

End-to-end Graph-based TAG Parsing with Neural Networks

Jungo Kasai[♣] Robert Frank[♣] Pauli Xu[♣]
William Merrill[♣] Owen Rambow[♡]

[♣]Department of Linguistics, Yale University

[♡]Elemental Cognition, LLP

{jungo.kasai,bob.frank,pauli.xu,william.merrill}@yale.edu

owenr@elementalcognition.com

Abstract

We present a graph-based Tree Adjoining Grammar (TAG) parser that uses BiLSTMs, highway connections, and character-level CNNs. Our best end-to-end parser, which jointly performs supertagging, POS tagging, and parsing, outperforms the previously reported best results by more than 2.2 LAS and UAS points. The graph-based parsing architecture allows for global inference and rich feature representations for TAG parsing, alleviating the fundamental trade-off between transition-based and graph-based parsing systems. We also demonstrate that the proposed parser achieves state-of-the-art performance in the downstream tasks of Parsing Evaluation using Textual Entailments (PETE) and Unbounded Dependency Recovery. This provides further support for the claim that TAG is a viable formalism for problems that require rich structural analysis of sentences.

1 Introduction

Tree Adjoining Grammar (TAG, Joshi and Schabes (1997)) and Combinatory Categorical Grammar (CCG, Steedman and Baldridge (2011)) are both mildly context-sensitive grammar formalisms that are lexicalized: every elementary structure (elementary tree for TAG and category for CCG) is associated with exactly one lexical item, and every lexical item of the language is associated with a finite set of elementary structures in the grammar (Rambow and Joshi, 1994). In TAG and CCG, the task of parsing can be decomposed into two phases (e.g. TAG: Bangalore and Joshi (1999); CCG: Clark and Curran (2007)): *supertagging*, where elementary units or *supertags* are assigned to each lexical item and *parsing* where these supertags are combined together. The first phase of supertagging can be considered as “almost parsing” because supertags for a sentence almost always determine a unique parse (Bangalore and

Joshi, 1999). This near uniqueness of a parse given a gold sequence of supertags has been confirmed empirically (TAG: Bangalore et al. (2009); Chung et al. (2016); Kasai et al. (2017); CCG: Lewis et al. (2016)).

We focus on TAG parsing in this work. TAG differs from CCG in having a more varied set of supertags. Concretely, the TAG-annotated version of the WSJ Penn Treebank (Marcus et al., 1993) that we use (Chen et al., 2005) includes 4727 distinct supertags (2165 occur once) while the CCG-annotated version (Hockenmaier and Steedman, 2007) only includes 1286 distinct supertags (439 occur once). This large set of supertags in TAG presents a severe challenge in supertagging and causes a large discrepancy in parsing performance with gold supertags and predicted supertags (Bangalore et al., 2009; Chung et al., 2016; Kasai et al., 2017).

In this work, we present a supertagger and a parser that substantially improve upon previously reported results. We propose crucial modifications to the bidirectional LSTM (BiLSTM) supertagger in Kasai et al. (2017). First, we use character-level Convolutional Neural Networks (CNNs) for encoding morphological information instead of suffix embeddings. Secondly, we perform concatenation after each BiLSTM layer. Lastly, we explore the impact of adding additional BiLSTM layers and highway connections. These techniques yield an increase of 1.3% in accuracy. For parsing, since the derivation tree in a lexicalized TAG is a type of dependency tree (Rambow and Joshi, 1994), we can directly apply dependency parsing models. In particular, we use the biaffine graph-based parser proposed by Dozat and Manning (2017) together with our novel techniques for supertagging.

In addition to these architectural extensions for supertagging and parsing, we also explore multi-task learning approaches for TAG parsing. Specif-

ically, we perform POS tagging, supertagging, and parsing using the same feature representations from the BiLSTMs. This joint modeling has the benefit of avoiding a time-consuming and complicated pipeline process, and instead produces a full syntactic analysis, consisting of supertags and the derivation that combines them, simultaneously. Moreover, this multi-task learning framework further improves performance in all three tasks. We hypothesize that our multi-task learning yields feature representations in the LSTM layers that are more linguistically relevant and that generalize better (Caruana, 1997). We provide support for this hypothesis by analyzing syntactic analogies across induced vector representations of supertags (Kasai et al., 2017; Friedman et al., 2017). The end-to-end TAG parser substantially outperforms the previously reported best results.

Finally, we apply our new parsers to the downstream tasks of Parsing Evaluation using Textual Entailments (PETE, Yuret et al. (2010)) and Unbounded Dependency Recovery (Rimell et al., 2009). We demonstrate that our end-to-end parser outperforms the best results in both tasks. These results illustrate that TAG is a viable formalism for tasks that benefit from the assignment of rich structural descriptions to sentences.

2 Our Models

TAG parsing can be decomposed into *supertagging* and *parsing*. Supertagging assigns to words elementary trees (supertags) chosen from a finite set, and parsing determines how these elementary trees can be combined to form a derivation tree that yield the observed sentence. The combinatory operations consist of *substitution*, which inserts obligatory arguments, and *adjunction*, which is responsible for the introduction of modifiers, function words, as well as the derivation of sentences involving long-distance dependencies. In this section, we present our supertagging models, parsing models, and joint models.

2.1 Supertagging Model

Recent work has explored neural network models for supertagging in TAG (Kasai et al., 2017) and CCG (Xu et al., 2015; Lewis et al., 2016; Vaswani et al., 2016; Xu, 2016), and has shown that such models substantially improve performance beyond non-neural models. We extend previously proposed BiLSTM-based models (Lewis

et al., 2016; Kasai et al., 2017) in three ways: 1) we add character-level Convolutional Neural Networks (CNNs) to the input layer, 2) we perform concatenation of both directions of the LSTM not only after the final layer but also after each layer, and 3) we use a modified BiLSTM with highway connections.

2.1.1 Input Representations

The input for each word is represented via concatenation of a 100-dimensional embedding of the word, a 100-dimensional embedding of a predicted part of speech (POS) tag, and a 30-dimensional character-level representation from CNNs that have been found to capture morphological information (Santos and Zdrozny, 2014; Chiu and Nichols, 2016; Ma and Hovy, 2016). The CNNs encode each character in a word by a 30 dimensional vector and 30 filters produce a 30 dimensional vector for the word. We initialize the word embeddings to be the pre-trained GloVe vectors (Pennington et al., 2014); for words not in GloVe, we initialize their embedding to a zero vector. The other embeddings are randomly initialized. We obtain predicted POS tags from a BiLSTM POS tagger with the same configuration as in Ma and Hovy (2016).

2.1.2 Deep Highway BiLSTM

The core of the supertagging model is a deep bidirectional Long Short-Term Memory network (Graves and Schmidhuber, 2005). We use the following formulas to compute the activation of a single LSTM cell at time step t :

$$i_t = \sigma(W_i[x_t; h_{t-1}] + b_i) \quad (1)$$

$$f_t = \sigma(W_f[x_t; h_{t-1}] + b_f) \quad (2)$$

$$\tilde{c}_t = \tanh(W_c[x_t; h_{t-1}] + b_c) \quad (3)$$

$$o_t = \sigma(W_o[x_t; h_{t-1}] + b_o) \quad (4)$$

$$c_t = f \odot c_{t-1} + i_t \odot \tilde{c}_t \quad (5)$$

$$h_t = o \odot \tanh(c_t) \quad (6)$$

Here a semicolon ; means concatenation, \odot is element-wise multiplication, and σ is the sigmoid function. In the first BiLSTM layer, the input x_t is the vector representation of word t . (The sequence is reversed for the backwards pass.) In all subsequent layers, x_t is the corresponding output from the previous BiLSTM; the output of a BiLSTM at timestep t is equal to $[h_t^f; h_t^b]$, the concatenation of hidden state corresponding to input t in the forward and backward pass. This concatenation af-

ter each layer differs from [Kasai et al. \(2017\)](#) and [Lewis et al. \(2016\)](#), where concatenation happens only after the final BiLSTM layer. We will show in a later section that concatenation after each layer contributes to improvement in performance.

We also extend the models in [Kasai et al. \(2017\)](#) and [Lewis et al. \(2016\)](#) by allowing highway connections between LSTM layers. A highway connection is a gating mechanism that combines the current and previous layer outputs, which can prevent the problem of vanishing/exploding gradients ([Srivastava et al., 2015](#)). Specifically, in networks with highway connections, we replace Eq. 6 by:

$$r_t = \sigma(W_r[x_t; h_{t-1}] + b_r)$$

$$h_t = r_t \odot o_t \odot \tanh(c_t) + (1 - r_t) \odot W_h x_t$$

Indeed, our experiments will show that highway connections play a crucial role as we add more BiLSTM layers.

We generally follow the hyperparameters chosen in [Lewis et al. \(2016\)](#) and [Kasai et al. \(2017\)](#). Specifically, we use BiLSTMs layers with 512 units each. Input, layer-to-layer, and recurrent ([Gal and Ghahramani, 2016](#)) dropout rates are all 0.5. For the CNN character-level representation, we used the hyperparameters from [Ma and Hovy \(2016\)](#).

We train this network, including the embeddings, by optimizing the negative log-likelihood of the observed sequences of supertags in a mini-batch stochastic fashion with the Adam optimization algorithm with batch size 100 and $\ell = 0.01$ ([Kingma and Ba, 2015](#)). In order to obtain predicted POS tags and supertags of the training data for subsequent parser input, we also perform 10-fold jackknife training. After each training epoch, we test the supertagger on the dev set. When classification accuracy does not improve on five consecutive epochs, training ends.

2.2 Parsing Model

Until recently, TAG parsers have been grammar based, requiring as input a set of elementary trees (supertags). For example, [Bangalore et al. \(2009\)](#) proposes the MICA parser, an Earley parser that exploits a TAG grammar that has been transformed into a variant of a probabilistic CFG. One advantage of such a parser is that its parses are guaranteed to be well-formed according to the TAG grammar provided as input.

More recent work, however, has shown that data-driven transition-based parsing systems outperform such grammar-based parsers ([Chung et al., 2016](#); [Kasai et al., 2017](#); [Friedman et al., 2017](#)). [Kasai et al. \(2017\)](#) and [Friedman et al. \(2017\)](#) achieved state-of-the-art TAG parsing performance using an unlexicalized shift-reduce parser with feed-forward neural networks that was trained on a version of the Penn Treebank that had been annotated with TAG derivations. Here, we pursue this data-driven approach, applying a graph-based parser with deep biaffine attention ([Dozat and Manning, 2017](#)) that allows for global training and inference.

2.2.1 Input Representations

The input for each word is the concatenation of a 100-dimensional embedding of the word and a 30-dimensional character-level representation obtained from CNNs in the same fashion as in the supertagger.¹ We also consider adding 100-dimensional embeddings for a predicted POS tag ([Dozat and Manning, 2017](#)) and a predicted supertag ([Kasai et al., 2017](#); [Friedman et al., 2017](#)). The ablation experiments in [Kiperwasser and Goldberg \(2016\)](#) illustrated that adding predicted POS tags boosted performance in Stanford Dependencies. In Universal Dependencies, [Dozat et al. \(2017\)](#) empirically showed that their dependency parser gains significant improvements by using POS tags predicted by a Bi-LSTM POS tagger. Indeed, [Kasai et al. \(2017\)](#) and [Friedman et al. \(2017\)](#) demonstrated that their unlexicalized neural network TAG parsers that only get as input predicted supertags can achieve state-of-the-art performance, with lexical inputs providing no improvement in performance. We initialize word embeddings to be the pre-trained GloVe vectors as in the supertagger. The other embeddings are randomly initialized.

2.2.2 Biaffine Parser

We train our parser to predict edges between lexical items in an LTAG derivation tree. Edges are labeled by the operations together with the deep syntactic roles of substitution sites (0=underlying subject, 1=underlying direct object, 2=underlying indirect object, 3,4=oblique arguments, CO=co-head for prepositional/particle verbs, and adj=all adjuncts). Figure 1 shows our biaffine parsing ar-

¹We fix the embedding of the ROOT token to be a 0-vector.

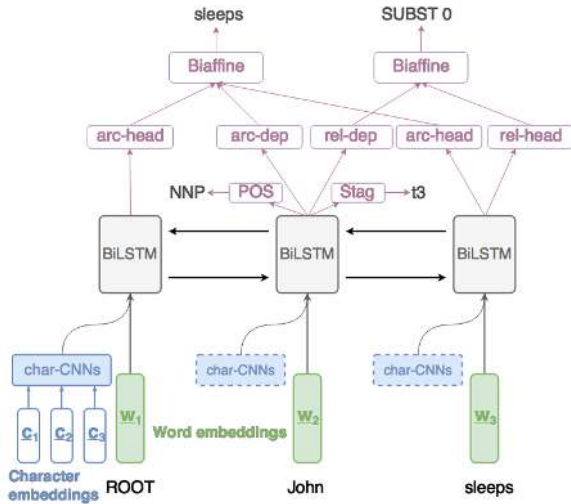


Figure 1: Biaffine parsing architecture. For the dependency from *John* to *sleeps* in the sentence *John sleeps*, the parser first predicts the head of *John* and then predicts the dependency label by combining the dependent and head representations. In the joint setting, the parser also predicts POS tags and supertags.

chitecture. Following Dozat and Manning (2017) and Kiperwasser and Goldberg (2016), we use BiLSTMs to obtain features for each word in a sentence. We add highway connections in the same fashion as our supertagging model.

We first perform unlabeled arc-factored scoring using the final output vectors from the BiLSTMs, and then label the resulting arcs. Specifically, suppose that we score edges coming into the i th word in a sentence i.e. assigning scores to the potential parents of the i th word. Denote the final output vector from the BiLSTM for the k th word by h_k and suppose that h_k is d -dimensional. Then, we produce two vectors from two separate multilayer perceptrons (MLPs) with the ReLU activation:

$$\begin{aligned} h_k^{\text{arc-dep}} &= \text{MLP}^{\text{(arc-dep)}}(h_k) \\ h_k^{\text{arc-head}} &= \text{MLP}^{\text{(arc-head)}}(h_k) \end{aligned}$$

where $h_k^{\text{arc-dep}}$ and $h_k^{\text{arc-head}}$ are d_{arc} -dimensional vectors that represent the k th word as a dependent and a head respectively. Now, suppose the k th row of matrix $H^{\text{(arc-head)}}$ is $h_k^{\text{arc-head}}$. Then, the probability distribution s_i over the potential heads of the i th word is computed by

$$s_i = \text{softmax}(H^{\text{(arc-head)}}W^{\text{(arc)}}h_i^{\text{arc-dep}} + H^{\text{(arc-head)}}b^{\text{(arc)}}) \quad (7)$$

where $W^{\text{(arc)}} \in \mathbb{R}^{d_{\text{arc}} \times d_{\text{arc}}}$ and $b^{\text{(arc)}} \in \mathbb{R}^{d_{\text{arc}}}$. In training, we simply take the greedy maximum

probability to predict the parent of each word. In the testing phase, we use the heuristics formulated by Dozat and Manning (2017) to ensure that the resulting parse is single-rooted and acyclic.

Given the head prediction of each word in the sentence, we assign labeling scores using vectors obtained from two additional MLP with ReLU. For the k th word, we obtain:

$$\begin{aligned} h_k^{\text{rel-dep}} &= \text{MLP}^{\text{(rel-dep)}}(h_k) \\ h_k^{\text{rel-head}} &= \text{MLP}^{\text{(rel-head)}}(h_k) \end{aligned}$$

where $h_k^{\text{rel-dep}}, h_k^{\text{rel-head}} \in \mathbb{R}^{d_{\text{rel}}}$. Let p_i be the index of the predicted head of the i th word, and r be the number of dependency relations in the dataset. Then, the probability distribution ℓ_i over the possible dependency relations of the arc pointing from the p_i th word to the i th word is calculated by:

$$\begin{aligned} \ell_i &= \text{softmax}(h_{p_i}^{T(\text{rel-head})}U^{(\text{rel})}h_i^{\text{(rel-dep)}} \\ &+ W^{(\text{rel})}(h_i^{\text{(rel-head)}} + h_{p_i}^{\text{(rel-head)}}) + b^{(\text{rel})}) \end{aligned} \quad (8)$$

where $U^{(\text{rel})} \in \mathbb{R}^{d_{\text{rel}} \times d_{\text{rel}} \times r}$, $W^{(\text{rel})} \in \mathbb{R}^{r \times d_{\text{rel}}}$, and $b^{(\text{rel})} \in \mathbb{R}^r$.

We generally follow the hyperparameters chosen in Dozat and Manning (2017). Specifically, we use BiLSTMs layers with 400 units each. Input, layer-to-layer, and recurrent dropout rates are all 0.33. The depths of all MLPs are all 1, and the MLPs for unlabeled attachment and those for labeling contain 500 (d_{arc}) and 100 (d_{rel}) units respectively. For character-level CNNs, we use the hyperparameters from Ma and Hovy (2016).

We train this model with the Adam algorithm to minimize the sum of the cross-entropy losses from head predictions (s_i from Eq. 7) and label predictions (ℓ_i from Eq. 8) with $\ell = 0.01$ and batch size 100 (Kingma and Ba, 2015). After each training epoch, we test the parser on the dev set. When labeled attachment score (LAS)² does not improve on five consecutive epochs, training ends.

2.3 Joint Modeling

The simple BiLSTM feature representations for parsing presented above are conducive to joint modeling of POS tagging and supertagging; rather than using POS tags and supertags to predict a derivation tree, we can instead use the BiLSTM hidden vectors derived from lexical inputs alone

²We disregard pure punctuation when evaluating LAS and UAS, following prior work (Bangalore et al., 2009; Chung et al., 2016; Kasai et al., 2017; Friedman et al., 2017).

to predict POS tags and supertags along with the TAG derivation tree.

$$h_k^{\text{pos}} = \text{MLP}^{(\text{pos})}(h_k)$$

$$h_k^{\text{stag}} = \text{MLP}^{(\text{stag})}(h_k)$$

where $h_k^{\text{pos}} \in \mathbb{R}^{d_{\text{pos}}}$ and $h_k^{\text{stag}} \in \mathbb{R}^{d_{\text{stag}}}$. We obtain probability distribution over the POS tags and supertags by:

$$\text{softmax}(W^{(\text{pos})}h_k^{\text{pos}} + b^{(\text{pos})}) \quad (9)$$

$$\text{softmax}(W^{(\text{stag})}h_k^{\text{stag}} + b^{(\text{stag})}) \quad (10)$$

where $W^{(\text{pos})}$, $b^{(\text{pos})}$, $W^{(\text{stag})}$, and $b^{(\text{stag})}$ are in $\mathbb{R}^{n_{\text{pos}} \times d_{\text{pos}}}$, $\mathbb{R}^{n_{\text{pos}}}$, $\mathbb{R}^{n_{\text{stag}} \times d_{\text{stag}}}$, and $\mathbb{R}^{n_{\text{stag}}}$ respectively, with n_{pos} and n_{stag} the numbers of possible POS tags and supertags respectively.

We use the same hyperparameters as in the parser. The MLPs for POS tagging and supertagging both contain 500 units. We again train this model with the Adam algorithm to minimize the sum of the cross-entropy losses from head predictions (s_i from Eq. 7), label predictions (ℓ_i from Eq. 8), POS predictions (Eq. 9), and supertag predictions (Eq. 10) with $\ell = 0.01$ and batch size 100. After each training epoch, we test the parser on the dev set and compute the percentage of each token that is assigned the correct parent, relation, supertag, and POS tag. When the percentage does not improve on five consecutive epochs, training ends.

This joint modeling has several advantages. First, the joint model yields a full syntactic analysis simultaneously without the need for training separate models or performing jackknife training. Secondly, joint modeling introduces a bias on the hidden representations that could allow for better generalization in each task (Caruana, 1997). Indeed, in experiments described in a later section, we show empirically that predicting POS tags and supertags does indeed benefit performance on parsing (as well as the tagging tasks).

3 Results and Discussion

We follow the protocol of Bangalore et al. (2009), Chung et al. (2016), Kasai et al. (2017), and Friedman et al. (2017); we use the grammar and the TAG-annotated WSJ Penn Tree Bank extracted by Chen et al. (2005). Following that work, we use Sections 01-22 as the training set, Section 00 as the dev set, and Section 23 as the test set. The training, dev, and test sets comprise 39832, 1921,

and 2415 sentences, respectively. We implement all of our models in TensorFlow (Abadi et al., 2016).³

3.1 Supertaggers

Our BiLSTM POS tagger yielded 97.37% and 97.53% tagging accuracy on the dev and test sets, performance on par with the state-of-the-art (Ling et al., 2015; Ma and Hovy, 2016).⁴ Seen in the middle section of Table 1 is supertagging performance obtained from various model configurations. ‘‘Final concat’’ in the model name indicates that vectors from forward and backward pass are concatenated only after the final layer. Concatenation happens after each layer otherwise. Numbers immediately after BiLSTM indicate the numbers of layers. CNN, HW, and POS denote respectively character-level CNNs, highway connections, and pipeline POS input from our BiLSTM POS tagger. Firstly, the differences in performance between BiLSTM2 (final concat) and BiLSTM2 and between BiLSTM2 and BiLSTM2-CNN suggest an advantage to performing concatenation after each layer and adding character-level CNNs. Adding predicted POS to the input somewhat helps supertagging though the difference is small. Adding a third BiLSTM layer helps only if there are highway connections, presumably because deeper BiLSTMs are more vulnerable to the vanishing/exploding gradient problem. Our supertagging model (BiLSTM3-HW-CNN-POS) that performs best on the dev set achieves an accuracy of 90.81% on the test set, outperforming the previously best result by more than 1.3%.

3.2 Parsers

Table 3 shows parsing results on the dev set. Abbreviations for models are as before with one addition: Stag denotes pipeline supertag input from our best supertagger (BiLSTM3-HW-CNN-POS in Table 1). As with supertagging, we observe a gain from adding character-level CNNs. Interestingly, adding predicted POS tags or supertags deteriorates performance with BiLSTM3. These results suggest that morphological information and word information from character-level CNNs and word embeddings overwhelm the in-

³Our code is available online for easy replication of our results at https://github.com/jungokasai/graph_parser.

⁴We cannot directly compare these results because the data split is different in the POS tagging literature.

Supertagger	Dev	Test
Bangalore et al. (2009)	88.52	86.85
Chung et al. (2016)	87.88	–
Kasai et al. (2017)	89.32	89.44
BiLSTM2 (final concat)	88.96	–
BiLSTM2	89.60	–
BiLSTM2-CNN	89.97	–
BiLSTM2-CNN-POS	90.03	–
BiLSTM2-HW-CNN-POS	90.12	–
BiLSTM3-CNN-POS	90.12	–
BiLSTM3-HW-CNN-POS	90.45	90.81
BiLSTM4-CNN-POS	89.99	–
BiLSTM4-HW-CNN-POS	90.43	–
Joint (Stag)	90.51	–
Joint (POS+Stag)	90.67	91.01

Table 1: Supertagging Results. Joint (Stag) and Joint (POS+Stag) indicate joint parsing models that perform supertagging, and POS tagging and supertagging respectively.

POS tagger	Dev	Test
BiLSTM	97.37	97.53
Joint (POS+Stag)	97.54	97.73

Table 2: POS tagging results.

formation from predicted POS tags and supertags. Again, highway connections become crucial as the number of layers increases. We finally evaluate the parsing model with the best dev performance (BiLSTM4-HW-CNN) on the test set (Table 3). It achieves 91.37 LAS points and 92.77 UAS points, improvements of 1.8 and 1.7 points respectively from the state-of-the-art.

3.3 Joint Models

We provide joint modeling results for supertagging and parsing in Tables 2 and 3. For these joint models, we employed the best parsing configuration (4 layers of BiLSTMs, character-level CNNs, and highway connections), with and without POS tagging added as an additional task. We can observe that our full joint model that performs




Figure 2: F1 Score with Dependency Length.

Parser	Dev		Test	
	UAS	LAS	UAS	LAS
Bangalore et al. (2009)	87.60	85.80	86.66	84.90
Chung et al. (2016)	89.96	87.86	–	–
Friedman et al. (2017)	90.36	88.91	90.31	88.96
Kasai et al. (2017)	90.88	89.39	90.97	89.68
BiLSTM3	91.75	90.22	–	–
BiLSTM3-CNN	92.27	90.76	–	–
BiLSTM3-CNN-POS	92.07	90.53	–	–
BiLSTM3-CNN-Stag	92.15	90.65	–	–
BiLSTM3-HW-CNN	92.29	90.71	–	–
BiLSTM4-CNN	92.11	90.66	–	–
BiLSTM4-HW-CNN	92.78	91.26	92.77	91.37
BiLSTM5-CNN	92.34	90.77	–	–
BiLSTM5-HW-CNN	92.64	91.11	–	–
Joint (Stag)	92.97	91.48	–	–
Joint (POS+Stag)	93.22	91.80	93.26	91.89
Joint (Shuffled Stag)	92.23	90.56	–	–

Table 3: Parsing results on the dev and test sets.

POS tagging, supertagging, and parsing further improves performance in all of the three tasks, yielding the test result of 91.89 LAS and 93.26 UAS points, an improvement of more than 2.2 points each from the state-of-the-art.

Figures 2 and 3 illustrate the relative performance of the feed-forward neural network shift-reduce TAG parser (Kasai et al., 2017) and our joint graph-based parser with respect to two of the measures explored by McDonald and Nivre (2011), namely dependency length and distance between a dependency and the root of a parse. The graph-based parser outperforms the shift-reduce parser across all conditions. Most interesting is the fact that the graph-based parser shows less of an effect of dependency length. Since the shift-reduce parser builds a parse sequentially with one parsing action depending on those that come before it, we would expect to find a propagation of errors made in establishing shorter dependencies to the establishment of longer dependencies.

Lastly, it is worth noting our joint parsing ar-




Figure 3: F1 Score with Distance to Root.

chitecture has a substantial advantage regarding parsing speed. Since POS tagging, supertagging, and parsing decisions are made independently for each word in a sentence, our system can parallelize computation once the sentence is encoded in the BiLSTM layers. Our current implementation processes 225 sentences per second on a single Tesla K80 GPU, an order of magnitude faster than the MICA system (Bangalore et al., 2009).⁵

4 Joint Modeling and Network Representations

Given the improvements we have derived from the joint models, we analyze the nature of inductive bias that results from multi-task training and attempt to provide an explanation as to why joint modeling improves performance.

4.1 Noise vs. Inductive Bias

One might argue that joint modeling improves performance merely because it adds noise to each task and prevents over-fitting. If the introduction of noise were the key, we would still expect to gain an improvement in parsing even if the target supertag were corrupted, say by shuffling the order of supertags for the entire training data (Caruana, 1997). We performed this experiment, and the result is shown as “Joint (Shuffled Stag)” in Table 3. Parsing performance falls behind the best non-joint parser by 0.7 LAS points. This suggests that inducing the parser to create representations to predict both supertags and a parse tree is beneficial for both tasks, beyond a mere introduction of noise.

4.2 Syntactic Analogies

We next analyze the induced vector representations in the output projection matrices of our supertagger and joint parsers using the syntactic analogy framework (Kasai et al., 2017). Consider, for instance, the analogy that an elementary tree representing a clause headed by a transitive verb (t27) is to a clause headed by an intransitive verb (t81) as a subject relative clause headed by a transitive verb (t99) is to a subject relative headed by an intransitive verb (t109). Following the ideas in Mikolov et al. (2013) for word analogies, we can express this structural analogy as $t27 - t81 +$

⁵While such computational resources were not available in 2009, our parser differs from the MICA chart parser in being able to better exploit parallel computation enabled by modern GPUs.

$t109 = t99$ and test it by cosine similarity. Table 4 shows the results of the analogy test with 246 equations involving structural analogies with only the 300 most frequent supertags in the training data. While the embeddings (projection matrix) from the independently trained supertagger do not appear to reflect the syntax, those obtained from the joint models yield linguistic structure despite the fact that the supertag embeddings (projection matrix) is trained without any *a priori* syntactic knowledge about the elementary trees.

The best performance is obtained by the supertag representations obtained from the training of the transition-based parser Kasai et al. (2017) and Friedman et al. (2017). For the transition-based parser, it is beneficial to share statistics among the input supertags that differ only by a certain operation or property (Kasai et al., 2017) during the training phase, yielding the success in the analogy task. For example, a transitive verb supertag whose object has been filled by substitution should be treated by the parser in the same way as an intransitive verb supertag. In our graph-based parsing setting, we do not have a notion of parse history or partial derivations that directly connect intransitive and transitive verbs. However, syntactic analogies still hold to a considerable degree in the vector representations of supertags induced by our joint models, with average rank of the correct answer nearly the same as that obtained in the transition-based parser.

This analysis bolsters our hypothesis that joint training biases representation learning toward linguistically sensible structure. The supertagger is just trained to predict linear sequences of supertags. In this setting, many intervening supertags can occur, for instance, between a subject noun and its verb, and the supertagger might not be able to systematically link the presence of the two in the sequence. In the joint models, on the other hand, parsing actions will explicitly guide the network to associate the two supertags.

5 Downstream Tasks

Previous work has applied TAG parsing to the downstream tasks of syntactically-oriented textual entailment (Xu et al., 2017) and semantic role labeling (Chen and Rambow, 2003). In this work, we apply our parsers to the textual entailment and unbounded dependency recovery tasks and achieve state-of-the-art performance. These re-

Parser / Supertagger	%correct	Avg. rank
Transition-based	67.07	2.36
Our Supertagger	0.00	152.46
Our Joint (Stag)	29.27	2.55
Our Joint (POS+Stag)	30.08	2.57

Table 4: Syntactic analogy test results on the 300 most frequent supertags. Avg. rank is the average position of the correct choice in the ranked list of the closest neighbors; the top line indicates the result of using supertag embeddings that are trained jointly with a transition based parser (Friedman et al., 2017).

sults bolster the significance of the improvements gained from our joint parser and the utility of TAG parsing for downstream tasks.

5.1 PETE

Parser Evaluation using Textual Entailments (PETE) is a shared task from the SemEval-2010 Exercises on Semantic Evaluation (Yuret et al., 2010). The task was intended to evaluate syntactic parsers across different formalisms, focusing on entailments that could be determined entirely on the basis of the syntactic representations of the sentences that are involved, without recourse to lexical semantics, logical reasoning, or world knowledge. For example, syntactic knowledge alone tells us that the sentence *John, who loves Mary, saw a squirrel* entails *John saw a squirrel* and *John loves Mary* but not, for instance, that *John knows Mary* or *John saw an animal*. Prior work found the best performance was achieved with parsers using grammatical frameworks that provided rich linguistic descriptions, including CCG (Rimell and Clark, 2010; Ng et al., 2010), Minimal Recursion Semantics (MRS) (Lien, 2014), and TAG (Xu et al., 2017). Xu et al. (2017) provided a set of linguistically-motivated transformations to use TAG derivation trees to solve the PETE task. We follow their procedures and evaluation for our new parsers.

We present test results from two configurations in Table 5. One configuration is a pipeline approach that runs our BiLSTM POS tagger, supertagger, and parser. The other one is a joint approach that only uses our full joint parser. The joint method yields 78.1% in accuracy and 76.4% in F1, improvements of 2.4 and 2.7 points over the previously reported best results.

System	%A	%P	%R	F1
Rimell and Clark (2010)	72.4	79.6	62.8	70.2
Ng et al. (2010)	70.4	68.3	80.1	73.7
Lien (2014)	70.7	88.6	50.0	63.9
Xu et al. (2017)	75.7	88.1	61.5	72.5
Our Pipeline Method	77.1	86.6	66.0	74.9
Our Joint Method	78.1	86.3	68.6	76.4

Table 5: PETE test results. Precision (P), recall (R), and F1 are calculated for “entails.”

5.2 Unbounded Dependency Recovery

The unbounded dependency corpus (Rimell et al., 2009) specifically evaluates parsers on unbounded dependencies, which involve a constituent moved from its original position, where an unlimited number of clause boundaries can intervene. The corpus comprises 7 constructions: object extraction from a relative clause (ObRC), object extraction from a reduced relative clause (ObRed), subject extraction from a relative clause (SbRC), free relatives (Free), object wh-questions (ObQ), right node raising (RNR), and subject extraction from an embedded clause (SbEm).

Because of variations across formalisms in their representational format for unbounded dependencies, past work has conducted manual evaluation on this corpus (Rimell et al., 2009; Nivre et al., 2010). We instead conduct an automatic evaluation using a procedure that converts TAG parses to structures directly comparable to those specified in the unbounded dependency corpus. To this end, we apply two types of structural transformation in addition to those used for the PETE task:⁶ 1) a more extensive analysis of coordination, 2) resolution of differences in dependency representations in cases involving copula verbs and co-anchors (e.g., verbal particles). See Appendix A for details. After the transformations, we simply check if the resulting dependency graphs contain target labeled arcs given in the dataset.

Table 6 shows the results. Our joint parser outperforms the other parsers, including the neural network shift-reduce TAG parser (Kasai et al., 2017). Our data-driven parsers yield relatively low performance in the ObQ and RNR constructions. Performance on ObQ is low, we expect, because of their rarity in the data on which the parser is

⁶One might argue that since the unbounded dependency evaluation is recall-based, we added too many edges by the transformations. However, it turns out that applying all the transformations for the corpus even improves performance on PETE (77.6 F1 score), which considers precision and recall, verifying that our transformations are reasonable.

System	ObRC	ObRed	SbRC	Free	ObQ	RNR	SbEm	Total	Avg
C&C (CCG)	59.3	62.6	80.0	72.6	72.6	49.4	22.4	53.6	61.1
Enju (HPSG)	47.3	65.9	82.1	76.2	32.5	47.1	32.9	54.4	54.9
Stanford (PCFG)	22.0	1.1	74.7	64.3	41.2	45.4	10.6	38.1	37.0
MST (Stanford Dependencies)	34.1	47.3	78.9	65.5	41.2	45.4	37.6	49.7	50.0
MALT (Stanford Dependencies)	40.7	50.5	84.2	70.2	31.2	39.7	23.5	48.0	48.5
NN Shift-Reduce TAG Parser	60.4	75.8	68.4	79.8	53.8	45.4	44.7	59.4	61.2
Our Joint Method	72.5	78.0	81.1	85.7	56.3	47.1	49.4	64.9	67.0

Table 6: Parser accuracy on the unbounded dependency corpus. The results of the first five parsers are taken from Rimell et al. (2009) and Nivre et al. (2010). The Total and Avg columns indicate the percentage of correctly recovered dependencies out of all dependencies and the average of accuracy on the 7 constructions.

trained.⁷ For RNR, rarity may be an issue as well as the limits of the TAG analysis of this construction. Nonetheless, we see that the rich structural representations that a TAG parser provides enables substantial improvements in the extraction of unbounded dependencies. In the future, we hope to evaluate state-of-the-art Stanford dependency parsers automatically.

6 Related Work

The two major classes of data-driven methods for dependency parsing are often called transition-based and graph-based parsing (Kübler et al., 2009). Transition-based parsers (e.g. MALT (Nivre, 2003)) learn to predict the next transition given the input and the parse history. Graph-based parsers (e.g. MST (McDonald et al., 2005)) are trained to directly assign scores to dependency graphs.

Empirical studies have shown that a transition-based parser and a graph-based parser yield similar overall performance across languages (McDonald and Nivre, 2011), but the two strands of data-driven parsing methods manifest the fundamental trade-off of parsing algorithms. The former prefers rich feature representations with parsing history over global training and exhaustive search, and the latter allows for global training and inference at the expense of limited feature representations (Kübler et al., 2009).

Recent neural network models for transition-based and graph-based parsing can be viewed as remedies for the aforementioned limitations. Andor et al. (2016) developed a transition-based parser using feed-forward neural networks that performs global training approximated by beam search. The globally normalized objective addresses the label bias problem and makes global

⁷The substantially better performance of the C&C parser is in fact the result of additions that were made to the training data.

training effective in the transition-based parsing setting. Kiperwasser and Goldberg (2016) incorporated a dynamic oracle (Goldberg and Nivre, 2013) in a BiLSTM transition-based parser that remedies global error propagation. Kiperwasser and Goldberg (2016) and Dozat and Manning (2017) proposed graph-based parsers that have access to rich feature representations obtained from BiLSTMs.

Previous work integrated CCG supertagging and parsing using belief propagation and dual decomposition approaches (Auli and Lopez, 2011). Nguyen et al. (2017) incorporated a graph-based dependency parser (Kiperwasser and Goldberg, 2016) with POS tagging. Our work followed these lines of effort and improved TAG parsing performance.

7 Conclusion and Future Work

In this work, we presented a state-of-the-art TAG supertagger, a parser, and a joint parser that performs POS tagging, supertagging, and parsing. The joint parser has the benefit of giving a full syntactic analysis of a sentence simultaneously. Furthermore, the joint parser achieved the best performance, an improvement of over 2.2 LAS points from the previous state-of-the-art. We have also seen that the joint parser yields state-of-the-art in textual entailment and unbounded dependency recovery tasks, and raised the possibility that TAG can provide useful structural analysis of sentences for other NLP tasks. We will explore more applications of our TAG parsers in future work.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.

- Daniel Andor, Chris Alberti, David Weiss, Aliaksei Severyn, Alessandro Presta, Kuzman Ganchev, Slav Petrov, and Michael Collins. 2016. Globally normalized transition-based neural networks. In *ACL*. Association for Computational Linguistics, Berlin, Germany, pages 2442–2452. <http://www.aclweb.org/anthology/P16-1231>.
- Michael Auli and Adam Lopez. 2011. A comparison of loopy belief propagation and dual decomposition for integrated CCG supertagging and parsing. In *ACL*. Association for Computational Linguistics, pages 470–480. <http://www.aclweb.org/anthology/P11-1048>.
- Srinivas Bangalore, Pierre Boullier, Alexis Nasr, Owen Rambow, and Benoît Sagot. 2009. MICA: A probabilistic dependency parser based on tree insertion grammars (application note). In *NAACL-HLT (short)*. Association for Computational Linguistics, Boulder, Colorado, pages 185–188. <http://www.aclweb.org/anthology/N09/N09-2047>.
- Srinivas Bangalore and Aravind K. Joshi. 1999. Supertagging: An Approach to Almost Parsing. *Computational Linguistics* 25:237–266.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*. OReilly Media.
- Rich Caruana. 1997. Multitask learning. *Machine Learning* 28:41–75.
- John Chen, Srinivas Bangalore, and K. Vijay-Shanker. 2005. Automated extraction of tree-adjoining grammars from treebanks. *Natural Language Engineering* 12(3):251–299.
- John Chen and Owen Rambow. 2003. Use of deep linguistic features for the recognition and labeling of semantic arguments. In *EMNLP*. pages 41–48. <http://aclanthology.coli.uni-saarland.de/pdf/W/W03/W03-1006.pdf>.
- Jason Chiu and Eric Nichols. 2016. Named entity recognition with bidirectional LSTM-CNNs. *TACL* 4:357–370. <https://transacl.org/ojs/index.php/tacl/article/view/792>.
- Wonchang Chung, Suhas Siddhesh Mhatre, Alexis Nasr, Owen Rambow, and Srinivas Bangalore. 2016. Revisiting supertagging and parsing: How to use supertags in transition-based parsing. In *TAG+*. pages 85–92. <http://www.aclweb.org/anthology/W16-3309>.
- Stephen Clark and James R Curran. 2007. Wide-coverage efficient statistical parsing with CCG and log-linear models. *Computational Linguistics* 33(4):493–552. <http://www.newdesign.aclweb.org/anthology-new/J/J07/J07-4004.pdf>.
- Timothy Dozat and Christopher Manning. 2017. Deep biaffine attention for neural dependency parsing. In *ICLR*.
- Timothy Dozat, Peng Qi, and Christopher D. Manning. 2017. Stanford’s graph-based neural dependency parser at the CoNLL 2017 shared task. In *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, Vancouver, Canada, pages 20–30. <http://www.aclweb.org/anthology/K17-3002>.
- Dan Friedman, Jungo Kasai, R. Thomas McCoy, Robert Frank, Forrest Davis, and Owen Rambow. 2017. Linguistically rich vector representations of supertags for TAG parsing. In *TAG+*. Association for Computational Linguistics, Umeå, Sweden, pages 122–131. <http://www.aclweb.org/anthology/W17-6213>.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *NIPS*.
- Yoav Goldberg and Joakim Nivre. 2013. Training deterministic parsers with non-deterministic oracles. *TACL* 1:403–414. <https://www.aclweb.org/anthology/Q/Q13/Q13-1033.pdf>.
- Alex Graves and Jürgen Schmidhuber. 2005. Frame-wise phoneme classification with bidirectional LSTM and other neural network architectures. *Neural Networks* 18(5):602–610.
- Julia Hockenmaier and Mark Steedman. 2007. CCG-bank: a corpus of CCG derivations and dependency structures extracted from the Penn Treebank. *Computational Linguistics* 33(3):355–396.
- Aravind K. Joshi and Yves Schabes. 1997. Tree-adjoining grammars. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages, Volume 3: Beyond Words*, Springer, New York, pages 69–124.
- Jungo Kasai, Robert Frank, Tom McCoy, Owen Rambow, and Alexis Nasr. 2017. TAG parsing with neural networks and vector representations of supertags. In *EMNLP*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1713–1723. <https://www.aclweb.org/anthology/D17-1180>.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. ADAM: A Method for Stochastic Optimization. In *ICLR*.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional lstm feature representations. *TACL* 4:313–327. <https://transacl.org/ojs/index.php/tacl/article/view/885>.
- Sandra Kübler, Ryan McDonald, and Joakim Nivre. 2009. *Dependency Parsing*. Morgan & Claypool Publishers.

- Mike Lewis, Kenton Lee, and Luke Zettlemoyer. 2016. LSTM CCG parsing. In *HLT-NAACL*. pages 221–231.
- Elisabeth Lien. 2014. Using minimal recursion semantics for entailment recognition. In *Proceedings of the Student Research Workshop at EACL*. Gothenburg, Sweden, page 7684.
- Wang Ling, Tiago Luís, Luís Marujo, Ramón Fernández Astudillo, Silvio Amir, Chris Dyer, Alan W Black, and Isabel Trancoso. 2015. Finding function in form: Compositional character models for open vocabulary word representation. In *EMNLP*.
- Xuezhe Ma and Eduard Hovy. 2016. [End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF](#). In *ACL*. Association for Computational Linguistics, Berlin, Germany, pages 1064–1074. <http://www.aclweb.org/anthology/P16-1101>.
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The Penn Treebank. *Computational Linguistics* 19(2):313–330.
- Ryan McDonald and Joakim Nivre. 2011. [Analyzing and integrating dependency parsers](#). *Computational Linguistics* 37(1):197–230. http://dx.doi.org/10.1162/coli_a_00039.
- Ryan McDonald, Fernando Pereira, Kiril Ribarov, and Jan Hajic. 2005. [Non-projective dependency parsing using spanning tree algorithms](#). In *EMNLP*. Association for Computational Linguistics, Vancouver, British Columbia, Canada, pages 523–530. <http://www.aclweb.org/anthology/H/H05/H05-1066>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *NIPS*, Curran Associates, Inc., pages 3111–3119.
- Dominick Ng, James W.D. Constable, Matthew Honnibal, and James R. Curran. 2010. SCHWA: PETE using CCG dependencies with the C&C parser. In *SemEval*. page 313316.
- Dat Quoc Nguyen, Mark Dras, and Mark Johnson. 2017. [A novel neural network model for joint pos tagging and graph-based dependency parsing](#). In *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, pages 134–142. <https://doi.org/10.18653/v1/K17-3014>.
- Joakim Nivre. 2003. An efficient algorithm for projective dependency parsing. In *IWPT*.
- Joakim Nivre, Laura Rimell, Ryan McDonald, and Carlos Gómez Rodríguez. 2010. [Evaluation of dependency parsers on unbounded dependencies](#). In *COLING*. Coling 2010 Organizing Committee, Beijing, China, pages 833–841. <http://www.aclweb.org/anthology/C10-1094>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *EMNLP*. pages 1532–1543.
- Owen Rambow and Aravind Joshi. 1994. A formal look at dependency grammars and phrase-structure grammars, with special consideration of word-order phenomena. In Leo Wanner, editor, *Recent Trends in Meaning-Text Theory*, Amsterdam and Philadelphia, pages 167–190.
- Laura Rimell and Stephen Clark. 2010. Cambridge: Parser evaluation using textual entailment by grammatical relation comparison. In *SemEval*. pages 268–271.
- Laura Rimell, Stephen Clark, and Mark Steedman. 2009. Unbounded dependency recovery for parser evaluation. In *EMNLP*. Singapore, page 813821.
- Cicero D. Santos and Bianca Zadrozny. 2014. [Learning character-level representations for part-of-speech tagging](#). In Tony Jebara and Eric P. Xing, editors, *ICML*. JMLR Workshop and Conference Proceedings, pages 1818–1826. <http://jmlr.org/proceedings/papers/v32/santos14.pdf>.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. *arXiv preprint arXiv:1505.00387*.
- Mark Steedman and Jason Baldrige. 2011. Combinatory categorial grammar. In Robert Borsley and Kersti Börjars, editors, *Non-Transformational Syntax: Formal and Explicit Models of Grammar*, Wiley-Blackwell.
- Ashish Vaswani, Yonatan Bisk, Kenji Sagae, and Ryan Musa. 2016. [Supertagging with LSTMs](#). In *NAACL*. Association for Computational Linguistics, San Diego, California, pages 232–237. <http://www.aclweb.org/anthology/N16-1027>.
- Pauli Xu, Robert Frank, Jungo Kasai, and Owen Rambow. 2017. [TAG parser evaluation using textual entailments](#). In *TAG+*. Association for Computational Linguistics, Umeå, Sweden, pages 132–141. <http://www.aclweb.org/anthology/W17-6214>.
- Wenduan Xu. 2016. [LSTM shift-reduce CCG parsing](#). In *EMNLP*. Association for Computational Linguistics, Austin, Texas, pages 1754–1764. <https://aclweb.org/anthology/D16-1181>.
- Wenduan Xu, Michael Auli, and Stephen Clark. 2015. [CCG supertagging with a recurrent neural network](#). In *ACL*. Association for Computational Linguistics,

Beijing, China, pages 250–255. <http://www.aclweb.org/anthology/P15-2041>.

Deniz Yuret, Aydin Han, and Zehra Turgut. 2010. *SemEval-2010 task 12: Parser evaluation using textual entailments*. In *SemEval*. Association for Computational Linguistics, Uppsala, Sweden, pages 51–56. <http://www.aclweb.org/anthology/S10-1009>.

A Transformations for Unbounded Dependency Recovery Corpus

For automatic evaluation on the unbounded dependency recovery corpus (UDR, Rimell et al. (2009)), we run simple conversion of dependency labels in UDR to those in our TAG grammar (See Table 7) with a couple of exceptions.

- Change arcs from verbs to wh-adverbs as in “where is the city located?” to adjunction.
- Reflect causative-inchoative alternation in the subject embedded construction. Concretely, change the role of “door” in “hold the door shut” from the subject to the object of “shut.”

We then transform TAG dependency trees. Finally, we simply check if the resulting dependency graphs contain target labeled arcs given in the dataset.

Below is a full description of transformations. This set of structural transformations is applied in the order in which we will present it, so that the output of previous transformations can feed subsequent ones. In the following, we denote an arc pointing from node B to node A with label C as (A, B, C) where A and B are called the child (dependent) and the parent (head) in the relation.

A.1 Transformations from PETE

We apply three types of transformation from Xu et al. (2017) to interpret the TAG parses.

Relative Clauses When an elementary tree of a relative clause adjoins into a noun, we add a reverse arc with the label reflecting the type of the relative clause elementary tree. For a subject relative, we add a 0-labeled arc, for an object relative, we add a 1-labeled arc, and so forth.

UDR Labels	TAG Labels
nsubj, cop	0
dobj, pobj, obj2, nsubjpass	1
others (advmod etc)	ADJ

Table 7: UD to TAG label conversion.

Sentential Complements Sentential complementation in TAG derivations can be analyzed via either adjoining the higher clause into the embedded clause (necessarily so in cases of long-distance extraction from the embedded clause) or substituting the embedded clause in the higher clause. In order to normalize this divergence, for an adjunction arc involving a predicative auxiliary elementary tree (supertag), we add a reverse arc involving the 1 relation (sentential complements).

A.2 Coordination

We roughly follow the method presented in Xu et al. (2017) with extensions. Under the TAG analysis, VP coordination involves a VP-recursive auxiliary tree headed by the coordinator that includes a VP substitution node (for the second conjunct) with label 1. In order to allow the first clauses subject argument (as well as modal verbs and negations) to be shared by the second verb, we add the relevant relations to the second verb. In addition, we analyze sentential coordination cases. Sentence coordination in our TAG grammar usually happens between two complete sentences and no modifiers or arguments are shared, and therefore it can be analyzed via substituting a sentence into the coordinator with label 1. However, when sentential coordination happens between two relative clause modifiers, our TAG grammar analyzes the second clause as a complete sentence, meaning that we need to recover the extracted argument by consulting the property of the first clause. Furthermore, the deep syntactic role of the extracted argument can be different in the two relative clauses. For instance, in the sentence, “... the same stump which had impaled the car of many a guest in the past thirty years and which he refused to have removed,” we need to recover an arc from removed to stump with label 1 whereas the arc from impaled to stump has label 0. To resolve this issue, when there is coordination of two relative clause modifiers, we add an edge from the head of the second clause to the modified noun with the same label as the label that under which the relative pronoun is attached to the head.

A.3 Resolving Differences in Dependency Representations

Small Clauses The UDR corpus has inconsistency with regards to small clauses. UDR gives an analysis that a small clause contains a subject and a complement as in (nsubj, guy, liar) in

“the guy who I call a liar.” in the subject embedded constructions. However, in the object question and object free relative constructions, a small clause is analyzed as two arguments of the verb. For instance, UDR specifies (what, adopted, dobj) in “we adopted what I would term pseudo-capitalism.” To solve this problem we add an arc from the head of the matrix clause to the subject in a small clause with label 1.

Co-anchors In our TAG grammar, Co-anchor attachment represents the substitution into a node that is construed as a co-head of an elementary tree. For instance, “for” is deemed as a co-anchor to “hope” in the sentence “*that is exactly what I’m hoping for*” (Figure 4). In this case, UDR would pick the relation (what, hope, pobj). Therefore, when there is a co-anchor to a head tree, we add all arcs that involve the head tree to the co-anchor tree.

Wh-determiners and Wh-adverbs Our TAG grammar analyzes a wh-determiner via adjoining the noun into the wh-determiner (Figure 5). This is also true for cases where a wh-adverb is followed by an adjective and a noun as in *how many battles did she win?* In contrast, UDR corpus gives an analysis that the noun is the head of the constituent. In order to resolve this discrepancy, when a word adjoins into a wh-word,⁸ we pick all arcs with the wh-word as the child and add the arcs obtained from such arcs by replacing the wh-word child by the word adjoining into the wh-word.

Copulas A copula is usually treated as a dependent to the predicate both in our TAG grammar (adjunction) and UDR. However, we found two situations where they differ from each other. First, when wh-extraction happens on the complement, as in “obviously there has been no agreement on what American conservatism is, or rather, what it should be,” the TAG grammar analyzes it via substituting the wh-word (“what”) into the copula (“is”). To reconcile this disagreement between the TAG grammar and UDR, when substitution happens into a *be* verb, we add the substitution into




Figure 4: Co-anchor case from a sentence “*that is exactly what I’m hoping for*”. The UDR gives the red arc (what, for, pobj). The blue arc (what, for, 1) is obtained from (what, hope, 1).




Figure 5: Wh-determiner case from a sentence *What songs did he sing?* The UDR gives the red arc (songs, sing, dobj). The blue arc (song, sing, 1) is obtained from (what, sing, 1) and (songs, what, ADJ).

⁸We considered imposing a more strict condition that the word adjoining into the wh-word is a noun, but we found cases that this method fails to cover; for example, UDR gives (dobj, get, much) for a sentence “opinion is mixed on how **much** of a boost the overall stock market would **get** even if dividend growth continues at double-digit levels.”

the copula.⁹ Second, UDR treats non-*be* copulas differently than *be* verbs. An example is the UDR relation (those, stayed, nsubj) “in the other hemisphere it is growing colder and nymphs, those who stayed alive through the summer, are being brought into nests for quickening and more growing” where our parser yields (those, alive, 0). For this reason, when a lemma of a verb is a non-*be* copula,¹⁰ we add arcs involving the word to the copula adjoining into the copula.

PP attachment with multiple noun candidates

We observed that PP attachment with multiple noun candidates is often at stake in UDR.¹¹ For instance, UDR provides (part, had, nsubj) and (several, tried, nsubj) for the sentences “... there is no part of the earth that has not had them” and “there were several on the Council who tried to live like Christians” while the TAG parser outputs (earth, had, nsubj) and (Council, tried, nsubj) respectively. While we count these cases as “wrong” since they manifest certain disambiguation (though not purely unbounded dependency recovery), we ignore superficial (conventional) differences in head selection. In our TAG grammar “a lot of people” would be headed by “lot” whereas UDR would recognize “people” as the head. Hence, when “lot/lots/kind/kinds/none of” occurs, we add all arcs with “lot/lots/kind/kinds/none” to the head of the phrase that is the object of “of.”

Modals In the UDR corpus, a modal depends on an auxiliary verb following the modal, if there is one. For example, “Rosie reinvented this man, who may or may not have known about his child” is given the relation (may, have, aux). In the TAG grammar, both “may” and “have” adjoin into “known.” Therefore, when the head of a modal has another child with adjunction, we add an arc from the child to the modal.

Existential *there* UDR gives the “cop” relation between an existential *there* and the *be* verb. For example, it gives (be, legislation, cop) in “... on how much social legislation there should be.” On the other hand, our TAG grammar analyzes that

“there” is attached to “be” with label 0.¹² To resolve this issue, for arcs that point into an existential *there* with label 0, we add a reverse edge with label 0.

Determiner modifying a sentence Finally, when a determiner followed by an adjective modifies a sentence via adjunction in our TAG as in “the more highly placed they are – that is, the more they know – the more concerned they have become,” we add an edge from the verb to the adjective with label 1.

⁹We use the nltk lemmatizer (Bird et al., 2009) to identify *be* verbs.

¹⁰We chose “stay,” “become,” “seem,” and “remain.”

¹¹This is indeed one of the problems with UDR. Performance on UDR is not purely reflective of unbounded dependency recovery.

¹²Usually, “there” is attached to the noun, not the *be* verb, but in this case, extraction is happening on the noun, so the *be* verb becomes the head. See the discussion on copulas above.