

Fast-R2D2: A Pretrained Recursive Neural Network based on Pruned CKY for Grammar Induction and Text Representation

Xiang Hu[†] Haitao Mi^{†*} Liang Li[‡] Gerard de Melo[§]

Ant Group[†]

{aaron.hx, haitao.mi}@alibaba-inc.com[†]

School of Cyber Science and Technology, Shandong University, China /

Key Laboratory of Cryptologic Technology and

Information Security of Ministry of Education, Shandong University /

Quancheng Laboratory, China[‡]

li.liang@sdu.edu.cn[‡]

Hasso Plattner Institute / University of Potsdam[§]

gdm@demelo.org[§]

Abstract

Chart-based models have shown great potential in unsupervised grammar induction, running recursively and hierarchically, but requiring $O(n^3)$ time-complexity. The Recursive Transformer based on Differentiable Trees (R2D2) makes it possible to scale to large language model pretraining even with a complex tree encoder, by introducing a heuristic pruning method. However, its rule-based pruning process suffers from local optima and slow inference. In this paper, we propose a unified R2D2 method that overcomes these issues. We use a top-down unsupervised parser as a model-guided pruning method, which also enables parallel encoding during inference. Our parser casts parsing as a split point scoring task by first scoring all split points for a given sentence and then using the highest-scoring one to recursively split a span into two parts. The reverse order of the splits is considered as the order of pruning in the encoder. We optimize the unsupervised parser by minimizing the Kullback–Leibler distance between tree probabilities from the parser and the R2D2 model. Our experiments show that our Fast-R2D2 significantly improves the grammar induction quality and achieves competitive results in downstream tasks.¹

1 Introduction

Compositional, hierarchical, and recursive processing are widely believed to be essential traits

^{*}Work done while at Ant Group. To contact Haitao, haitaomi@global.tencent.com

¹The code is available at: https://github.com/alipay/StructuredLM_RTDT

of human language across diverse linguistic theories (Chomsky, 1956, 2014). Chart-based models (Maillard et al., 2017; Kim et al., 2019a; Drozdov et al., 2019; Hu et al., 2021) have made promising progress in both grammar induction and hierarchical encoding in recent years. The differential CKY encoding architecture of Maillard et al. (2017) simulates the hierarchical and recursive process explicitly by introducing an energy function to combine all possible derivations when constructing each cell representation. However, this entails a cubic time complexity, which makes it impossible to scale to large language model training like BERT (Devlin et al., 2018). Simultaneously, its cubic memory cost also limits the tree encoder’s ability to draw on huge parameter models as a backbone.

Hu et al. (2021) introduced a heuristic pruning method, successfully reducing the time complexity to a linear number of compositions. The experiments show that chart-based models exhibit great potential for grammar induction and representation learning when applying a sophisticated tree encoder such as Transformers with large corpus pretraining, leading to a Recursive Transformer based on Differentiable Trees, or R2D2 for short. However, R2D2’s heuristic pruning approach is rule-based and only considers certain composition probabilities. Thus, trees constructed in this way are not guaranteed to be globally optimal. Moreover, as each step during pruning is based on previous decisions, the entire encoding process is sequential and thus slow in the inference stage.

In this work, we resolve these issues by proposing a unified method with a new global pruning

strategy based on a lightweight and fast top-down parser. We cast parsing as split point scoring, where we first encode the input sentence with a bi-directional LSTM, and score all split points in parallel. Specifically, for a given sentence, the parser first scores each split point between words in parallel by looking at its left and right contexts, and then recursively splits a span (starting with the whole sentence) into two sub-spans by picking the highest-scoring split point among the current split candidates. Subsequently, the reverse order of the sorted split points can serve as the merge order to guide the pruning of the CKY encoder, which enables the encoder to search for more reasonable trees. As the gradient of the pretrained component cannot be back-propagated to the parser, inspired by URNNG (Kim et al., 2019b), we optimize the parser by taking trees sampled from the CKY chart table generated by the encoder as ground truth. Thus, the parser and the chart-based encoder promote each other in this way back and force just like the strategy and value networks in AlphaZero (Silver et al., 2017). Additionally, the pretrained tree encoder can compose sequences recursively in parallel according to the trees generated by the parser, which makes Fast-R2D2 a Recursive Neural Network (Pollack, 1990; Socher et al., 2013) variant.

In this paper, we make the following main contributions:

1. We propose an architecture to jointly pretrain parser and encoder of a recursive network in linear memory cost. Experiments show that our pretrained parser outperforms models custom-tailored for grammar induction.
2. By encoding in parallel following trees generated by the top-down parser, Fast-R2D2 significantly improves the inference speed 30 to 50 fold compared to R2D2.
3. We pre-train Fast-R2D2 on a large corpus and evaluate it on downstream tasks. The experiments demonstrate that a pretrained recursive model based on an unsupervised parser significantly outperforms pretrained sequential Transformers (Vaswani et al., 2017) with the same parameter size in single sentence classification tasks.

2 Preliminaries

2.1 R2D2 Architecture

Differentiable Trees. R2D2 follows the work of Maillard et al. (2017) in defining a CKY-

style (Cocke, 1969; Kasami, 1966; Younger, 1967) encoder. For a sentence $\mathbf{S} = \{s_1, s_2, \dots, s_n\}$ with n words or word-pieces, it defines a chart table as illustrated in Figure 1. In the table, each cell $\mathcal{T}_{i,j}$ is a tuple $\langle e_{i,j}, p_{i,j}, \tilde{p}_{i,j} \rangle$, where $e_{i,j}$ is a vector representation, $p_{i,j}$ is the probability of a single composition step, and $\tilde{p}_{i,j}$ is the probability of the subtree for the span $[i, j]$ over the sub-string $s_{i:j}$. When i equals j , the table has terminal nodes $\mathcal{T}_{i,i}$ with $e_{i,i}$ initialized with the embeddings of input tokens s_i , while $p_{i,i}$ and $\tilde{p}_{i,i}$ are set to one. When $j > i$, the representation $e_{i,j}$ is a weighted sum of intermediate combinations $c_{i,j}^k$, defined as:

$$c_{i,j}^k, p_{i,j}^k = f(e_{i,k}, e_{k+1,j}) \quad (1)$$

$$\tilde{p}_{i,j}^k = p_{i,j}^k \tilde{p}_{i,k} \tilde{p}_{k+1,j} \quad (2)$$

$$\alpha_{i,j} = \text{GUMBEL}(\log(\tilde{\mathbf{p}}_{i,j})) \quad (3)$$

$$e_{i,j} = [c_{i,j}^i, c_{i,j}^{i+1}, \dots, c_{i,j}^{j-1}] \alpha_{i,j} \quad (4)$$

$$[p_{i,j}, \tilde{p}_{i,j}] = \alpha_{i,j}^\top [p_{i,j}, \tilde{p}_{i,j}] \quad (5)$$

k is a split point from i to $j - 1$, $f(\cdot)$ is an n -layer Transformer encoder, $p_{i,j}^k$ and $\tilde{p}_{i,j}^k$ denote the single step combination probability and the subtree probability, respectively, at split point k , $\mathbf{p}_{i,j}$ and $\tilde{\mathbf{p}}_{i,j}$ are the concatenation of all $p_{i,j}^k$ or $\tilde{p}_{i,j}^k$ values, and GUMBEL is the Straight-Through Gumbel-Softmax operation of Jang et al. (2017) with temperature set to one. As GUMBEL picks the optimal splitting point k at each cell in practice, it is straightforward to recover the complete derivation tree from the root node $\mathcal{T}_{1,n}$ in a top-down manner recursively.




Figure 1: Chart data structure. There are two alternative ways of generating $\mathcal{T}_{1,3}$: combining either $(\mathcal{T}_{1,2}, \mathcal{T}_{3,3})$ or $(\mathcal{T}_{1,1}, \mathcal{T}_{2,3})$.

Heuristic pruning. As shown in Figure 2, R2D2 starts to prune if all cells beneath height m have been encoded. The heuristic rules work as follows:

1. Recover the maximum sub-tree for each cell at the m -th level, and collect all cells at the 2nd level that appear in any sub-tree.




Figure 2: Example of chart pruning and encoding process. With R2D2’s original heuristic pruning, cells to merge are selected according to local composition probabilities. For better model-based pruning, we propose selecting cells according to the merge order estimated by a top-down parser.

2. Rank candidates in Step 1 by the composition probability $p_{i,j}$, and pick the highest-scoring cell as a non-splittable span (e.g., $\mathcal{T}_{1,2}$).
3. Remove any invalid cells that would break the now non-splittable span from Step 2, e.g., the dark cells in (c), and reorganize the chart table much like in the Tetris game as in (d).
4. Encode the blank cells at the m -th level, e.g., the cell highlighted with stripes in (d), and go back to Step 1 until the root cell has been encoded.

Pretraining. To learn meaningful structures without gold trees, Hu et al. (2021) propose a self-supervised pretraining objective. Similar to the bidirectional masked language model task, R2D2 reconstructs a given token s_i based on its context representation $e_{1:i-1}$ and $e_{i+1:n}$. The probability of each token is estimated by the tree encoder defined in R2D2. The final objective is:

$$\min_{\theta} \sum_{i=1}^n -\log p_{\theta}(s_i | e_{1:i-1}, e_{i+1:n}) \quad (6)$$

3 Methodology

3.1 Global Pruning Strategy

We propose a top-down parser based on syntactic distance (Shen et al., 2018) to evaluate scores for all split points in a sentence and generate a merge order according to the scores.

Top-down parser. Given a sentence $\mathbf{S} = \{s_1, s_2, \dots, s_n\}$, there are $n-1$ split points between words. We define a top-down parser by giving confidence scores to all split points as follows:

$$\mathbf{v} = [v_1, v_2, \dots, v_{n-1}] = f(\mathbf{S}; \theta) \quad (7)$$

To keep it simple and rigorously maintain linear complexity, we select bidirectional LSTMs as the backbone, though Transformers are also an option. As shown in Figure 3, first, a bi-directional LSTM

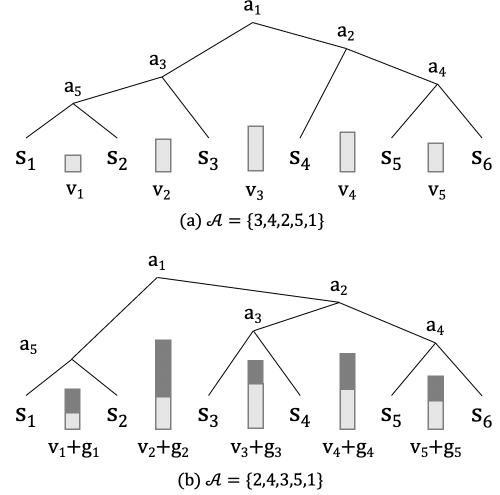


Figure 3: (a) A parsed tree obtained by sorting split scores (v_i). (b) A sampled tree by adding Gumbel noise (g_i in dark vertical bars).

encodes the sentence, and then, for each split point, an MLP over the concatenation of the left and right context representations yields the final split scores. Formally, we have:

$$\begin{aligned} \vec{\mathbf{h}}, \overleftarrow{\mathbf{h}} &= \text{BiLSTM}(\mathbf{E}; \theta) \\ v_i &= \text{LayerNorm}(\text{MLP}(\vec{\mathbf{h}}_i \oplus \overleftarrow{\mathbf{h}}_{i+1})) \end{aligned} \quad (8)$$

Here, \mathbf{E} is the embedding of the input sentence \mathbf{S} , while $\vec{\mathbf{h}}$ and $\overleftarrow{\mathbf{h}}$ denote the forward and reverse representation, respectively. v_i is the score of the i -th split point, whose left and right context representations are $\vec{\mathbf{h}}_i$ and $\overleftarrow{\mathbf{h}}_{i+1}$. Given scores $[v_1, v_2, \dots, v_{n-1}]$, one can easily recover the binary tree shown in Figure 3: We recursively split a span (starting with the entire sentence) into two subspans by picking the split point with the highest score in the current span. Taking the sentence in Figure 3 (a) as an example, we split the overall sentence at split point 3 in the first step, which leads to two sub-trees over $s_{1:3}$ and $s_{4:6}$. Then we split $s_{1:3}$

at 2 and $s_{4,6}$ at 4. We can continue this procedure until the complete tree has been recovered.

Tree sampling. In the training stage, we perform sampling over the computed scores $[v_1, v_2, \dots, v_{n-1}]$ in order to increase the robustness and exploration of our model. Let \mathcal{P}^t denote the list of split points at time t in ascending order, which is $\{1, 2, 3, \dots, n-1\}$ in the first step. Then a particular split point a_t is selected from \mathcal{P}^t by sampling based on the probabilities estimated by stacking of split points scores. The sampled $\{a_1, a_2, \dots, a_{n-1}\}$ together form the final split point sequence \mathcal{A} . At each time step, we remove a_t from \mathcal{P}^t when a_t is selected, then sample the next split point until the set of remaining split points is empty. Formally, we have:

$$a_t \sim \text{softmax}(\mathbf{v}^t) \quad (9)$$

$$\mathcal{P}^{t+1} = \mathcal{P}^t \setminus \{a_t\} \quad (10)$$

where \mathbf{v}^t is concatenation of v_i in \mathcal{P}^t . As the Gumbel-Max trick (E.J.Gumbel, 1954; Maddison et al., 2014) provides a simple and efficient way to draw samples from a categorical distribution with class probabilities, we can obtain a_t via the Gumbel-Max trick as:

$$a_t = \underset{i}{\operatorname{argmax}} [v_i + g_i], i \in \mathcal{P}^t, \quad (11)$$

where g_i is the Gumbel noise for the i -th split point. Therefore, the aforementioned process is equivalent to sorting the original sequence of split points scores with added Gumbel noise. Figure 3 (b) shows a sampled tree with respect to the split point scores. The split point sequence \mathcal{A} can hence be obtained simply as:

$$\mathcal{A} = \underset{i}{\operatorname{argsort}}(\mathbf{v} + \mathbf{g}) \quad (12)$$

Here, $\operatorname{argsort}$ sorts the array in descending order and returns the indices of the original array. The sampled \mathcal{A} is $\{2, 4, 3, 5, 1\}$ in Figure 3 (b).

Span Constraints. As word-pieces (Wu et al., 2016) and Byte-Pair Encoding (BPE) are commonly used in pretrained language models, it is straightforward to incorporate multiple word-piece constraints into the top-down parser to reduce word-level parsing errors. We denote a list of span constraints composed of beginning and end positions of non-split-table spans as \mathcal{C} , defined as $\mathcal{C} = \{(b_1, e_1), (b_2, e_2), \dots, (b_n, e_n)\}$. For each

(b_i, e_i) in \mathcal{C} , there should be a sub-tree for a span covering the sub-string $s_{b_i:e_i}$. This goal can be achieved by simply adjusting the scores of all splits within the spans in \mathcal{C} by $-\delta$. To make them smaller than the scores of span boundaries, δ could be defined as $(\max(\mathbf{v}) - \min(\mathbf{v}) + c)$, where c could be any positive number.

Model-based Pruning. We denote the reverse order of the split point sequence \mathcal{A} as \mathcal{M} and then treat \mathcal{M} as a bottom-up merge order inferred by the top-down parser based on the global context. Subsequently, the complete pruning process is as follows:

1. Pick the next merge index by invoking Alg 1.
2. Perform Steps 3 and 4 in the heuristic pruning part in Section 2.1

As shown in Figure 2, we still retain the threshold and the pruning logic of R2D2, but we select cells to merge according to \mathcal{M} instead of following heuristic rules. Specifically, given a shrinking chart table, we select the next merge index among the second row by popping and modifying \mathcal{M} in Algorithm 1.

Algorithm 1 Next merge index in the second row

```

1: function NEXT-INDEX( $\mathcal{M}$ )
2:    $i = \text{pop}(\mathcal{M})$  ▷ Index
3:   for  $j \in 1$  to  $\mathcal{M}.\text{len}$  do
4:     if  $\mathcal{M}_j > i$  then ▷ Merging at left
5:        $\mathcal{M}_j = \mathcal{M}_j - 1$  ▷ Shift left
6:   return  $i$ 

```

Take the example in Figure 3 (b) for instance: \mathcal{M} starts with $\{1, 5, 3, 4, 2\}$. Then we merge the first cell in the second row in Figure 2 (b), and obtain a new $\mathcal{M} = \{4, 2, 3, 1\}$. In the next round, we treat the 4th cell covering $s_{5,6}$ as a non-splittable cell in Figure 2 (e), and \mathcal{M} becomes $\{2, 3, 1\}$.

3.2 Optimization

We denote the tree probabilities estimated by the top-down parser and R2D2 as $p_\theta(\mathbf{z}|\mathbf{S})$, $q_\phi(\mathbf{z}|\mathbf{S})$, respectively. The difficulty here is that while $q_\phi(\mathbf{z}|\mathbf{S})$ may be optimized by the objective defined in Equation 6, there is no gradient feedback for $p_\theta(\mathbf{z}|\mathbf{S})$. To make $p_\theta(\mathbf{z}|\mathbf{S})$ learnable, an intuitive solution is to fit $p_\theta(\mathbf{z}|\mathbf{S})$ to $q_\phi(\mathbf{z}|\mathbf{S})$ by minimizing their Kullback–Leibler distance. While the tree probabilities of both distributions are discrete and not exhaustive, inspired by URNNG (Kim et al., 2019b), a Monte Carlo estimate for the gradient with respect

to θ can be defined as:

$$\begin{aligned} & \nabla_{\theta} \text{KL}[q_{\phi}(\mathbf{z}|\mathbf{S}) \parallel p_{\theta}(\mathbf{z}|\mathbf{S})] \\ &= \nabla_{\theta} \mathbf{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{S})} \left[\log \frac{q_{\phi}(\mathbf{z}|\mathbf{S})}{p_{\theta}(\mathbf{z}|\mathbf{S})} \right] \\ &\approx - \nabla_{\theta} \frac{1}{K} \sum_{k=1}^K \log p_{\theta}(\mathbf{z}^{(k)}|\mathbf{S}) \end{aligned} \quad (13)$$

with samples $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(K)}$ from $q_{\phi}(\mathbf{z}|\mathbf{S})$. Algorithm 2 shows the complete sampling process from $q_{\phi}(\mathbf{z}|\mathbf{S})$. Specifically, we sample split points with corresponding span boundaries recursively as in previous work (Goodman, 1998; Finkel et al., 2006; Kim et al., 2019b) with respect to the intermediate tree probabilities calculated during hierarchical encoding.

Algorithm 2 Top-down tree sampling for R2D2

```

1: function SAMPLE( $\mathcal{T}_{1,n}$ ) ▷ Root cell
2:    $Q = [\mathcal{T}_{1,n}]$ 
3:    $K = []$ 
4:   while  $Q$  is not empty do
5:      $\mathcal{T} = \text{pop}(Q)$ 
6:      $i, j = \mathcal{T}.i, \mathcal{T}.j$  ▷ Start/end indices
7:      $L = \mathcal{T}.\text{splits}$  ▷  $m$  splits at most
8:      $\tau = 0$ 
9:     for  $k \in 1$  to  $\text{len}(L)$  do
10:       $w_k = \tilde{p}_{i,j}^{L[k]}$  ▷ Using Equation 2
11:       $\tau = \tau + w_k$  ▷ Sum up all  $w_k$ 
12:       $idx \sim \text{Cat}([w_1/\tau, \dots, w_{\text{len}(L)}/\tau])$ 
13:      ▷ Sample a split point
14:      push( $K, (L[idx], i, j)$ )
15:      ▷ Keep the split point and span boundary
16:      if  $L[idx] > i$  then ▷ Add left child
17:        push( $Q, \mathcal{T}_{i, L[idx]}$ )
18:      if  $L[idx] + 1 < j$  then ▷ Add right child
19:        push( $Q, \mathcal{T}_{L[idx]+1, j}$ )
20:   return  $K$ 

```

A sequence of split points and corresponding spans is returned by the sampler. For the k -th sample $\mathbf{z}^{(k)}$, let $p_{\theta}(a_t^k|\mathbf{S})$ denote the probability of taking a_t^k as split from span (i_t^k, j_t^k) at the t -th step. Formally, we have:

$$\begin{aligned} p_{\theta}(a_t^k|\mathbf{S}) &= \frac{e^{v_{a_t^k}}}{e^{v_{i_t^k}} + \dots + e^{v_{j_t^k}}} \\ \log p_{\theta}(\mathbf{z}^{(k)}|\mathbf{S}) &= \sum_{t=1}^{n-1} \log p_{\theta}(a_t^k|\mathbf{S}), \end{aligned} \quad (14)$$

where i_t^k and j_t^k denote the start and end of the corresponding span. Please note here that the v_i are not adjusted by span constraints.

3.3 Downstream Tasks

Inference. In this paper, we mainly focus on classification tasks as downstream tasks. We consider the root representation as representing the entire

sentence. As we have two models pre-trained in our framework – an R2D2 encoder and a top-down parser – we have two ways of generating the representations:

- a) Run forced encoding over the binary tree from the top-down parser with the R2D2 encoder.
- b) Use the binary tree to guide the pruning of the R2D2 encoder, and take the root representation $e_{1,n}$.

It is obvious that the first approach is much faster than the latter one, as the R2D2 encoder only runs $n - 1$ times in forced encoding, and can run in parallel layer by layer, e.g., we may run compositions at a_5, a_3 , and a_4 in parallel in Figure 3 (b). We explore both of these approaches in our experiments.

Training Objectives. As suggested in prior work (Radford et al., 2018; Howard and Ruder, 2018; Gururangan et al., 2020), given a pretrained model, continued pretraining on an in-domain corpus with the same pretraining objective can yield a better generalization ability. Thus, we simply combine our pretraining objectives via summation in all downstream tasks. At the same time, as the downstream task may guide R2D2 to more reasonable tree structures, we still maintain the KL loss to enable the parser to continuously update. For the two inference methods, we uniformly select the root representation $e_{1,n}$ as the representation for a given sentence followed by an MLP, and estimate the cross-entropy loss, denoted as $\mathcal{L}_{\text{forced}}$ and \mathcal{L}_{cky} , respectively. Let \mathcal{L}_{KL} denote the KL loss described in Section 3.2 and $\mathcal{L}_{\text{bilm}}$ denote the bidirectional language model loss described in Eq 6. The final loss is:

$$\mathcal{L} = \mathcal{L}_{\text{forced}} + \mathcal{L}_{\text{cky}} + \mathcal{L}_{\text{bilm}} + \mathcal{L}_{\text{KL}} \quad (15)$$

4 Experiments

4.1 Unsupervised Grammar Induction

4.1.1 Setup

Baselines and Evaluation. For comparison, we include six recent strong models for unsupervised parsing with available open source implementations: StructFormer (Shen et al., 2021), Ordered Neurons (Shen et al., 2019b), URNNG (Kim et al., 2019b), DIORA (Drozdov et al., 2019), C-PCFG (Kim et al., 2019a), and R2D2 (Hu et al., 2021). To observe the marginal gain from pretraining, we also include Fast-R2D2 without pretraining

denoted as Fast-R2D2_{w/o}. Following Htut et al. (2018), we train all systems on a training set consisting only of raw text, and evaluate and report the results on an annotated test set. As an evaluation metric, we adopt sentence-level unlabeled F_1 computed using the script from Kim et al. (2019a). We compare against the non-binarized gold trees per convention. The results of Fast-R2D2 are obtained from 3 runs of each model with different random seeds in pre-training. The best checkpoint for each system is picked based on scores on the validation set. Fast-R2D2 is pretrained with span constraints for the word level but without span constraints for the word-piece level. To support word-piece level evaluation, we convert gold trees to word-piece level trees by simply breaking each terminal node into a non-terminal node with its word-pieces as terminals, e.g., (NN discrepancy) into (NN (WP disc) (WP ##re) (WP ##pan) (WP ##cy)).

Environment. EFLOPS (Dong et al., 2020) is a highly scalable distributed training system designed by Alibaba. With its optimized hardware architecture and co-designed supporting software tools, including ACCL (Dong et al., 2021) and KSpeed (the high-speed data-loading service), it could easily be extended to 10K nodes (GPUs) with linear scalability.

Hyperparameters. The tree encoder of our model uses 4-layer Transformers with 768-dimensional embeddings, 3,072-dimensional hidden layer representations, and 12 attention heads. The top-down parser of our model uses a 4-layer bidirectional LSTM with 128-dimensional embeddings and 256-dimensional hidden layer. The sampling number K is set to be 256. Training is conducted using Adam optimization with weight decay using a learning rate of 5×10^{-5} for the tree encoder and 1×10^{-2} for the top-down parser. The batch size is set to 64 per GPU for $m=4$, though we also limit the maximum total length for each batch, such that excess sentences are moved to the next batch. The limit is set to 1,536. It takes about 120 hours for 60 epochs of training with $m=4$ on 8 A100 GPUs.

Data. For English, to fully leverage the scalability of Fast-R2D2, we pretrain Fast-R2D2 on WikiText103 (Merity et al., 2017) and then fine-tune the model on the Penn Treebank (PTB) (Marcus et al., 1993) for 10 epochs with the same objective. WikiText103 is split at the sentence level, and

Model	eval gran.	mem. cplx	WSJ $F_1(\mu)$	CTB $F_1(\mu)$
Left Branching (W)	WD	$O(n)$	8.15	11.28
Right Branching (W)	WD	$O(n)$	39.62	27.53
Random Trees (W)	WD	$O(n)$	17.76	20.17
URNNG (W)	WD	$O(n^3)$	45.4 [†]	—
ON-LSTM (W)	WD	$O(n)$	47.7 [†]	24.73
DIORA (W)	WD	$O(n^3)$	51.4	—
StructFormer (W)	WD	$O(n^2)$	54.0 [‡]	—
C-PCFG (W)	WD	$O(n^3)$	55.2 [†]	49.95
R2D2 (WP)	WD	$O(n)$	48.11	44.85
Fast-R2D2*(W) _{w/o}	WD	$O(n)$	48.24	45.24
Fast-R2D2*(WP) _{w/o}	WD	$O(n)$	48.89	45.26
Fast-R2D2*(WP)	WD	$O(n)$	57.22	53.13
R2D2 (WP)	WP	$O(n)$	52.28	63.94
Fast-R2D2(WP)	WP	$O(n)$	50.20	67.79
Fast-R2D2*(WP)	WP	$O(n)$	53.88	67.74

Table 1: Unsupervised parsing results with words (W) or word-pieces (WP) as input. “eval gran.” is short for evaluation granularity. Values marked with [†] are taken from Kim et al. (2019a), while [‡] denotes values taken from Shen et al. (2021). The bottom three systems are all pre-trained or trained at the word-piece level **without** span constraints and are measured against word-piece level golden trees. w/o means without pretraining.

sentences longer than 200 after tokenization are discarded (about 0.04% of the original data). The total number of sentences is 4,089,500, and the average sentence length is 26.97. For Chinese, we use a subset of Chinese Wikipedia (Simplified Characters) for pretraining, specifically the first 10,000,000 sentences shorter than 150 characters and then fine-tune on Chinese Penn Treebank (CTB) 8.0 (Xue et al., 2005). We test our approach on PTB WSJ data with the standard splits (2–21 for training, 22 for validation, 23 for test) and the same preprocessing as in recent work (Kim et al., 2019a), where we discard punctuation and lower-case all tokens. To explore the universality of the model across languages, we further evaluate using the CTB, on which we also remove punctuation. Note that in all settings, the training and fine-tuning is conducted entirely on raw unannotated text.

4.1.2 Results and Discussion

Table 1 shows the results of all systems with words (W) and word-pieces (WP) as input on the WSJ and CTB test sets. When we evaluate all systems on word-level golden trees, our Fast-R2D2 performs

substantially better than R2D2 across both datasets. We denote as Fast-R2D2 the method of using the parser to guide the pruning and selecting the best tree using the chart table and as Fast-R2D2* the system that uses the top-down parser for tree induction with subsequent R2D2 encoding. Interestingly, the results suggest that Fast-R2D2* outperforms Fast-R2D2, especially on the WSJ test set. Additionally, pretrained Fast-R2D2* outperforms the models specifically designed for grammar induction.

	Model	WD	NNP	VP	SBAR
WSJ	DIORA (WP)	94.63	77.83	17.30	22.16
	C-PCFG (W)	—	—	41.7 [†]	56.1 [†]
	C-PCFG (WP)	87.35	66.44	23.63	40.40
	R2D2 (WP)	99.76	86.76	24.74	39.81
	Fast-R2D2* (WP)	97.67	83.44	63.80	65.68
CTB	C-PCFG(WP)	89.34	46.74	39.53	—
	R2D2 (WP)	97.16	67.19	37.90	—
	Fast-R2D2* (WP)	97.80	68.57	46.59	—

Table 2: Recall of constituents and words. WD means word. Values with [†] are taken from Kim et al. (2019a).

Following Kim et al. (2019b) and Drozdov et al. (2020), we also compute the recall of constituents when evaluating on word-piece level golden trees. Besides standard constituents, we also compare the recall of word-piece chunks and proper noun chunks. Proper noun chunks are extracted by finding adjacent unary nodes with the same parent and tag NNP. Table 2 reports the recall scores for constituents and words on the WSJ and CTB test sets. Compared with the R2D2 baseline, our Fast-R2D2 performs slightly worse for small semantic units, but significantly better over larger semantic units (such as VP and SBAR) on the WSJ test set. On the CTB test set, our Fast-R2D2 outperforms R2D2 on all constituents.

From Tables 1 and 2, we conclude that Fast-R2D2 overall obtains better results than R2D2 on CTB, while faring slightly worse than R2D2 only for small semantic units on WSJ. We conjecture that this difference stems from differences in tokenization between Chinese and English. Chinese is a character-based language without complex morphology, where collocations of characters are consistent with the language, making it easier for the top-down parser to learn them well. In contrast, word-pieces for English are built based on statistics, and individual word-pieces are not necessarily

natural semantic units. Thus, there may not be sufficient semantic self-consistency, such that it is harder for a top-down parser with a small number of parameters to fit it well.

4.2 Downstream Tasks

We next consider the effectiveness of Fast-R2D2 in downstream tasks. This experiment is not intended to advance the state-of-the-art on the GLUE benchmark but rather to assess to what extent our approach performs respectably against the dominant inductive bias as in conventional sequential Transformers.

4.2.1 Setup

Data and Baseline. We fine-tune pretrained models on several datasets, including SST-2, CoLA, QQP, and MNLI from the GLUE benchmark (Wang et al., 2018). As sequential Transformers with their dominant inductive bias remain the norm for numerous NLP tasks, we mainly compare Fast-R2D2 with BERT (Devlin et al., 2018) as a representative pretrained model based on a sequential Transformer. We did not include recursive models such as Gumbel-Tree-LSTMs (Choi et al., 2018) and CRvNN (Chowdhury and Caragea, 2021) among our baselines, as they are not pretrained models. In order to compare the two forms of inductive bias fairly and efficiently, we pretrain BERT models with 4 layers and 12 layers as well as our Fast-R2D2 from scratch on the WikiText103 corpus following Section 4.1.1. Considering that longer inputs in the pre-training stage are helpful for BERT’s downstream task performance, we use the original corpus that is not split into sentences as inputs. For simplicity, Fast-R2D2 is fine-tuned without span constraints. Following the common settings, we add an MLP layer over the root representation of the R2D2 encoder for single-sentence classification. For cross-sentence tasks such as QQP and MNLI, we feed the root representations of the two sentences into the pretrained tree encoder of R2D2 as left and right inputs, and also add a new task ID as another input term to the R2D2 encoder. Then we feed the hidden output of the new task ID into another MLP layer to predict the final label. We train all systems across the four datasets for 10 epochs with a learning rate of 5×10^{-5} , batch size 64, and maximum input length 200. We validate each model in each epoch and report the best results on development sets.

Model	Para.	Single sent.		Cross sent.	
		SST-2 (Acc.)	CoLA (McC.)	QQP (F1)	MNLI m/mm (Acc.)
BERT (4L)	52M	84.98	17.07	84.01	73.73/74.63
BERT (12L)	116M	90.25	40.72	87.13	80.00/80.41
R2D2	52M	89.33	34.79	84.27	69.35/68.72
Fast-R2D2 [†]		87.50	8.67	83.97	69.53/69.50
Fast-R2D2* [†]	52M/	88.30	10.14	84.07	69.36/69.11
Fast-R2D2	10M	90.25	38.45	84.35	69.36/68.80
Fast-R2D2*		90.71	40.11	84.32	69.64/69.57

Table 3: Downstream results. All systems are pretrained from scratch on WikiText103. Para. describes the number of parameters for each model. Fast-R2D2 contains the R2D2 encoder and top-down parser, two components with 52M and 10M parameters, respectively. Mcc. stands for Matthew’s correlation coefficient. Fast-R2D2 with [†] are models fine-tuned without $\mathcal{L}_{\text{bilm}}$ for an ablation study.

4.2.2 Results and Discussion

Table 3 shows the corresponding scores on SST-2, CoLA, QQPI, and MNLI. In terms of the parameter size, our Fast-R2D2 model has 52M and 10M parameters for the R2D2 encoder and top-down parser, respectively. It is clear that 12-layer BERT is significantly better than 4-layer BERT. As mentioned in Section 3.3, Fast-R2D2 has two options to construct the final tree and representation for a given input sentence: Fast-R2D2* uses the output tree from the top-down parser, while Fast-R2D2 uses the best tree inferred by the R2D2 encoder. Similar to the results for unsupervised parsing, Fast-R2D2* in classification tasks again outperforms Fast-R2D2. We hypothesize that trees generated by the top-down parser without Gumbel noise are more stable and reasonable. Fast-R2D2 significantly outperforms 4-layer BERT and achieves competitive results compared to 12-layer BERT in single sentence classification tasks such as SST-2 and CoLA, but still performs significantly worse in the cross-sentence tasks. We believe this is an expected result, as there is no cross-attention mechanism in the inductive bias of Fast-R2D2. However, the performance of Fast-R2D2 on classification tasks shows that the inductive bias of R2D2 has higher parameter utilization than sequentially applied Transformers. Importantly, we demonstrate that a Recursive Neural Network variant with an unsupervised parser can achieve comparable results

Model	Sequence Length Ranges			
	0–50	50–100	100–200	200–500
BERT (12L)	1.36	1.46	1.62	2.38
R2D2	38.06	173.74	555.95	—
Fast-R2D2	4.67	14.91	39.73	150.26
Fast-R2D2*	1.28	2.96	5.56	10.70

Table 4: Inference time in seconds for various systems to process 1,000 sentences with a batch size of 50.

to pretrained sequential Transformers even with fewer parameters and interpretable intermediate results. Hence, our Fast-R2D2 framework provides an alternative for NLP tasks.

4.3 Speed Evaluation

To assess the time cost, we mainly compare sequential Transformers and Fast-R2D2 in forced encoding on various sequence length ranges. We randomly select 1,000 sentences for each range from WikiText103 and report the average time consumption on a single A100 GPU. BERT is based on the open source Transformers library² and R2D2 is based on the official code in Hu et al. (2021).³

Table 4 shows the inference time in seconds for different systems to process 1,000 sentences with a batch size of 50. Running R2D2 is time-consuming, since the heuristic pruning method involves substantial memory exchanges between GPU and CPU. In Fast-R2D2, we alleviate this problem by using model-guided pruning to accelerate the chart table processing, in conjunction with a code implementation in CUDA, Fast-R2D2 reduces the inference time significantly. Fast-R2D2* further improves the inference speed by running forced encoding in parallel over the binary tree generated by the parser, which is about 30–50 times faster than R2D2 in various ranges. Although there is still a gap in speed compared to sequential Transformers, Fast-R2D2* is sufficiently fast for most NLP tasks while producing interpretable intermediate representations.

5 Related Work

Many attempts have been done in grammar induction and hierarchical encoding. Clark (2001) and Klein and Manning (2002) provided some of the first successful statistical approaches to grammar induction. There have been multiple recent papers

²<https://github.com/huggingface/transformers>

³https://github.com/alipay/StructuredLM_RTDT/tree/r2d2

that focus on structure induction based on language modeling objectives (Shen et al., 2019a,b, 2021; Kim et al., 2019a). Pollack (1990) proposed to use RvNN as a recursive architecture to encode text hierarchically, and Socher et al. (2013) showed the effectiveness of RvNNs with gold trees for sentiment analysis. In this work, we focus on models that are capable of learning meaningful structures in an unsupervised way and encoding text over the induced tree hierarchically.

In the line of work on learning a sentence representation with structures, Yogatama et al. (2017) jointly train their shift-reduce parser and sentence embedding components without gold trees. As their parser is not differentiable, they have to resort to reinforcement training, resulting in increased variance, which may easily collapse to trivial left or right branching trees. Gumbel-Tree-LSTMs (Choi et al., 2018) construct trees by recursively selecting two terminal nodes to merge and learning the composition probability via downstream tasks. CRvNN (Chowdhury and Caragea, 2021) makes the entire process end-to-end differentiable and parallel by introducing a continuous relaxation. URNNG (Kim et al., 2019b) propose the first architecture to jointly pretrain parser and encoder based on RNNNG (Dyer et al., 2016). However, it has $O(n^3)$ complexity and remains unable to improve upon a right-branching baseline when punctuation is removed. Maillard et al. (2017) propose an alternative approach, based on a differentiable CKY encoding. The algorithm is differentiable due to a soft-gating approach, which approximates discrete candidate selection by a probabilistic mixture of the constituents available in a given cell of the chart. While their work relies on annotated downstream tasks to learn structures, Drozdov et al. (2019) propose a novel auto-encoder-like pretraining objective based on the inside–outside algorithm (Baker, 1979; Casacuberta, 1994).

6 Conclusion

In this paper, we have presented Fast-R2D2, which improves the performance and inference speed of R2D2 by introducing a fast top-down parser to guide the pruning of the R2D2 encoder. Pretrained on the same corpus, Fast-R2D2 significantly outperforms sequential Transformers with a similar scale of parameters on classification tasks. Experimental results show that Fast-R2D2 is a promising and feasible way to learn hierarchical text representations, which is different from layer stacking

models and can also generate interpretable intermediate representations. As future work, we are investigating leveraging the intermediate representations in additional downstream tasks.

7 Limitations

Our approach has three major limitations. First, Fast-R2D2 has shortcomings with regard to cross-sentence tasks due to the lack of cross-attention between sentences. Second, Fast-R2D2 requires greater memory resources for pretraining compared to sequential Transformers. At each invocation, the composition function takes four inputs and runs on m candidates, which means the total number of calls to the MLP is $4mn$. Hence, the pretraining time of Fast-R2D2 is about 3 to 4 times that of BERT with 12 layers. Finally, our model does not beat most of the baselines in grammar induction when trained on WSJ only. A side effect of the pruning strategy is that the chart-table actually is a sparse table, which means not all tokens are reconstructed based on complete context. This issue can be alleviated by pre-training on a large corpus, which is what our method is designed for, and why we introduce the ability to parallelize the computation.

8 Acknowledgement

We would like to thank the Aliyun EFLOPS team for their substantial support in designing and providing a cutting-edge training platform to facilitate fast experimentation in this work. We would also like to thank the Zhixiaobao team for their support in applying our model to real applications.

References

- James K. Baker. 1979. Trainable grammars for speech recognition. *Journal of the Acoustical Society of America*, 65.
- Francisco Casacuberta. 1994. Statistical estimation of stochastic context-free grammars using the inside-outside algorithm and a transformation on grammars. In *Grammatical Inference and Applications, Second International Colloquium, ICGI-94, Alicante, Spain, September 21-23, 1994, Proceedings*, volume 862 of *Lecture Notes in Computer Science*, pages 119–129. Springer.
- Jihun Choi, Kang Min Yoo, and Sang-goo Lee. 2018. Learning to compose task-specific tree structures. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5094–5101. AAAI Press.
- Noam Chomsky. 1956. Three models for the description of language. *IRE Trans. Inf. Theory*, 2(3):113–124.
- Noam Chomsky. 2014. *Aspects of the Theory of Syntax*, volume 11. MIT press.
- Jishnu Ray Chowdhury and Cornelia Caragea. 2021. Modeling hierarchical structures with continuous recursive neural networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 1975–1988. PMLR.
- Alexander Clark. 2001. Unsupervised induction of stochastic context-free grammars using distributional clustering. In *Proceedings of the ACL 2001 Workshop on Computational Natural Language Learning, CoNLL 2001, Toulouse, France, July 6-7, 2001*. ACL.
- John Cocke. 1969. *Programming Languages and Their Compilers: Preliminary Notes*. New York University, USA.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jianbo Dong, Zheng Cao, Tao Zhang, Jianxi Ye, Shaochuang Wang, Fei Feng, Li Zhao, Xiaoyong Liu, Liuyihan Song, Liwei Peng, Yiqun Guo, Xiaowei Jiang, Lingbo Tang, Yin Du, Yingya Zhang, Pan Pan, and Yuan Xie. 2020. EFLOPS: algorithm and system co-design for a high performance distributed training platform. In *IEEE International Symposium on High Performance Computer Architecture, HPCA 2020, San Diego, CA, USA, February 22-26, 2020*, pages 610–622. IEEE.
- Jianbo Dong, Shaochuang Wang, Fei Feng, Zheng Cao, Heng Pan, Lingbo Tang, Pengcheng Li, Hao Li, Qianyan Ran, Yiqun Guo, Shanyuan Gao, Xin Long, Jie Zhang, Yong Li, Zhisheng Xia, Liuyihan Song, Yingya Zhang, Pan Pan, Guohui Wang, and Xiaowei Jiang. 2021. ACCL: architecting highly scalable distributed training systems with highly efficient collective communication library. *IEEE Micro*, 41(5):85–92.
- Andrew Drozdov, Subendhu Rongali, Yi-Pei Chen, Tim O’Gorman, Mohit Iyyer, and Andrew McCallum. 2020. Unsupervised parsing with S-DIORA: Single tree encoding for deep inside-outside recursive autoencoders. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4832–4845, Online. Association for Computational Linguistics.
- Andrew Drozdov, Patrick Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised latent tree induction with deep inside-outside recursive auto-encoders. 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 1129–1141, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent neural network grammars. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.
- E.J.Gumbel. 1954. Statistical theory of extreme values and some practical applications: a series of lectures.
- Jenny Rose Finkel, Christopher D. Manning, and Andrew Y. Ng. 2006. Solving the problem of cascading errors: Approximate bayesian inference for linguistic annotation pipelines. In *EMNLP 2006, Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing, 22-23 July 2006, Sydney, Australia*, pages 618–626. ACL.
- Joshua Goodman. 1998. Parsing inside-out. *CoRR*, cmp-lg/9805007.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.

- Phu Mon Htut, Kyunghyun Cho, and Samuel Bowman. 2018. [Grammar induction with neural language models: An unusual replication](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4998–5003, Brussels, Belgium. Association for Computational Linguistics.
- Xiang Hu, Haitao Mi, Zujie Wen, Yafang Wang, Yi Su, Jing Zheng, and Gerard de Melo. 2021. [R2D2: Recursive transformer based on differentiable tree for interpretable hierarchical language modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4897–4908, Online. Association for Computational Linguistics.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Tadao Kasami. 1966. An efficient recognition and syntax-analysis algorithm for context-free languages. *Coordinated Science Laboratory Report no. R-257*.
- Yoon Kim, Chris Dyer, and Alexander Rush. 2019a. [Compound probabilistic context-free grammars for grammar induction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2369–2385, Florence, Italy. Association for Computational Linguistics.
- Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019b. [Unsupervised recurrent neural network grammars](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1105–1117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dan Klein and Christopher D. Manning. 2002. [A generative constituent-context model for improved grammar induction](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 128–135. ACL.
- Chris J. Maddison, Daniel Tarlow, and Tom Minka. 2014. [A* sampling](#). In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3086–3094.
- Jean Maillard, Stephen Clark, and Dani Yogatama. 2017. [Jointly learning sentence embeddings and syntax with unsupervised tree-lstms](#). *CoRR*, abs/1705.09189.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Jordan B. Pollack. 1990. [Recursive distributed representations](#). *Artif. Intell.*, 46(1-2):77–105.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordani, Aaron C. 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, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 1171–1180. Association for Computational Linguistics.
- Yikang Shen, Shawn Tan, Seyed Arian Hosseini, Zhouhan Lin, Alessandro Sordani, and Aaron C. Courville. 2019a. [Ordered memory](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5038–5049.
- Yikang Shen, Shawn Tan, Alessandro Sordani, and Aaron C. Courville. 2019b. [Ordered neurons: Integrating tree structures into recurrent neural networks](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Yikang Shen, Yi Tay, Che Zheng, Dara Bahri, Donald Metzler, and Aaron C. Courville. 2021. [Structformer: Joint unsupervised induction of dependency and constituency structure from masked language modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 7196–7209. Association for Computational Linguistics.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy P. Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. 2017. [Mastering the game of go without human knowledge](#). *Nat.*, 550(7676):354–359.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*,

EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIG-DAT, a Special Interest Group of the ACL, pages 1631–1642. ACL.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.


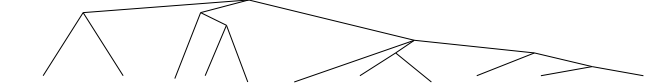
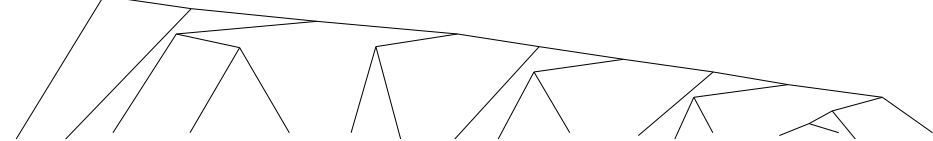

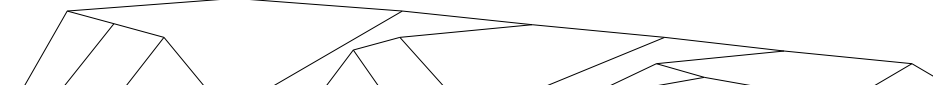

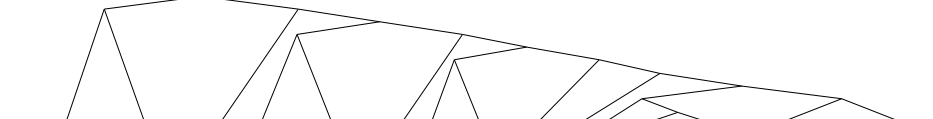
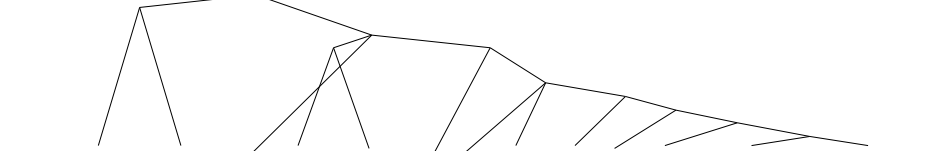
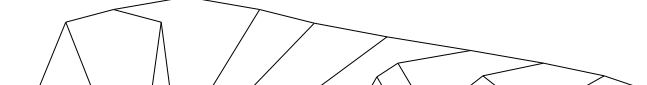

Naiwen Xue, Fei Xia, Fu-dong Chiou, and Marta Palmer. 2005. [The penn chinese treebank: Phrase structure annotation of a large corpus](#). *Nat. Lang. Eng.*, 11(2):207–238.

Dani Yogatama, Phil Blunsom, Chris Dyer, Edward Grefenstette, and Wang Ling. 2017. [Learning to compose words into sentences with reinforcement learning](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.

Daniel H Younger. 1967. Recognition and parsing of context-free languages in time n^3 . *Information and control*, 10(2):189–208.

9 Appendix

9.1 Tree Examples

System	Tree
FAST-R2D2	 <p data-bbox="533 450 1235 483">pricing cycles to be sure are nothing new for plastics producers</p>
GOLD	 <p data-bbox="533 584 1235 618">pricing cycles to be sure are nothing new for plastics producers</p>
FAST-R2D2	 <p data-bbox="384 779 1385 813">we were all wonderful heroes last year says an executive at one of quantum ' s competitors</p>
GOLD	 <p data-bbox="384 936 1385 969">we were all wonderful heroes last year says an executive at one of quantum ' s competitors</p>
FAST-R2D2	 <p data-bbox="384 1061 1385 1095">a quick turn ##around is crucial to quantum because its cash requirements remain heavy</p>
GOLD	 <p data-bbox="384 1211 1385 1245">A quick turnaround is crucial to Quantum because its cash requirements remain heavy</p>
FAST-R2D2	 <p data-bbox="384 1391 1385 1424">some analysts saw the payment as an effort also to di ##sp ##el takeover speculation</p>
GOLD	 <p data-bbox="384 1608 1385 1641">Some analysts saw the payment as an effort also to dispel takeover speculation</p>
FAST-R2D2	 <p data-bbox="533 1749 1235 1783">ford motor co . said it acquired 5 % of the shares in jaguar plc</p>
GOLD	 <p data-bbox="533 1906 1235 1939">Ford Motor Co. said it acquired 5 % of the shares in Jaguar PLC</p>