>	<u>93.45</u>	91.85
Liu and Zhang [23]	-	86.1	-	-
Liu and Zhang [23] (BERT) \dagger	<u>44.94</u>	91.81	-	-
Ours (BERT)	49.72 \pm 0.83	93.59 \pm 0.26	93.80 \pm 0.26	93.40 \pm 0.28

Table 2 summarizes our results on CTB. Our method outperforms existing parsers by a large margin (0.95 in F1). Compared to PTB, the CTB results have a larger SEM. The reason could be that CTB has a small testing set of only 348 examples, leading to less stable evaluation metrics. However, the SEM is still much smaller than our performance margin with existing parsers.

Parsing speed We measure parsing speed empirically using the wall time for parsing the 2,416 PTB testing sentences. Results are shown in Table 3. It takes 33.9 seconds for our method (with XLNet, without beam search), 37.3 seconds for Zhou and Zhao [49], and 40.8 seconds for Mrini et al. [27]. Our method is slightly faster, but the gap is small. About 50% of the time is spent on the XLNet encoder, which is shared among all three methods and explains their similar run time. These experiments were run on machines with 2 CPUs, 16GB memory, and one GTX 2080 Ti GPU.

Table 3: The wall time for parsing the PTB testing set. We run each experiment 5 times.

	Ours (XLNet)	Zhou and Zhao [50] (XLNet)	Mrini et al. [28] (XLNet)
Time (seconds)	33.9 \pm 0.3	37.3 \pm 0.2	40.8 \pm 0.9

Effect of the transition system Our method differs from Liu and Zhang [23] in not only the transition system but also the overall model architecture. To more closely compare our attach-juxtapose transition system with their In-order Shift-reduce System (ISR), we perform an ablation that only changes the transition system while keeping the other part of the model as close as possible.

We implement a baseline that generates ISR actions on top of our encoder and GCNs. To that end, we rely on Theorem 3 to interpret ISR states (stacks) as augmented partial trees. Specifically, given an ISR state $s \in \mathcal{U}_{ISR}$, we have $\varphi(s) \in \mathcal{U}'$ from Theorem 3. We know $\varphi(s) = (t, i)$, where t is a partial tree and i is an integer ranging from -1 to $L(t) - 1$. First, we encode t in the same way as before, using XLNet and GCNs. Then, we take the GCN feature of the i th node on the rightmost chain and use it to generate actions in ISR. We add a special node to the rightmost chain to handle $i = -1$.

Results are shown in Table 4, the ISR baseline achieves an average F1 score of 96.23 on PTB, which is lower than our method (96.34). This ablation demonstrates that the attach-juxtapose transition system contributes to the performance.

Table 4: Comparison on PTB between different transition systems. Both models use XLNet for encoding tokens and GCNs for learning graph features.

Transition system	EM	F1
ISR [23]	58.99 \pm 0.11	96.23 \pm 0.04
Attach-juxtapose	59.17 \pm 0.33	96.34 \pm 0.03

Effect of graph neural networks A key ingredient of our model is using GNNs to effectively leverage structural information in partial trees. We conduct an ablation study to demonstrate its importance. Specifically, we keep the encoder fixed and replace the graph-based decoder with a simple two-layer network. For each new token, it predicts an action from the token feature alone—no partial tree is built. It predicts `target_node` as an integer in $[0, 249]$. We increase the feature dimensions so that both models have the same number of parameters. Results are in Table 5; the graph-based decoder achieves better performance in all settings, which demonstrates the value explicit structural information.

Table 5: Ablation study comparing our graph-based action decoder with a sequence-based decoder that cannot leverage structural information in partial trees. The graph-based decoder leads to better performance in all settings, which demonstrates the importance of explicit structural information.

Decode	BERT (PTB)		XLNet (PTB)		BERT (CTB)	
	EM	F1	EM	F1	EM	F1
Sequence-based	53.26 \pm 0.45	94.89 \pm 0.03	55.84 \pm 0.53	95.54 \pm 0.07	44.14 \pm 0.88	90.98 \pm 0.35
Graph-based	57.29 \pm 0.57	95.79 \pm 0.05	59.17 \pm 0.33	96.34 \pm 0.03	49.72 \pm 0.83	93.59 \pm 0.26

6 Conclusion

We proposed the attach-juxtapose transition system for constituency parsing. It is inspired by the strong incrementality of human parsing discovered by psycholinguistics. We presented theoretical results characterizing its capability and its connections with existing shift-reduce systems. Further, we developed a parser based on it and achieved state-of-the-art performance on two standard benchmarks.

Broader Impact

We evaluated our method on constituency parsing for English and Chinese. They are the two most spoken languages, with more than two billion speakers across the globe. However, there are more than 7,000 languages in the world. And it is important to deliver parsing and other NLP technology to benefit speakers of diverse languages. Fortunately, our method can be applied to many languages with little additional effort. We developed the system using PTB, and when adding CTB, we only had to make a few minor changes in language-specific preprocessing.

However, a potential barrier is the lack of training data for low-resource languages. Our method relies on supervised learning with a large number of annotated parse trees, which are available only for some languages. A potential solution is to do joint multilingual training as in Kitaev et al. [20].

Acknowledgments and Disclosure of Funding

This work is partially supported by the National Science Foundation under Grant IIS-1903222 and the Office of Naval Research under Grant N00014-20-1-2634.

References

- [1] Salah Ait-Mokhtar and Jean-Pierre Chanod. Incremental finite-state parsing. In *Proceedings of the fifth conference on Applied natural language processing*, pages 72–79. Annual Meeting of the Association for Computational Linguistics (ACL), 1997.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. Graph convolutional encoders for syntax-aware neural machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
- [4] Daniel Beck, Gholamreza Haffari, and Trevor Cohn. Graph-to-sequence learning using gated graph neural networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [5] Do Kook Choe and Eugene Charniak. Parsing as language modeling. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [6] Michael Collins and Brian Roark. Incremental parsing with the perceptron algorithm. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2004.
- [7] Fabrizio Costa, Paolo Frasconi, Vincenzo Lombardo, and Giovanni Soda. Towards incremental parsing of natural language using recursive neural networks. *Applied Intelligence*, 19(1-2):9–25, 2003.
- [8] James Cross and Liang Huang. Span-based constituency parsing with a structure-label system and provably optimal dynamic oracles. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
- [9] Marie-Catherine De Marneffe, Bill MacCartney, Christopher D Manning, et al. Generating typed dependency parses from phrase structure parses. In *Lrec*, volume 6, pages 449–454, 2006.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019.
- [11] Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A Smith. Transition-based dependency parsing with stack long short-term memory. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.
- [12] Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A Smith. Recurrent neural network grammars. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2016.
- [13] Daniel Fried, Nikita Kitaev, and Dan Klein. Cross-domain generalization of neural constituency parsers. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [14] David Gaddy, Mitchell Stern, and Dan Klein. What’s going on in neural constituency parsers? an analysis. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.
- [15] Carlos Gómez-Rodríguez and David Vilares. Constituent parsing as sequence labeling. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [16] Zhijiang Guo, Yan Zhang, and Wei Lu. Attention guided graph convolutional networks for relation extraction. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [18] Tao Ji, Yuanbin Wu, and Man Lan. Graph-based dependency parsing with graph neural networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

- [19] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- [20] Nikita Kitaev, Steven Cao, and Dan Klein. Multilingual constituency parsing with self-attention and pre-training. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [21] Nikita Kitaev and Dan Klein. Constituency parsing with a self-attentive encoder. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [22] Jonathan K Kummerfeld, David Hall, James R Curran, and Dan Klein. Parser showdown at the wall street corral: An empirical investigation of error types in parser output. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1048–1059, 2012.
- [23] Jiangming Liu and Yue Zhang. In-order transition-based constituent parsing. *Transactions of the Association for Computational Linguistics*, 5:413–424, 2017.
- [24] Lema Liu, Muhua Zhu, and Shuming Shi. Improving sequence-to-sequence constituency parsing. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [25] Mitchell P Marcus and Mary Ann Marcinkiewicz. Building a large annotated corpus of english: The penn treebank. *Computational Linguistics*, 19(2).
- [26] William Marslen-Wilson. Linguistic structure and speech shadowing at very short latencies. *Nature*, 244(5417):522–523, 1973.
- [27] David Milward. Incremental interpretation of categorial grammar. In *European Chapter of the Association for Computational Linguistics (EACL)*, 1995.
- [28] Khalil Mrini, Franck Dernoncourt, Trung Bui, Walter Chang, and Ndapa Nakashole. Rethinking self-attention: Towards interpretability in neural parsing. *arXiv preprint arXiv:1911.03875*, 2020.
- [29] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *International Conference on Machine Learning (ICML)*, 2010.
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [31] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2018.
- [32] Brian Roark. Probabilistic top-down parsing and language modeling. *Computational Linguistics*, 27(2):249–276, 2001.
- [33] Kenji Sagae and Alon Lavie. A classifier-based parser with linear run-time complexity. In *International Workshop on Parsing Technology*, 2005.
- [34] Sunil Kumar Sahu, Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. Inter-sentence relation extraction with document-level graph convolutional neural network. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [35] Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordani, Aaron Courville, and Yoshua Bengio. Straight to the tree: Constituency parsing with neural syntactic distance. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [36] Edward P Stabler. The finite connectivity of linguistic structure. *Perspectives on Sentence Processing*, 2015.
- [37] Mitchell Stern, Jacob Andreas, and Dan Klein. A minimal span-based neural constituency parser. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017.
- [38] Patrick Sturt and Vincenzo Lombardo. Processing coordinated structures: Incrementality and connectedness. *Cognitive Science*, 29(2):291–305, 2005.
- [39] Jun Suzuki, Sho Takase, Hidetaka Kamigaito, Makoto Morishita, and Masaaki Nagata. An empirical study of building a strong baseline for constituency parsing. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [40] Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.
- [41] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2):26–31, 2012.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [43] Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. Grammar as a foreign language. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2015.
- [44] Taro Watanabe and Eiichiro Sumita. Transition-based neural constituent parsing. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.
- [45] Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, 1989.

- [46] Yinchuan Xu and Junlin Yang. Look again at the syntax: Relational graph convolutional network for gendered ambiguous pronoun resolution. In *Workshop on Gender Bias in Natural Language Processing*, 2019.
- [47] Naiwen Xue, Fei Xia, Fu-Dong Chiou, and Marta Palmer. The penn chinese treebank: Phrase structure annotation of a large corpus. *Natural Language Engineering*, 11(2), 2005.
- [48] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [49] Yuhao Zhang, Peng Qi, and Christopher D Manning. Graph convolution over pruned dependency trees improves relation extraction. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2018.
- [50] Junru Zhou and Hai Zhao. Head-driven phrase structure grammar parsing on penn treebank. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [51] Muhua Zhu, Yue Zhang, Wenliang Chen, Min Zhang, and Jingbo Zhu. Fast and accurate shift-reduce constituent parsing. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2013.

Appendix A: Proofs about Oracle Actions

Given a constituency tree without unary chains, we prove the existence of oracle actions (Theorem 1) by proving the correctness of an algorithm (Algorithm 1) for computing oracle actions. Further, we prove that the oracle action sequence is unique. Before diving into the theorems and proofs, we first define the relevant terms:

Definition 1 (Constituency tree). *Given a sentence $s = [w_0, w_1, \dots, w_{n-1}]$ of length n , we define a constituency tree for s as a rooted tree with arbitrary branching factors. It has n leaves labeled with tokens w_0, w_1, \dots, w_{n-1} from left to the right, whereas its internal nodes are labeled with syntactic categories. The root node must be an internal node. In the degenerated case of $n = 0$, we define a special constant `empty_tree` to be the constituency tree for $s = []$.*

Definition 2 (Unary chain). *Let T be a constituency tree. We say T contains unary chains if there exist two internal nodes x and y such that y is the only child of x . Conversely, if such x and y do not exist, we say T does not contain unary chains.*

Then we present Algorithm 1 for computing oracle actions. Given a constituency tree T without unary chains, it recursively finds and undoes the last action until T becomes `empty_tree`. The algorithm has a time complexity of $\mathcal{O}(n \log n)$, where n is the sentence length. It calls `last_action` n times, and each call needs $\mathcal{O}(\log n)$ for locating the last leaf of the tree.

Now we are ready to state and prove Theorem 1 and Theorem 2 in the main paper.

Theorem 1 (Existence of oracle actions). *Let T be a constituency tree for a sentence of length n . If T does not contain unary chains, there exists a sequence of actions a_0, a_1, \dots, a_{n-1} such that `empty_tree`(a_0)(a_1)...(a_{n-1}) = T . And this sequence of actions can be computed via Algorithm 1.*

Proof. We prove the correctness of Algorithm 1 by induction on the sentence length n .

When $n = 0$, we have $T = \text{empty_tree}$ (Definition 1), and `oracle_action_sequence`(T) returns an empty action sequence `[]`. The conclusion holds straightforwardly.

When $n > 0$, it is sufficient to prove T' is a valid constituency tree without unary chains for a sentence of length $n - 1$. We proceed by enumerating all possible execution traces in `last_action`. The function contains two conditional statements and therefore $2 \times 3 = 6$ execute traces. We use “Case i - j ” to denote the execution trace taking the i th branch in the first conditional statement and the j th branch in the second conditional statement.

- Case 1-1—`last_leaf` has siblings, and `last_subtree` is the root node.
We have `last_subtree` = `last_leaf` (the first conditional statement). So `last_leaf` is the root node while being a leaf, which contradicts with the assumption that T is a constituency tree (Definition 1).

Algorithm 1: Computing the oracle actions for a constituency tree without unary chains

```

1 def oracle_action_sequence( $T$ ):
2   if  $T == \text{empty\_tree}$ :
3     return []
4   else:
5      $a_{n-1} = \text{last\_action}(T)$ 
6      $T' = \text{Undo the last action } a_{n-1} \text{ on } T$ 
7     return oracle_action_sequence( $T'$ ) + [ $a$ ]
8
9 def last_action( $T$ ):
10   $\text{last\_leaf} = \text{The last (rightmost) leaf in } T$ 
11  if  $\text{last\_leaf}$  has siblings:
12     $\text{parent\_label} = \text{None}$ 
13     $\text{last\_subtree} = \text{last\_leaf}$ 
14  else:
15     $\text{parent\_label} = \text{The label of last\_leaf's parent}$ 
16     $\text{last\_subtree} = \text{last\_leaf's parent}$ 
17
18  if  $\text{last\_subtree}$  is the root node:
19    return attach(0,  $\text{parent\_label}$ )
20  elif  $\text{last\_subtree}$  has exactly one sibling and its sibling is an internal node:
21     $\text{new\_label} = \text{The label of last\_subtree's parent}$ 
22     $\text{target\_node} = \text{The index of last\_subtree's sibling}$ 
23    return juxtapose( $\text{target\_node}$ ,  $\text{parent\_label}$ ,  $\text{new\_label}$ )
24  else:
25     $\text{target\_node} = \text{The index of last\_subtree's parent}$ 
26    return attach( $\text{target\_node}$ ,  $\text{parent\_label}$ )
  
```




Figure A: Case 1-2. The last action is a juxtapose.

- Case 1-2— last_leaf has siblings; last_subtree is not the root node; last_subtree has exactly one sibling, and its sibling is an internal node.
We have $\text{last_subtree} = \text{last_leaf}$ (the first conditional statement). The local configuration of T looks like Fig. A *Right*, on its left is T' obtained from T by undoing action $\text{juxtapose}(\text{target_node}, \text{parent_label}, \text{new_label})$. T' is still a valid constituency tree without unary chains.
- Case 1-3— last_leaf has siblings; last_subtree is not the root node; last_subtree has either no sibling, one leaf node as its sibling, or more than one siblings.
We have $\text{last_subtree} = \text{last_leaf}$ (the first conditional statement). So last_subtree has either one leaf node or more than one nodes as its siblings. These two cases are shown separately in Fig. B. In both cases, T' is still a valid constituency tree without unary chains.
- Case 2-1— last_leaf has no sibling, and last_subtree is the root node.
We have $\text{last_subtree} = \text{last_leaf's parent}$ (the first conditional statement). As shown in Fig. C, T' is empty_tree in this case, which is also a valid constituency tree without unary chains.
- Case 2-2— last_leaf has no sibling; last_subtree is not the root node; last_subtree has exactly one sibling, and its sibling is an internal node.




Figure B: Case 1-3. There are two possible cases depending on the number of siblings of *last_subtree*. In both cases, the last action is an attach.




Figure C: Case 2-1. $T' = \text{empty_tree}$ and *last_subtree* is the root of T . The last action is an attach.

We have *last_subtree* = *last_leaf*'s parent (the first conditional statement). As Fig. D shows, T' is still a valid constituency tree without unary chains.

- Case 2-3—*last_leaf* has no sibling; *last_subtree* is not the root node; *last_subtree* has either no sibling, one leaf node as its sibling, or more than one siblings.
We have *last_subtree* = *last_leaf*'s parent (the first conditional statement). So *last_subtree* is an internal node. Since T does not contain unary chains, any non-root internal node must have siblings. As a result, *last_subtree* has either one leaf node or more than one nodes as its sibling. These two cases are shown separately in Fig. E. In both cases, T' is still a valid constituency tree without unary chains.

We have proved T' to be a valid constituency tree for a sentence of length $n-1$ no matter which execution trace *last_action* takes. Applying the induction hypothesis, we know *oracle_action_sequence*(T') outputs a sequence of actions a_0, a_1, \dots, a_{n-2} such that *empty_tree*(a_0)...(a_{n-2}) = T' . Since $T'(a_{n-1}) = T$, we have finally derived *empty_tree*(a_0)...(a_{n-1}) = T .

□

Theorem 2 (Uniqueness of oracle actions). *Let T be a constituency tree for a sentence of length n , and T does not contain unary chains. If *oracle_action_sequence*(T) = a_0, a_1, \dots, a_{n-1} , it is the only action sequence that satisfies *empty_tree*(a_0)(a_1)...(a_{n-1}) = T .*

Proof. We prove by contradiction. Assume there is different action sequence $a'_0, a'_1, \dots, a'_{n-1}$ that satisfies *empty_tree*(a'_0)(a'_1)...(a'_{n-1}) = T . We first prove $a'_{n-1} = a_{n-1}$, in other words, a_{n-1} is the only possible last action. Similar to Theorem 1, we prove by enumerating all execution traces.

- Case 1-1—*last_leaf* has siblings, and *last_subtree* is the root node.
Similarly, it contradicts with T being a constituency tree.
- Case 1-2—*last_leaf* has siblings; *last_subtree* is not the root node; *last_subtree* has exactly one sibling, and its sibling is an internal node.




Figure D: Case 2-2. The last action is an juxtapose.




Figure E: Case 2-3. There are two possible cases depending on the number of siblings of *last_subtree*. In both cases, the last action is an attach.

In Fig. A *Right*, it is apparent that $parent_label = \text{None}$. Also, a'_{n-1} must be a juxtapose action since otherwise the gray node will introduce a unary chain in T' . Therefore, $a'_{n-1} = a_{n-1}$.

- Case 1-3—*last_leaf* has siblings; *last_subtree* is not the root node; *last_subtree* has either no sibling, one leaf node as its sibling, or more than one siblings.

In Fig. B, $parent_label = \text{None}$. No matter how many siblings *last_subtree* has, a'_{n-1} must be an attach action. Therefore, $a'_{n-1} = a_{n-1}$.

The remaining three cases are similar, and we omit the details. Now that we have proved $a'_{n-1} = a_{n-1}$, it is straightforward to apply the same reasoning to derive $a'_{n-2} = a_{n-2}$, $a'_{n-3} = a_{n-3}$ and all the way until $a'_0 = a_0$. This contradicts with the assumption that $a'_0, a'_1, \dots, a'_{n-1}$ is different from a_0, a_1, \dots, a_{n-1} . Therefore, we have a unique sequence of oracle actions a_0, a_1, \dots, a_{n-1} . □

Appendix B: Proofs about Connections with In-order Shift-reduce System

We prove Theorem 3 and Theorem 4 in the main paper; they reveal the connections between our system with In-order Shift-reduce System (ISR) proposed by Liu and Zhang [23]. Before any formal derivation, we first illuminate the connections through an example to gain some intuition.

Fig. F shows a juxtapose action when parsing “Arthur is King of the Britons.” in our system, which can be translated into 4 actions in ISR: *reduce*, *PJ-NP*, *shift*, *PJ-PP*. In the figure, \mathcal{U}_{AJ} denotes

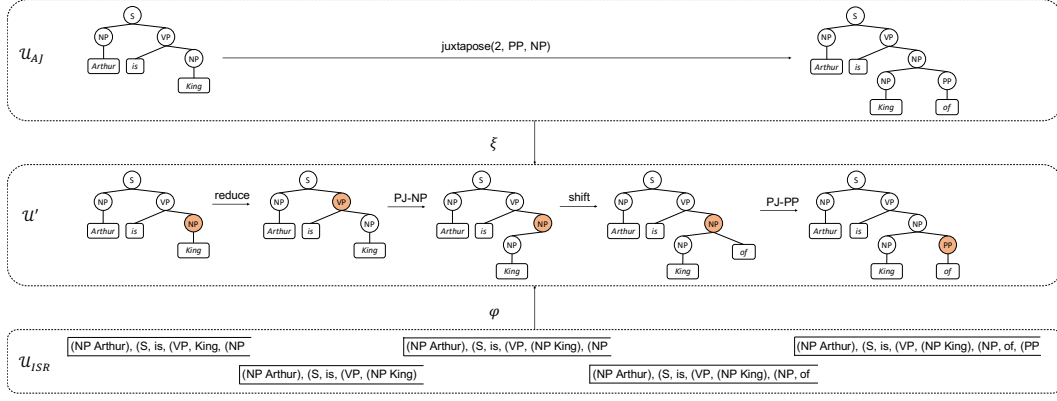


Figure F: An juxtapose action when parsing “Arthur is King of the Britons.” in our attach-juxtapose system. It can be translated into 4 actions in In-order Shift-reduce System.

our state space—the set of partial trees without unary chains. Whereas \mathcal{U}_{ISR} denotes ISR’s state space—the set of legal stack configurations. We represent a stack from left (stack bottom) to the right (stack top). “(X)” denotes a projected nonterminal X, while an S-expression such as “(NP Arthur)” denotes a subtree in the stack. \mathcal{U}' denotes the augmented space of partial trees; each element in \mathcal{U}' is a constituency tree that may have a node on the rightmost chain marked as special (orange). We observe a one-to-one correspondence between \mathcal{U}' and \mathcal{U}_{ISR} , which is denoted by the mapping φ .

We proceed to generalize this example to arbitrary state transitions in our system, which involves formulating and proving Theorem 3 and Theorem 4 in the main paper. First, we assume a fixed sentence and define the state spaces of our system and ISR:

Definition 3 (Space of partial trees). *Given the sentence $[w_0, w_1, \dots, w_{n-1}]$, we define the space of partial trees to be:*

$$\mathcal{U} = \{t \mid \exists 0 \leq m \leq n, \text{ s.t. } t \text{ is a constituency tree for } [w_0, w_1, \dots, w_{m-1}]\}. \quad (1)$$

From theorem 1, we know that the state space of our system is a subset of \mathcal{U} that does not contain unary chains:

$$\mathcal{U}_{AJ} = \{t \mid t \in \mathcal{U}, t \text{ does not contain unary chains}\}. \quad (2)$$

Definition 4 (Augmented space of partial trees). *Given the sentence $[w_0, w_1, \dots, w_{n-1}]$. We define the augmented space of partial trees to be:*

$$\mathcal{U}' = \{(t, i) \mid t \in \mathcal{U}, i \in \mathbb{Z}, -1 \leq i < L(t)\}, \quad (3)$$

where $L(t)$ denotes the number of internal nodes on the rightmost chain of t . We can define an injective mapping $\xi : \mathcal{U}_{AJ} \rightarrow \mathcal{U}'$:

$$\xi(t) = (t, L(t) - 1). \quad (4)$$

In ISR, the parser can be trapped in a state that will never lead to a complete tree no matter what actions it takes, e.g., the states where two tokens are at the bottom of the stack. We call such states illegal states and prove a lemma characterizing the set of legal states.

Lemma 1 (Legal states in ISR). *Let s be a stack configuration in In-order Shift-reduce System, it is a legal state if and only if you can end up with a single partial tree in the stack by repeatedly executing reduce.*

Proof. For the “if” part, we first keep executing reduce until there is only a single partial tree in the stack. Then we can arrive at a complete tree by executing one PJ-X, several shift to consume all remaining tokens, and one final reduce. Therefore, the stack s is legal.

For the “only if” part, s is legal. We prove by contradiction, assuming it is impossible to get a single partial tree by executing multiple reduce. Then we must be stuck somewhere. Referring to the definition of the reduce action [23], there could be several reasons for being stuck: (1) The stack has more than one element but no projected nonterminal; (2) the only projected nonterminal is at the bottom of the stack; (3) there are two consecutive projected nonterminals. For all three cases, the offending pattern must also exist in the original s , and no action sequence can remove them from s . Therefore, s must be illegal, which contradicts the assumption. \square

Although there are illegal states in ISR, it is possible to avoid them using heuristics in practice. So we do not consider them a problem for ISR. In the following derivations, we safely ignore illegal states and assume ISR's state space to consist of only legal states:

Definition 5 (Space of legal stack configurations). *Given the sentence $[w_0, w_1, \dots, w_{n-1}]$, we define \mathcal{U}_{ISR} as the set of legal stack configurations in In-order Shift-reduce System.*

As one of our main theoretical conclusions, ISR's state space is equivalent to the augmented space of partial trees; therefore, it is strictly larger than our state space:

Theorem 3 (Connection in state spaces). *There is a bijective mapping $\varphi : \mathcal{U}_{ISR} \rightarrow \mathcal{U}'$ between the legal states in In-order Shift-reduce System and the augmented space of partial trees.*

Proof. For a legal stack $s \in \mathcal{U}_{ISR}$, let $L(s)$ be the number of projected nonterminals in s (introduced by PJ-X actions in Liu and Zhang [23]). We are abusing the notation a little bit as we have used $L(t)$ to denote the number of internal nodes on the rightmost chain of a tree t . But as we will show, they are actually the same. If s is an empty stack, let $t = \text{empty_tree}$. Otherwise, let t be the tree produced by repeatedly executing `reduce` until there is only one partial tree remaining in the stack. This is always possible since s is a legal state (Lemma 1). Then we can define $\varphi(s) = (t, L(s) - 1)$. We prove φ is bijective by constructing its inverse mapping $\varphi^{-1} : \mathcal{U}' \rightarrow \mathcal{U}_{ISR}$.

Given any $(t, i) \in \mathcal{U}'$, t is a partial tree. We define the depth of a node in t as its distance to the root node. Further, we extend the definition of in-order traversal from binary trees to trees with arbitrary branching factors: the first subtree \rightarrow root node \rightarrow the second subtree \rightarrow the third subtree $\rightarrow \dots$

We define a mapping $\gamma : \mathcal{U}' \rightarrow \mathcal{U}_{ISR}$. Let $\gamma(t, i)$ be the stack obtained by starting with an empty stack and traversing the tree t in-order: At any subtree rooted at node x , (1) if node x is not on the rightmost chain or $\text{depth}(x) > i$, we push the entire subtree x onto the stack and skip traversing the nodes in it. (2) If node x is on the rightmost chain and $\text{depth}(x) \leq i$, we push node x 's label as a projected nonterminal and keep traversing the nodes in subtree x .

We now prove γ to be the inverse of φ , i.e. $\varphi \circ \gamma(t, i) = (t, i)$. It is straightforward that the stack $\gamma(t, i)$ has $i + 1$ projected nonterminals, corresponding to nodes on the rightmost chain with depth $0, 1, \dots, i$. So, $L(\gamma(t, i)) - 1 = i$, and we only have to prove t to be the tree obtained by repeatedly executing `reduce` on $\gamma(t, i)$.

In the trivial case of $t = \text{empty_tree}$, we have $L(t) = 0$, and i must be -1 . $\gamma(t, i)$ is an empty stack, and executing `reduce` on it will give `empty_tree`, which equals to t .

In the non-trivial case of $t \neq \text{empty_tree}$, we prove by induction on the number of `reduce` actions (k) executed on the stack $\gamma(t, i)$ to get a single tree.

When $k = 0$, $\gamma(t, i)$ contains a single tree, which must be t itself.

When $k > 0$, let $\gamma_1(t, i)$ be the stack after executing one `reduce` on $\gamma(t, i)$. We assert that $\gamma_1(t, i) = \gamma(t, i - 1)$. We can see this by comparing the in-order traversal of (t, i) and $(t, i - 1)$. The only difference is how we process the subtree rooted at the i th node on the rightmost chain. When traversing (t, i) , we push a projected nonterminal and proceed to nodes in the subtree. When traversing $(t, i - 1)$, we push the entire subtree and skip the nodes in it, which corresponds exactly to executing one `reduce` on $\gamma(t, i)$. Therefore, $\gamma_1(t, i) = \gamma(t, i - 1)$.

We only need $k - 1$ `reduce` actions to get a single tree from the stack $\gamma_1(t, i)$, or equivalently, $\gamma(t, i - 1)$. Applying the induction hypothesis, we will get the tree t by repeatedly applying `reduce` on $\gamma_1(t, i)$. Therefore, we will also get the same t by repeatedly applying `reduce` on $\gamma(t, i)$, i.e. $\varphi \circ \gamma(t, i) = (t, i)$.

Since (t, i) is arbitrary, we have proved γ to be the inverse mapping of φ , and φ is thus bijective. □

Corollary 1 (Connections in state spaces). $\varphi^{-1} \circ \xi : \mathcal{U}_{AJ} \rightarrow \mathcal{U}_{ISR}$ is an injective mapping from our state space to ISR's state space.

Proof. It is straightforward given that $\xi : \mathcal{U}_{AJ} \rightarrow \mathcal{U}'$ is injective (Definition 4) and $\varphi : \mathcal{U}_{ISR} \rightarrow \mathcal{U}'$ is bijective (Theorem 3). □

The mapping $\varphi^{-1} \circ \xi$ bridges our state space and ISR's state space. Not only is it injective, but it also preserves actions—each action in our system can be mapped to a combination of actions in ISR. We prove this for *attach* actions (Lemma 2) and *juxtapose* actions (Lemma 3) separately.

Lemma 2 (Translating *attach* actions to ISR). *Let t_1 and t_2 be two partial trees without unary chains, i.e., $t_1, t_2 \in \mathcal{U}_{aj}$. If $\text{attach}(i, X)$ brings t_1 to t_2 , The following action sequence in In-order Shift-reduce System will bring $\varphi^{-1} \circ \xi(t_1)$ to $\varphi^{-1} \circ \xi(t_2)$:*

$$\underbrace{\text{reduce}, \dots, \text{reduce}}_{L(t_1)-i-1}, \quad \text{shift}, \quad \underbrace{\text{PJ-X}}_{\text{if } X \neq \text{None}}, \quad (5)$$

where $L(t_1)$ is the number of internal nodes on the rightmost chain of t_1 . $X = \text{None}$ means the optional argument *parent_label* is not provided; in this case, we exclude *PJ-X*.

Proof. We know $\varphi^{-1} \circ \xi(t_1) = \varphi^{-1}(t_1, L(t_1) - 1)$ (Definition 4). Recall that in Theorem 3 we have proved that executing one *reduce* on $\varphi^{-1}(t, i)$ gives us $\varphi^{-1}(t, i - 1)$. Therefore, executing *reduce* $L(t_1) - i - 1$ times on $\varphi^{-1}(t_1, L(t_1) - 1)$ gives us $\varphi^{-1}(t_1, i)$. We still have to execute one *shift* and one optional *PJ-X*.

When $X = \text{None}$, we only have to execute a *shift*. The resulting stack is $\varphi^{-1}(t_1, i)$ plus a new token at the top. We prove that the new stack equals to $\varphi^{-1}(t_2, L(t_2) - 1)$. Since t_2 is a result of executing *attach*(*i*, *None*) on t_1 , we know $L(t_2) - 1 = i$ from the definition of the *attach* action. So, we only have to prove that the new stack equals to $\varphi^{-1}(t_2, i)$. We unfold the definition of φ^{-1} (in the proof of Theorem 3) and compare the in-order traversal of (t_2, i) and (t_1, i) . When visiting the subtree rooted at *target_node* i in t_2 , we have the new token as the rightmost child; it corresponds to the new token at the top of the stack. Therefore, $\varphi^{-1}(t_2, i)$ equals to $\varphi^{-1}(t_1, i)$ plus the new token. We have proved the new stack to be $\varphi^{-1}(t_2, L(t_2) - 1)$ and therefore it equals to $\varphi^{-1} \circ \xi(t_2)$ (Definition 4).

When $X \neq \text{None}$, we have to execute a *shift* and a *PJ-X*. The resulting stack is $\varphi^{-1}(t_1, i)$ plus a new token and a projected nonterminal X at the top. Similarly, we want to prove the new stack to equal to $\varphi^{-1}(t_2, L(t_2) - 1)$. The reasoning is similar, by comparing the in-order traversal of (t_1, i) and (t_2, i) . We thus omit the details.

Therefore, in ISR, we can arrive at the state $\varphi^{-1} \circ \xi(t_2)$ from the state $\varphi^{-1} \circ \xi(t_1)$ by executing $L(t_1) - i - 1$ *reduce* actions, one *shift* action and one optional *PJ-X* action. \square

Lemma 3 (Translating *juxtapose* actions to ISR). *Let t_1 and t_2 be two partial trees without unary chains, i.e., $t_1, t_2 \in \mathcal{U}_{aj}$. If $\text{juxtapose}(i, X, Y)$ brings t_1 to t_2 , The following action sequence in In-order Shift-reduce System will bring $\varphi^{-1} \circ \xi(t_1)$ to $\varphi^{-1} \circ \xi(t_2)$:*

$$\underbrace{\text{reduce}, \dots, \text{reduce}}_{L(t_1)-i}, \quad \text{PJ-Y}, \quad \text{shift}, \quad \underbrace{\text{PJ-X}}_{\text{if } X \neq \text{None}}. \quad (6)$$

Proof. Similar to Lemma 2, we first execute $L(t_1) - i$ *reduce* actions on $\varphi^{-1} \circ \xi(t_1) = \varphi^{-1}(t_1, L(t_1) - 1)$ to get $\varphi^{-1}(t_1, i - 1)$. We still have to execute one *PJ-Y*, one *shift* and one optional *PJ-X*.

When $X = \text{None}$, we only have to execute a *PJ-Y* and a *shift*. The resulting stack is $\varphi^{-1}(t_1, i - 1)$ plus a projected nonterminal Y and a new token. We prove that the new stack equals to $\varphi^{-1}(t_2, L(t_2) - 1)$. Since t_2 is a result of executing *juxtapose*(*i*, *None*, Y) on t_1 , we know $L(t_2) - 1 = i$ from the definition of the *juxtapose* action. So, we only need the new stack to equal to $\varphi^{-1}(t_2, i)$, which can be proved by comparing the in-order traversal of (t_2, i) and $(t_1, i - 1)$: In t_2 , the subtree rooted at Y has 2 children; the left child corresponds to the i th subtree on the rightmost chain of t_1 , whereas the right child is a single leaf. When we reach Y in the in-order traversal of t_2 , we first push its left subtree onto the stack. Now the stack equals to $\varphi^{-1}(t_1, i - 1)$. Then, we visit the node Y and its right child, which push the additional projected nonterminal Y and a new token onto the stack. Therefore, $\varphi^{-1}(t_2, i)$ is the new stack.

When $X \neq \text{None}$, we have to execute a *PJ-Y*, a *shift*, and a *PJ-X*. The derivation is similar, so we omit the details. \square

Theorem 4 (Connection in actions). *Let t_1 and t_2 be two partial trees without unary chains, i.e., $t_1, t_2 \in \mathcal{U}_{aj}$. If a is an attach-juxtapose action that brings t_1 to t_2 , there must exist a sequence of actions in In-order Shift-reduce System that brings $\varphi^{-1} \circ \xi(t_1)$ to $\varphi^{-1} \circ \xi(t_2)$.*

Proof. It follows straightforwardly from Lemma 2 and Lemma 3. \square

Corollary 2 (Connection in action sequences). *Let t_1 and t_2 be two partial trees without unary chains, i.e., $t_1, t_2 \in \mathcal{U}_{aj}$. If there exists a sequence of attach-juxtapose actions that brings t_1 to t_2 , there must exist a sequence of actions in In-order Shift-reduce System that brings $\varphi^{-1} \circ \xi(t_1)$ to $\varphi^{-1} \circ \xi(t_2)$.*

Proof. It is straightforward to prove using Theorem 4 and induction on the number of actions for bringing t_1 to t_2 . Details are omitted. \square

Appendix C: Separating Content and Position Features in GCN Layers

We use the encoder in Kitaev and Klein [21]. They showed that separating content and position features improves parsing performance. Specifically, the input token features to the encoder are the concatenation of content features (e.g., from BERT [10] or XLNet [48]) and position features from a learnable position embedding matrix P . Kitaev and Klein [21] proposed a variant of self-attention that processes two types of features independently. As a result, the output token features are also the concatenation of content and position.

We extend this idea to GCN layers. Vanilla GCN layers perform a linear transformation on the input node feature (i.e. $y = \Theta x + b$) before normalizing and aggregating the neighbors³. We assume the node features to be concatenation of content and position: $x = [x_c, x_p]$, and perform linear transformations for them separately: $y = [y_c, y_p] = [\Theta_c x + b_c, \Theta_p x + b_p]$. As a result, the node features at each GCN layer are also concatenation of content and position. As stated in the main paper, we merge the two types of features when generating attention weights w for nodes on the rightmost chain.

To study the effect of the separation, we present an ablation experiment. We compare our model with vanilla GCN layers and with GCN layers that separate content and position. We also scale the feature dimensions so that both models have approximately the same number of parameters. Results are summarized in Table A. We run each experiment 5 times and report the mean and its standard error (SEM). For PTB, we use XLNet as the pre-trained embeddings. Results show that separating content and position improves performance in all settings, which is consistent with prior work [21].

Table A: Ablation study separating content and positions.

GCN layer	PTB		CTB	
	EM	F1	EM	F1
Vanilla GCN [19]	58.90 \pm 0.23	96.25 \pm 0.06	49.54 \pm 0.65	93.49 \pm 0.27
Separating content and position	59.17 \pm 0.33	96.34 \pm 0.03	49.72 \pm 0.83	93.59 \pm 0.26

Appendix D: Error Categorization using Berkeley Parser Analyzer

We use Berkeley Parser Analyzer [22] to categorize the errors of Mrini et al. [28] and our model (with XLNet) on PTB. Results are shown in Table B. Two methods have the same relative ordering of error categories. The 3 most frequent categories are “PP Attachment”, “Single Word Phrase”, and “Unary”. Compared to Mrini et al., our method has more “PP Attachment” (342 vs. 320) and “UNSET move” (33 vs. 23), but fewer “Clause Attachment” (110 vs. 122) and “XoverX Unary” (48 vs. 56).

³https://pytorch-geometric.readthedocs.io/en/latest/notes/create_gnn.html

Table B: Categorization of parsing errors made by Mrini et al. [28] and our model (with XLNet). For each category, we show its occurrence and the number of brackets attributed to it.

Error category	Mrini et al. [28]		Ours (XLNet)	
	#Errors	#Brackets	#Errors	#Brackets
PP Attachment	320	746	342	804
Single Word Phrase	267	325	259	324
Unary	237	237	235	235
NP Internal Structure	200	230	202	236
Different label	197	394	197	394
Modifier Attachment	137	251	131	283
Clause Attachment	122	398	110	334
UNSET add	85	85	88	88
UNSET remove	78	78	71	71
Co-ordination	68	138	65	120
XoverX Unary	56	56	48	48
NP Attachment	44	160	39	135
UNSET move	23	62	33	113
VP Attachment	12	37	17	56