

# Compound Probabilistic Context-Free Grammars for Grammar Induction

**Yoon Kim**

Harvard University  
Cambridge, MA, USA

yoonykim@seas.harvard.edu

**Chris Dyer**

DeepMind  
London, UK

cdyer@google.com

**Alexander M. Rush**

Harvard University  
Cambridge, MA, USA

srush@seas.harvard.edu

## Abstract

We study a formalization of the grammar induction problem that models sentences as being generated by a compound probabilistic context free grammar. In contrast to traditional formulations which learn a single stochastic grammar, our context-free rule probabilities are modulated by a per-sentence continuous latent variable, which induces marginal dependencies beyond the traditional context-free assumptions. Inference in this grammar is performed by collapsed variational inference, in which an amortized variational posterior is placed on the continuous variable, and the latent trees are marginalized with dynamic programming. Experiments on English and Chinese show the effectiveness of our approach compared to recent state-of-the-art methods for grammar induction from words with neural language models.

## 1 Introduction

Grammar induction is the task of inducing hierarchical syntactic structure from data. Statistical approaches to grammar induction require specifying a probabilistic grammar (e.g. formalism, number and shape of rules), and fitting its parameters through optimization. Early work found that it was difficult to induce probabilistic context-free grammars (PCFG) from natural language data through direct methods, such as optimizing the log likelihood with the EM algorithm (Lari and Young, 1990; Carroll and Charniak, 1992). While the reasons for the failure are manifold and not completely understood, two major potential causes are the ill-behaved optimization landscape and the overly strict independence assumptions of PCFGs. More successful approaches to grammar induction have thus resorted to carefully-crafted auxiliary objectives (Klein and Manning, 2002), priors or

non-parametric models (Kurihara and Sato, 2006; Johnson et al., 2007; Liang et al., 2007; Wang and Blunsom, 2013), and manually-engineered features (Huang et al., 2012; Golland et al., 2012) to encourage the desired structures to emerge.

We revisit these aforementioned issues in light of advances in model parameterization and inference. First, contrary to common wisdom, we find that parameterizing a PCFG’s rule probabilities with neural networks over distributed representations makes it possible to induce linguistically meaningful grammars by simply optimizing log likelihood. While the optimization problem remains non-convex, recent work suggests that there are optimization benefits afforded by over-parameterized models (Arora et al., 2018; Xu et al., 2018; Du et al., 2019), and we indeed find that this neural PCFG is significantly easier to optimize than the traditional PCFG. Second, this factored parameterization makes it straightforward to incorporate side information into rule probabilities through a sentence-level continuous latent vector, which effectively allows different contexts in a derivation to coordinate. In this compound PCFG—continuous mixture of PCFGs—the context-free assumptions hold conditioned on the latent vector but not unconditionally, thereby obtaining longer-range dependencies within a tree-based generative process.

To utilize this approach, we need to efficiently optimize the log marginal likelihood of observed sentences. While compound PCFGs break efficient inference, if the latent vector is known the distribution over trees reduces to a standard PCFG. This property allows us to perform grammar induction using a collapsed approach where the latent trees are marginalized out exactly with dynamic programming. To handle the latent vector, we employ standard amortized inference using reparameterized samples from a variational

Code: <https://github.com/harvardnlp/compound-pcfg>

posterior approximated from an inference network (Kingma and Welling, 2014; Rezende et al., 2014).

On standard benchmarks for English and Chinese, the proposed approach is found to perform favorably against recent neural network-based approaches to grammar induction (Shen et al., 2018, 2019; Drozdov et al., 2019; Kim et al., 2019).

## 2 Probabilistic Context-Free Grammars

We consider context-free grammars (CFG) consisting of a 5-tuple  $\mathcal{G} = (S, \mathcal{N}, \mathcal{P}, \Sigma, \mathcal{R})$  where  $S$  is the distinguished start symbol,  $\mathcal{N}$  is a finite set of nonterminals,  $\mathcal{P}$  is a finite set of preterminals,<sup>1</sup>  $\Sigma$  is a finite set of terminal symbols, and  $\mathcal{R}$  is a finite set of rules of the form,

$$\begin{aligned} S &\rightarrow A, & A &\in \mathcal{N} \\ A &\rightarrow BC, & A &\in \mathcal{N}, B, C \in \mathcal{N} \cup \mathcal{P} \\ T &\rightarrow w, & T &\in \mathcal{P}, w \in \Sigma. \end{aligned}$$

A probabilistic context-free grammar (PCFG) consists of a grammar  $\mathcal{G}$  and rule probabilities  $\pi = \{\pi_r\}_{r \in \mathcal{R}}$  such that  $\pi_r$  is the probability of the rule  $r$ . Letting  $\mathcal{T}_{\mathcal{G}}$  be the set of all parse trees of  $\mathcal{G}$ , a PCFG defines a probability distribution over  $\mathbf{t} \in \mathcal{T}_{\mathcal{G}}$  via  $p_{\pi}(\mathbf{t}) = \prod_{r \in \mathcal{R}} \pi_r$  where  $\mathbf{t}_{\mathcal{R}}$  is the set of rules used in the derivation of  $\mathbf{t}$ . It also defines a distribution over string of terminals  $\mathbf{x} \in \Sigma^*$  via

$$p_{\pi}(\mathbf{x}) = \sum_{\mathbf{t} \in \mathcal{T}_{\mathcal{G}}(\mathbf{x})} p_{\pi}(\mathbf{t}),$$

where  $\mathcal{T}_{\mathcal{G}}(\mathbf{x}) = \{\mathbf{t} \mid \text{yield}(\mathbf{t}) = \mathbf{x}\}$ , i.e. the set of trees  $\mathbf{t}$  such that  $\mathbf{t}$ 's leaves are  $\mathbf{x}$ . We will use  $p_{\pi}(\mathbf{t} \mid \mathbf{x}) \triangleq p_{\pi}(\mathbf{t} \mid \text{yield}(\mathbf{t}) = \mathbf{x})$  to denote the posterior distribution over latent trees given the observed sentence  $\mathbf{x}$ .

**Parameterization** The standard way to parameterize a PCFG is to simply associate a scalar to each rule  $\pi_r$  with the constraint that they form valid probability distributions, i.e. each nonterminal is associated with a fully-parameterized categorical distribution over its rules. This *direct parameterization* is algorithmically convenient since the M-step in the EM algorithm (Dempster et al., 1977) has a closed form. However, there is a long history of work showing that it is difficult to learn meaningful grammars from natural language data with this parameterization (Carroll and

<sup>1</sup>Since we will be inducing a grammar directly from words,  $\mathcal{P}$  is roughly the set of part-of-speech tags and  $\mathcal{N}$  is the set of constituent labels. However, to avoid issues of label alignment, evaluation is only on the tree topology.

Charniak, 1992).<sup>2</sup> Successful approaches to unsupervised parsing have therefore modified the model/learning objective by guiding potentially unrelated rules to behave similarly.

Recognizing that sharing among rule types is beneficial, we propose a *neural parameterization* where rule probabilities are based on distributed representations. We associate embeddings with each symbol, introducing input embeddings  $\mathbf{w}_N$  for each symbol  $N$  on the left side of a rule (i.e.  $N \in \{S\} \cup \mathcal{N} \cup \mathcal{P}$ ). For each rule type  $r$ ,  $\pi_r$  is parameterized as follows,

$$\begin{aligned} \pi_{S \rightarrow A} &= \frac{\exp(\mathbf{u}_A^{\top} f_1(\mathbf{w}_S))}{\sum_{A' \in \mathcal{N}} \exp(\mathbf{u}_{A'}^{\top} f_1(\mathbf{w}_S))}, \\ \pi_{A \rightarrow BC} &= \frac{\exp(\mathbf{u}_{BC}^{\top} \mathbf{w}_A)}{\sum_{B'C' \in \mathcal{M}} \exp(\mathbf{u}_{B'C'}^{\top} \mathbf{w}_A)}, \\ \pi_{T \rightarrow w} &= \frac{\exp(\mathbf{u}_w^{\top} f_2(\mathbf{w}_T))}{\sum_{w' \in \Sigma} \exp(\mathbf{u}_{w'}^{\top} f_2(\mathbf{w}_T))}, \end{aligned}$$

where  $\mathcal{M}$  is the product space  $(\mathcal{N} \cup \mathcal{P}) \times (\mathcal{N} \cup \mathcal{P})$ , and  $f_1, f_2$  are MLPs with two residual layers (see appendix A.1 for the full parameterization). We will use  $\mathbf{E}_{\mathcal{G}} = \{\mathbf{w}_N \mid N \in \{S\} \cup \mathcal{N} \cup \mathcal{P}\}$  to denote the set of input symbol embeddings for a grammar  $\mathcal{G}$ , and  $\lambda$  to refer to the parameters of the neural network used to obtain the rule probabilities. A graphical model-like illustration of the neural PCFG is shown in Figure 1 (left).

It is clear that the neural parameterization does not change the underlying probabilistic assumptions. The difference between the two is analogous to the difference between count-based vs. feed-forward neural language models, where feed-forward neural language models make the same Markov assumptions as the count-based models but are able to take advantage of shared, distributed representations.

## 3 Compound PCFGs

A compound probability distribution (Robbins, 1951) is a distribution whose parameters are themselves random variables. These distributions generalize mixture models to the continuous case, for example in factor analysis which assumes the following generative process,

$$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad \mathbf{x} \sim \mathcal{N}(\mathbf{W}\mathbf{z}, \Sigma).$$

Compound distributions provide the ability to model rich generative processes, but marginalizing over the latent parameter can be computationally expensive unless conjugacy can be exploited.

<sup>2</sup>In preliminary experiments we were indeed unable to learn linguistically meaningful grammars with this PCFG.

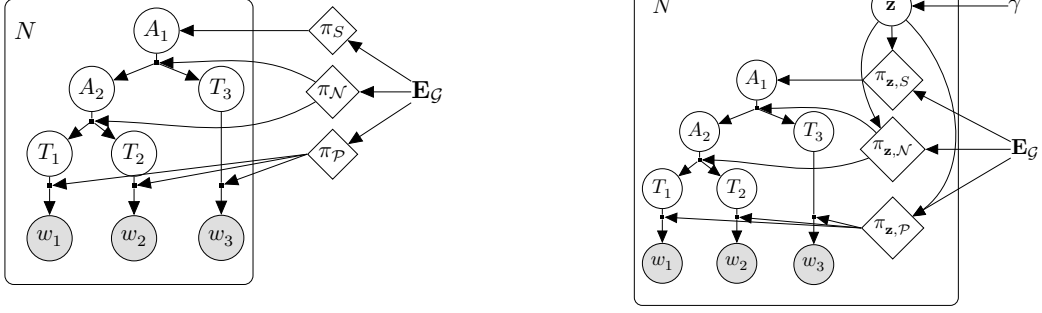


Figure 1: A graphical model-like diagram for the neural PCFG (left) and the compound PCFG (right) for an example tree structure. In the above,  $A_1, A_2 \in \mathcal{N}$  are nonterminals,  $T_1, T_2, T_3 \in \mathcal{P}$  are preterminals,  $w_1, w_2, w_3 \in \Sigma$  are terminals. In the neural PCFG, the global rule probabilities  $\pi = \pi_S \cup \pi_N \cup \pi_P$  are the output from a neural net run over the symbol embeddings  $\mathbf{E}_G$ , where  $\pi_N$  are the set of rules with a nonterminal on the left hand side ( $\pi_S$  and  $\pi_P$  are similarly defined). In the compound PCFG, we have per-sentence rule probabilities  $\pi_z = \pi_{z,S} \cup \pi_{z,N} \cup \pi_{z,P}$  obtained from running a neural net over a random vector  $\mathbf{z}$  (which varies across sentences) and global symbol embeddings  $\mathbf{E}_G$ . In this case, the context-free assumptions hold conditioned on  $\mathbf{z}$ , but they do not hold unconditionally: e.g. when conditioned on  $\mathbf{z}$  and  $A_2$ , the variables  $A_1$  and  $T_1$  are independent; however when conditioned on just  $A_2$ , they are not independent due to the dependence path through  $\mathbf{z}$ . Note that the rule probabilities are random variables in the compound PCFG but deterministic variables in the neural PCFG.

In this work, we study compound probabilistic context free grammars whose distribution over trees arises from the following generative process: we first obtain rule probabilities via

$$\mathbf{z} \sim p_\gamma(\mathbf{z}), \quad \pi_{\mathbf{z}} = f_\lambda(\mathbf{z}, \mathbf{E}_G),$$

where  $p_\gamma(\mathbf{z})$  is a prior with parameters  $\gamma$  (spherical Gaussian in this paper), and  $f_\lambda$  is a neural network that concatenates the input symbol embeddings with  $\mathbf{z}$  and outputs the sentence-level rule probabilities  $\pi_{\mathbf{z}}$ ,

$$\begin{aligned} \pi_{\mathbf{z}, S \rightarrow A} &\propto \exp(\mathbf{u}_A^\top f_1([\mathbf{w}_S; \mathbf{z}])), \\ \pi_{\mathbf{z}, A \rightarrow BC} &\propto \exp(\mathbf{u}_{BC}^\top [\mathbf{w}_A; \mathbf{z}]), \\ \pi_{\mathbf{z}, T \rightarrow w} &\propto \exp(\mathbf{u}_w^\top f_2([\mathbf{w}_T; \mathbf{z}])), \end{aligned}$$

where  $[\mathbf{w}; \mathbf{z}]$  denotes vector concatenation. Then a tree/sentence is sampled from a PCFG with rule probabilities given by  $\pi_{\mathbf{z}}$ ,

$$t \sim \text{PCFG}(\pi_{\mathbf{z}}), \quad x = \text{yield}(t).$$

This can be viewed as a continuous mixture of PCFGs, or alternatively, a Bayesian PCFG with a prior on sentence-level rule probabilities parameterized by  $\mathbf{z}, \lambda, \mathbf{E}_G$ .<sup>3</sup> Importantly, under this generative model the context-free assumptions hold *conditioned on  $\mathbf{z}$* , but they do not hold unconditionally. This is shown in Figure 1 (right) where there is a dependence path through  $\mathbf{z}$  if it is not conditioned upon. Compound PCFGs give rise to a marginal distribution over parse trees  $t$  via

$$p_\theta(t) = \int p(t | \mathbf{z}) p_\gamma(\mathbf{z}) d\mathbf{z},$$

<sup>3</sup>Under the Bayesian PCFG view,  $p_\gamma(\mathbf{z})$  is a distribution over  $\mathbf{z}$  (a subset of the prior), and is thus a hyperprior.

where  $p_\theta(t | \mathbf{z}) = \prod_{r \in \mathcal{R}} \pi_{\mathbf{z}, r}$ . The subscript in  $\pi_{\mathbf{z}, r}$  denotes the fact that the rule probabilities depend on  $\mathbf{z}$ . Compound PCFGs are clearly more expressive than PCFGs as each sentence has its own set of rule probabilities. However, it still assumes a tree-based generative process, making it possible to learn latent tree structures.

Our motivation for the compound PCFG is based on the observation that for grammar induction, first-order context-free assumptions are generally made not because they represent an adequate model of natural language, but because they allow for tractable training.<sup>4</sup> Higher-order PCFGs can introduce dependencies between children and ancestors/siblings through, for example, vertical/horizontal Markovization (Johnson, 1998; Klein and Manning, 2003). However such dependencies complicate training due to the rapid increase in the number of rules. Under this view, we can interpret the compound PCFG as a restricted version of some higher-order PCFG where a child can depend on its ancestors and siblings through a shared latent vector. We hypothesize that this dependence among siblings is especially useful in grammar induction from words, where (for example) if we know that *watched* is used as a verb

<sup>4</sup>A piece of evidence for the misspecification of first-order PCFGs as a statistical model of natural language is that if one pretrains a first-order PCFG on supervised data and continues training with the unsupervised objective (i.e. log marginal likelihood), the resulting grammar deviates significantly from the supervised initial grammar while the log marginal likelihood improves (Johnson et al., 2007). Similar observations have been made for part-of-speech induction with Hidden Markov Models (Merialdo, 1994).

then the noun phrase is likely to be a movie.

In contrast to the usual Bayesian treatment of PCFGs which places priors on global rule probabilities (Kurihara and Sato, 2006; Johnson et al., 2007; Wang and Blunsom, 2013), the compound PCFG assumes a prior on local, sentence-level rule probabilities. It is therefore closely related to the Bayesian grammars studied by Cohen et al. (2009) and Cohen and Smith (2009), who also sample local rule probabilities from a logistic normal prior for training dependency models with valence (DMV) (Klein and Manning, 2004).

**Inference in Compound PCFGs** The expressivity of compound PCFGs comes at a significant challenge in learning and inference. Letting  $\theta = \{\mathbf{E}_G, \lambda\}$  be the parameters of the generative model, we would like to maximize the log marginal likelihood of the observed sentence  $\log p_\theta(\mathbf{x})$ . In the neural PCFG the log marginal likelihood  $\log p_\theta(\mathbf{x}) = \log \sum_{\mathbf{t} \in \mathcal{T}_G(\mathbf{x})} p_\theta(\mathbf{t})$  can be obtained by summing out the latent tree structure using the inside algorithm (Baker, 1979), which is differentiable and thus amenable to gradient-based optimization. In the compound PCFG, the log marginal likelihood is given by

$$\log p_\theta(\mathbf{x}) = \log \left( \int \sum_{\mathbf{t} \in \mathcal{T}_G(\mathbf{x})} p_\theta(\mathbf{t} | \mathbf{z}) p_\gamma(\mathbf{z}) d\mathbf{z} \right).$$

Notice that while the integral over  $\mathbf{z}$  makes this quantity intractable, when we condition on  $\mathbf{z}$ , we can tractably perform the inner summation as before using the inside algorithm. We therefore resort to collapsed amortized variational inference. We first obtain a sample  $\mathbf{z}$  from a variational posterior distribution (given by an amortized inference network), then perform the inner marginalization conditioned on this sample. The evidence lower bound  $\text{ELBO}(\theta, \phi; \mathbf{x})$  is then given by,

$$\mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log p_\theta(\mathbf{x} | \mathbf{z})] - \text{KL}[q_\phi(\mathbf{z} | \mathbf{x}) || p_\gamma(\mathbf{z})],$$

and we can calculate  $p_\theta(\mathbf{x} | \mathbf{z}) = \sum_{\mathbf{t} \in \mathcal{T}_G(\mathbf{x})} p(\mathbf{t} | \mathbf{z})$  with the inside algorithm given a sample  $\mathbf{z}$  from a variational posterior  $q_\phi(\mathbf{z} | \mathbf{x})$ . For the variational family we use a diagonal Gaussian where the mean/log-variance vectors are given by an affine layer over max-pooled hidden states from an LSTM over  $\mathbf{x}$ . We can obtain low-variance estimators for the gradient  $\nabla_{\theta, \phi} \text{ELBO}(\theta, \phi; \mathbf{x})$  by using the reparameterization trick for the expected reconstruction likelihood and the analytical expression for the KL term (Kingma and Welling, 2014).

We remark that under the Bayesian PCFG view, since the parameters of the prior (i.e.  $\theta$ ) are estimated from the data, our approach can be seen as an instance of empirical Bayes (Robbins, 1956).<sup>5</sup>

**MAP Inference** After training, we are interested in comparing the learned trees against an annotated treebank. This requires inferring the most likely tree given a sentence, i.e.  $\text{argmax}_{\mathbf{t}} p_\theta(\mathbf{t} | \mathbf{x})$ . For the neural PCFG we can obtain the most likely tree by using the Viterbi version of the inside algorithm (CKY algorithm). For the compound PCFG, the  $\text{argmax}$  is intractable to obtain exactly, and hence we estimate it with the following approximation,

$$\begin{aligned} & \text{argmax}_{\mathbf{t}} \int p_\theta(\mathbf{t} | \mathbf{x}, \mathbf{z}) p_\theta(\mathbf{z} | \mathbf{x}) d\mathbf{z} \\ & \approx \text{argmax}_{\mathbf{t}} p_\theta(\mathbf{t} | \mathbf{x}, \boldsymbol{\mu}_\phi(\mathbf{x})), \end{aligned}$$

where  $\boldsymbol{\mu}_\phi(\mathbf{x})$  is the mean vector from the inference network. The above approximates the true posterior  $p_\theta(\mathbf{z} | \mathbf{x})$  with  $\delta(\mathbf{z} - \boldsymbol{\mu}_\phi(\mathbf{x}))$ , the Dirac delta function at the mode of the variational posterior.<sup>6</sup> This quantity is tractable to estimate as in the PCFG case. Other approximations are possible: for example we could use  $q_\phi(\mathbf{z} | \mathbf{x})$  as an importance sampling distribution to estimate the first integral. However we found the above approximation to be efficient and effective in practice.

## 4 Experimental Setup

**Data** We test our approach on the Penn Treebank (PTB) (Marcus et al., 1993) with the standard splits (2-21 for training, 22 for validation, 23 for test) and the same preprocessing as in recent works (Shen et al., 2018, 2019), where we discard punctuation, lowercase all tokens, and take the top 10K most frequent words as the vocabulary. This task is more challenging than traditional setups, which usually experiment on shorter sentences and use gold part-of-speech tags. We further experiment on Chinese with version 5.1 of the Chinese Penn Treebank (CTB) (Xue et al., 2005), with the same splits as in Chen and Manning (2014). On CTB we also remove punctuation and keep the top 10K words.

**Hyperparameters** Our PCFG uses 30 nonterminals and 60 preterminals, with 256-dimensional

<sup>5</sup>See Berger (1985) (chapter 4), Zhang (2003), and Cohen (2016) (chapter 3) for further discussion on compound models and empirical Bayes.

<sup>6</sup>Since  $p_\theta(\mathbf{t} | \mathbf{x}, \mathbf{z})$  is continuous with respect to  $\mathbf{z}$ , we have  $\int p_\theta(\mathbf{t} | \mathbf{x}, \mathbf{z}) \delta(\mathbf{z} - \boldsymbol{\mu}_\phi(\mathbf{x})) d\mathbf{z} = p_\theta(\mathbf{t} | \mathbf{x}, \boldsymbol{\mu}_\phi(\mathbf{x}))$ .

symbol embeddings. The compound PCFG uses 64-dimensional latent vectors. The bidirectional LSTM inference network has a single layer with 512 dimensions, and the mean and the log variance vector for  $q_\phi(\mathbf{z} | \mathbf{x})$  are given by max-pooling the hidden states of the LSTM and passing it through an affine layer. Model parameters are initialized with Xavier uniform initialization. For training we use Adam (Kingma and Ba, 2015) with  $\beta_1 = 0.75$ ,  $\beta_2 = 0.999$  and learning rate of 0.001, with a maximum gradient norm limit of 3. We train for 10 epochs with batch size equal to 4. We employ a curriculum learning strategy (Bengio et al., 2009) where we train only on sentences of length up to 30 in the first epoch, and increase this length limit by 1 each epoch. This slightly improved performance and similar strategies have used in the past for grammar induction (Spitkovsky et al., 2012). During training we perform early stopping based on validation perplexity.<sup>7</sup> To mitigate against overfitting to PTB, experiments on CTB utilize the same hyperparameters from PTB.

**Baselines and Evaluation** We observe that even on PTB, there is enough variation in setups across prior work on grammar induction to render a meaningful comparison difficult. Some important dimensions along which prior works vary include, (1) lexicalization: earlier work on grammar induction generally assumed gold (or induced) part-of-speech tags (Klein and Manning, 2004; Smith and Eisner, 2004; Bod, 2006; Snyder et al., 2009), while more recent works induce grammar directly from words (Spitkovsky et al., 2013; Shen et al., 2018); (2) use of punctuation: even within papers that induce a grammar directly from words, some papers employ heuristics based on punctuation as punctuation is usually a strong signal for start/end of constituents (Seginer, 2007; Ponvert et al., 2011; Spitkovsky et al., 2013), some train with punctuation (Jin et al., 2018; Drozdov et al., 2019; Kim et al., 2019), while others discard punctuation altogether for training (Shen et al., 2018, 2019); (3) train/test data: some works do not explicitly separate out train/test sets (Reichart and Rappoport, 2010; Golland et al., 2012) while some do (Huang et al., 2012; Parikh et al., 2014; Htut

<sup>7</sup>However, we used  $F_1$  against validation trees on PTB to select some hyperparameters (e.g. grammar size), as is sometimes done in grammar induction. Hence our PTB results are arguably not fully unsupervised in the strictest sense of the term. The hyperparameters of the PRPN/ON baselines are also tuned using validation  $F_1$  for fair comparison.

et al., 2018). Maintaining train/test splits is less of an issue for unsupervised structure learning, however in this work we follow the latter and separate train/test data. (4) evaluation: for unlabeled  $F_1$ , almost all works ignore punctuation (even approaches that use punctuation during training typically ignore them during evaluation), but there is some variance in discarding trivial spans (width-one and sentence-level spans) and using corpus-level versus sentence-level  $F_1$ .<sup>8</sup> In this paper we discard trivial spans and evaluate on sentence-level  $F_1$  per recent work (Shen et al., 2018, 2019).

Given the above, we mainly compare our approach against two recent, strong baselines with open source code: Parsing Predict Reading Network (PRPN)<sup>9</sup> (Shen et al., 2018) and Ordered Neurons (ON)<sup>10</sup> (Shen et al., 2019). These approaches train a neural language model with gated attention-like mechanisms to induce binary trees, and achieve strong unsupervised parsing performance even when trained on corpora where punctuation is removed. Since the original results were on both language modeling and grammar induction, their hyperparameters were presumably tuned to do well on both and thus may not be optimal for just unsupervised parsing. We therefore tune the hyperparameters of these baselines for unsupervised parsing only (i.e. on validation  $F_1$ ).

## 5 Results and Discussion

Table 1 shows the unlabeled  $F_1$  scores for our models and various baselines. All models soundly outperform right branching baselines, and we find that the neural PCFG/compound PCFG are strong models for grammar induction. In particular the compound PCFG outperforms other models by an appreciable margin on both English and Chinese. We again note that we were unable to induce meaningful grammars through a traditional PCFG with the scalar parameterization.<sup>11</sup> See appendix A.2 for the full results (including corpus-level  $F_1$ ) broken down by sentence length.

Table 2 analyzes the learned tree structures. We compare similarity as measured by  $F_1$  against gold, left, right, and “self” trees (top), where self  $F_1$  score is calculated by averaging over all 6 pairs

<sup>8</sup>Corpus-level  $F_1$  calculates precision/recall at the corpus level to obtain  $F_1$ , while sentence-level  $F_1$  calculates  $F_1$  for each sentence and averages across the corpus.

<sup>9</sup><https://github.com/yikangshen/PRPN>

<sup>10</sup><https://github.com/yikangshen/Ordered-Neurons>

<sup>11</sup>The training perplexity was much higher than in the neural case, indicating significant optimization issues.

Model	PTB		CTB	
	Mean	Max	Mean	Max
PRPN (Shen et al., 2018)	37.4	38.1	—	—
ON (Shen et al., 2019)	47.7	49.4	—	—
URNNG <sup>†</sup> (Kim et al., 2019)	—	45.4	—	—
DIORA <sup>†</sup> (Drozdov et al., 2019)	—	58.9	—	—
Left Branching	8.7		9.7	
Right Branching	39.5		20.0	
Random Trees	19.2	19.5	15.7	16.0
PRPN (tuned)	47.3	47.9	30.4	31.5
ON (tuned)	48.1	50.0	25.4	25.7
Neural PCFG	50.8	52.6	25.7	29.5
Compound PCFG	55.2	60.1	36.0	39.8
Oracle Trees	84.3		81.1	

Table 1: Unlabeled sentence-level  $F_1$  scores on PTB and CTB test sets. Top shows results from previous work while the rest of the results are from this paper. Mean/Max scores are obtained from 4 runs of each model with different random seeds. Oracle is the maximum score obtainable with binarized trees, since we compare against the non-binarized gold trees per convention. Results with <sup>†</sup> are trained on a version of PTB with punctuation, and hence not strictly comparable to the present work. For URNNG/DIORA, we take the parsed test set provided by the authors from their best runs and evaluate  $F_1$  with our evaluation setup.

obtained from 4 different runs. We find that PRPN is particularly consistent across multiple runs. We also observe that different models are better at identifying different constituent labels, as measured by label recall (Table 2, bottom). While left as future work, this naturally suggests an ensemble approach wherein the empirical probabilities of constituents (obtained by averaging the predicted binary constituent labels from the different models) are used either to supervise another model or directly as potentials in a CRF constituency parser. Finally, all models seemed to have some difficulty in identifying SBAR/VP constituents which typically span more words than NP constituents.

**Induced Trees for Downstream Tasks** While the compound PCFG has fewer independence assumptions than the neural PCFG, it is still a more constrained model of language than standard neural language models (NLM) and thus not competitive in terms of perplexity: the compound PCFG obtains a perplexity of 196.3 while an LSTM language model (LM) obtains 86.2 (Table 3).<sup>12</sup> In contrast, both PRPN and ON perform as well as an

<sup>12</sup>We did manage to almost match the perplexity of an NLM by additionally conditioning the terminal probabilities on previous history, i.e.  $\pi_{\mathbf{z}, T \rightarrow w_t} \propto \exp(\mathbf{u}_w^\top f_2([\mathbf{w}_T; \mathbf{z}; \mathbf{h}_t]))$  where  $\mathbf{h}_t$  is the hidden state from an LSTM over  $\mathbf{x}_{<t}$ . However the unsupervised parsing performance was far worse ( $\approx 25 F_1$  on the PTB).

	PRPN	ON	PCFG	Comp. PCFG
Gold	47.3	48.1	50.8	55.2
Left	1.5	14.1	11.8	13.0
Right	39.9	31.0	27.7	28.4
Self	82.3	71.3	65.2	66.8
SBAR	50.0%	51.2%	52.5%	56.1%
NP	59.2%	64.5%	71.2%	74.7%
VP	46.7%	41.0%	33.8%	41.7%
PP	57.2%	54.4%	58.8%	68.8%
ADJP	44.3%	38.1%	32.5%	40.4%
ADVP	32.8%	31.6%	45.5%	52.5%

Table 2: (Top) Mean  $F_1$  similarity against Gold, Left, Right, and Self trees. Self  $F_1$  score is calculated by averaging over all 6 pairs obtained from 4 different runs. (Bottom) Fraction of ground truth constituents that were predicted as a constituent by the models broken down by label (i.e. label recall).

LSTM LM while maintaining good unsupervised parsing performance.

We thus experiment to see if it is possible to use the induced trees to supervise a more flexible generative model that can make use of tree structures—namely, recurrent neural network grammars (RNNG) (Dyer et al., 2016). RNNGs are generative models of language that jointly model syntax and surface structure by incrementally generating a syntax tree and sentence. As with NLMs, RNNGs make no independence assumptions, and have been shown to outperform NLMs in terms of perplexity and grammaticality judgment when trained on gold trees (Kuncoro et al., 2018; Wilcox et al., 2019). We take the best run from each model and parse the training set,<sup>13</sup> and use the induced trees to supervise an RNNG for each model using the parameterization from Kim et al. (2019).<sup>14</sup> We are also interested in syntactic evaluation of our models, and for this we utilize the framework and dataset from Marvin and Linzen (2018), where a model is presented two minimally different sentences such as:

the senators near the assistant are old

\*the senators near the assistant is old

and must assign higher probability to grammatical sentence. Additionally, Kim et al. (2019) report perplexity improvements by fine-tuning an RNNG trained on gold trees with the *unsupervised* RNNG (URNNG)—whereas the RNNG is trained to maximize the joint likelihood  $\log p(\mathbf{x}, \mathbf{t})$ , the URNNG maximizes a lower bound on the log marginal likelihood  $\log \sum_{\mathbf{t}} p(\mathbf{x}, \mathbf{t})$  with a structured inference network that approximates the true

<sup>13</sup>The train/test  $F_1$  was similar for all models.

<sup>14</sup><https://github.com/harvardnlp/urnng>

	PPL	Syntactic Eval.	$F_1$
LSTM LM	86.2	60.9%	—
PRPN	87.1	62.2%	47.9
Induced RNNG	95.3	60.1%	47.8
Induced URNNG	90.1	61.8%	51.6
ON	87.2	61.6%	50.0
Induced RNNG	95.2	61.7%	50.6
Induced URNNG	89.9	61.9%	55.1
Neural PCFG	252.6	49.2%	52.6
Induced RNNG	95.8	68.1%	51.4
Induced URNNG	86.0	69.1%	58.7
Compound PCFG	196.3	50.7%	60.1
Induced RNNG	89.8	70.0%	58.1
Induced URNNG	83.7	76.1%	66.9
RNNG on Oracle Trees	80.6	70.4%	71.9
+ URNNG Fine-tuning	78.3	76.1%	72.8

Table 3: Results from training RNNGs on induced trees from various models (Induced RNNG). Induced URNNG indicates fine-tuning with the URNNG. We show perplexity (PPL), grammaticality judgment performance (Syntactic Eval.), and unlabeled  $F_1$ . PPL/ $F_1$  are on the PTB test set, while Syntactic Eval. is based on the dataset from [Marvin and Linzen \(2018\)](#). Note that the perplexity numbers here are not comparable to standard results on the PTB since our models are generative model of sentences and hence we do not carry information across sentence boundaries.

posterior. We experiment with a similar approach where we fine-tune RNNGs trained on induced trees with URNNGs. We perform early stopping for both RNNG and URNNG based on validation perplexity. See appendix A.3 for further details regarding the experimental setup.

The results are shown in Table 3. For perplexity, RNNGs trained on induced trees (Induced RNNG in Table 3) are unable to improve upon an LSTM LM,<sup>15</sup> in contrast to the supervised RNNG which does outperform the LSTM language model (Table 3, bottom). For grammaticality judgment however, the RNNG trained with compound PCFG trees outperforms the LSTM LM despite obtaining worse perplexity,<sup>16</sup> and performs on par with the RNNG trained on gold trees. Fine-tuning with the URNNG results in improvements in perplexity and grammaticality judgment across the board (Induced URNNG in Table 3). We also obtain large improvements on unsupervised parsing as measured by  $F_1$ , with the fine-tuned URNNGs outperforming the respective original models.<sup>17</sup> This is potentially due to an ensembling effect be-

<sup>15</sup>Under our RNNG parameterization, the LSTM LM is equivalent to an RNNG trained with right branching trees.

<sup>16</sup>[Kuncoro et al. \(2018\)](#) also find that lower perplexity does not always lead to better performance on syntactic evaluation.

<sup>17</sup>[Li et al. \(2019\)](#) similarly obtain improvements by refining a model trained on induced trees on classification tasks.

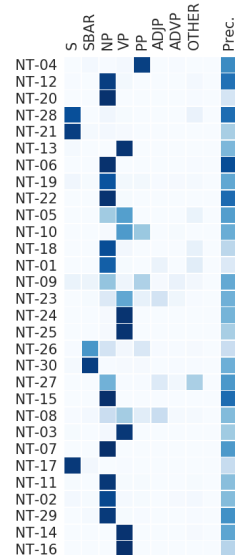


Figure 2: Alignment of induced nonterminals ordered from top based on predicted frequency (therefore NT-04 is the most frequently-predicted nonterminal). For each nonterminal we visualize the proportion of correctly-predicted constituents that correspond to particular gold labels. For reference we also show the precision (i.e. probability of correctly predicting unlabeled constituents) in the rightmost column.

tween the original model and the URNNG’s structured inference network, which is parameterized as a neural CRF constituency parser ([Durrett and Klein, 2015](#); [Liu et al., 2018](#)).<sup>18</sup>

**Model Analysis** We analyze our best compound PCFG model in more detail. Since we induce a full set of nonterminals in our grammar, we can analyze the learned nonterminals to see if they can be aligned with linguistic constituent labels. Figure 2 visualizes the alignment between induced and gold labels, where for each nonterminal we show the empirical probability that a predicted constituent of this type will correspond to a particular linguistic constituent in the test set, conditioned on its being a correct constituent (for reference we also show the precision). We observe that some of the induced nonterminals clearly align to linguistic nonterminals. More detailed results, including preterminal alignments to part-of-speech tags,<sup>19</sup> are shown in appendix A.4.

<sup>18</sup>While left as future work, it is possible to use the compound PCFG itself as an inference network. Also note that the  $F_1$  scores for the URNNGs in Table 3 are optimistic since we selected the best-performing runs of the original models based on validation  $F_1$  to parse the training set.

<sup>19</sup>As a POS induction system, the many-to-one performance of the compound PCFG using the preterminals is 68.0. A similarly-parameterized compound HMM with 60 hidden states (an HMM is a particularly type of PCFG) with 60 states obtains 63.2. This is still quite a bit lower than the state-of-the-art ([Tran et al., 2016](#); [He et al., 2018](#); [Stratos,](#)

<p><b>he retired as senior vice president finance and administration and chief financial officer of the company oct. N</b>  kenneth j. (unk) who was named president of this thrift holding company in august resigned citing personal reasons  the former president and chief executive eric w. (unk) resigned in june  (unk) 's president and chief executive officer john (unk) said the loss stems from several factors  mr. (unk) is executive vice president and chief financial officer of (unk) and will continue in those roles  charles j. lawson jr. N who had been acting chief executive since june N will continue as chairman</p>
<p><b>(unk) corp. received an N million army contract for helicopter engines</b>  boeing co. received a N million air force contract for developing cable systems for the (unk) missile  general dynamics corp. received a N million air force contract for (unk) training sets  grumman corp. received an N million navy contract to upgrade aircraft electronics  thomson missile products with about half british aerospace 's annual revenue include the (unk) (unk) missile family  already british aerospace and french (unk) (unk) (unk) on a british missile contract and on an air-traffic control radar system</p>
<p><b>meanwhile during the the s&amp;p trading halt s&amp;p futures sell orders began (unk) up while stocks in new york kept falling sharply</b>  but the (unk) of s&amp;p futures sell orders weighed on the market and the link with stocks began to fray again  on friday some market makers were selling again traders said  futures traders say the s&amp;p was (unk) that the dow could fall as much as N points  meanwhile two initial public offerings (unk) the (unk) market in their (unk) day of national over-the-counter trading friday  traders said most of their major institutional investors on the other hand sat tight</p>

Table 4: For each query sentence (bold), we show the 5 nearest neighbors based on cosine similarity, where we take the representation for each sentence to be the mean of the variational posterior.

We next analyze the continuous latent space. Table 4 shows nearest neighbors of some sentences using the mean of the variational posterior as the continuous representation of each sentence. We qualitatively observe that the latent space seems to capture topical information. We are also interested in the variation in the leaves due to  $\mathbf{z}$  when the variation due to the tree structure is held constant. To investigate this, we use the parsed dataset to obtain pairs of the form  $(\boldsymbol{\mu}_\phi(\mathbf{x}^{(n)}), \mathbf{t}_j^{(n)})$ , where  $\mathbf{t}_j^{(n)}$  is the  $j$ -th subtree of the (approximate) MAP tree  $\mathbf{t}^{(n)}$  for the  $n$ -th sentence. Therefore each mean vector  $\boldsymbol{\mu}_\phi(\mathbf{x}^{(n)})$  is associated with  $|\mathbf{x}^{(n)}| - 1$  subtrees, where  $|\mathbf{x}^{(n)}|$  is the sentence length. Our definition of subtree here ignores terminals, and thus each subtree is associated with many mean vectors. For a frequently occurring subtree, we perform PCA on the set of mean vectors that are associated with the subtree to obtain the top principal component. We then show the constituents that had the 5 most positive/negative values for this top principal component in Table 5. For example, a particularly common subtree—associated with 180 unique constituents—is given by

(NT-04 (T-13  $w_1$ ) (NT-12 (NT-20 (NT-20 (NT-07 (T-05  $w_2$ ) (T-45  $w_3$ )) (T-35  $w_4$ )) (T-40  $w_5$ )) (T-22  $w_6$ ))).

The top 5 constituents with the most negative/positive values are shown in the top left part of Table 5. We find that the leaves  $[w_1, \dots, w_6]$ , which form a 6-word constituent, vary in a regular manner as  $\mathbf{z}$  is varied. We also observe that root of this subtree (NT-04) aligns to prepositional phrases (PP) in Figure 2, and the leaves in Table 5 (top left) are indeed mostly PP. However, the

2019), though comparison is confounded by various factors such as preprocessing (e.g. we drop punctuation). A neural PCFG/HMM obtains 68.2 and 63.4 respectively.

model fails to identify ((T-40  $w_5$ ) (T-22  $w_6$ )) as a constituent in this case (as well as well in the bottom right example). See appendix A.5 for more examples. It is possible that the model is utilizing the subtrees to capture broad template-like structures and then using  $\mathbf{z}$  to fill them in, similar to recent works that also train models to separate “what to say” from “how to say it” (Wiseman et al., 2018; Peng et al., 2019; Chen et al., 2019a,b).

**Limitations** We report on some negative results as well as important limitations of our work. While distributed representations promote parameter sharing, we were unable to obtain improvements through more factorized parameterizations that promote even greater parameter sharing. In particular, for rules of the type  $A \rightarrow BC$ , we tried having the output embeddings be a function of the input embeddings (e.g.  $\mathbf{u}_{BC} = g([\mathbf{w}_B; \mathbf{w}_C])$  where  $g$  is an MLP), but obtained worse results. For rules of the type  $T \rightarrow w$ , we tried using a character-level CNN (dos Santos and Zadrozny, 2014; Kim et al., 2016) to obtain the output word embeddings  $\mathbf{u}_w$  (Jozefowicz et al., 2016; Tran et al., 2016), but found the performance to be similar to the word-level case.<sup>20</sup> We were also unable to obtain improvements through normalizing flows (Rezende and Mohamed, 2015; Kingma et al., 2016). However, given that we did not exhaustively explore the full space of possible parameterizations, the above modifications could eventually lead to improvements with the right setup. Relatedly, the models were quite sensitive to parameterization (e.g. it was important to use residual layers for  $f_1, f_2$ ), grammar size, and optimization method. Finally, despite vectorized GPU im-

<sup>20</sup>It is also possible to take advantage of pretrained word embeddings by using them to initialize output word embeddings or directly working with continuous emission distributions (Lin et al., 2015; He et al., 2018)



	<p><b>PC -</b> of the company 's capital structure in the company 's divestiture program by the company 's new board in the company 's core businesses on the company 's strategic plan</p> <p><b>PC +</b> above the treasury 's N-year note above the treasury 's seven-year note above the treasury 's comparable note above the treasury 's five-year note measured the earth 's ozone layer</p>		<p><b>PC -</b> purchased through the exercise of stock options circulated by a handful of major brokers higher as a percentage of total loans common with a lot of large companies surprised by the storm of sell orders</p> <p><b>PC +</b> brought to the u.s. against her will laid for the arrest of opposition activists uncertain about the magnitude of structural damage held after the assassination of his mother hurt as a result of the violations</p>
	<p><b>PC -</b> to terminate their contract with warner to support a coup in panama to suit the bureaucrats in brussels to thwart his bid for amr to prevent the pound from rising</p> <p><b>PC +</b> to change our strategy of investing to offset the growth of minimills to be a lot of art to change our way of life to increase the impact of advertising</p>		<p><b>PC -</b> raise the minimum grant for smaller states veto a defense bill with inadequate funding avoid an imminent public or private injury field a competitive slate of congressional candidates alter a longstanding ban on such involvement</p> <p><b>PC +</b> generate an offsetting profit by selling waves change an export loss to domestic plus expect any immediate problems with margin calls make a positive contribution to our earnings find a trading focus discouraging much participation</p>

Table 5: For each subtree, we perform PCA on the variational posterior mean vectors that are associated with that particular subtree and take the top principal component. We list the top 5 constituents that had the lowest (**PC -**) and highest (**PC +**) principal component values.

plementations, training was significantly more expensive (both in terms of time and memory) than NLM-based grammar induction systems due to the  $O(|\mathcal{R}||x|^3)$  dynamic program, which makes our approach potentially difficult to scale.

## 6 Related Work

Grammar induction has a long and rich history in natural language processing. Early work on grammar induction with pure unsupervised learning was mostly negative (Lari and Young, 1990; Carroll and Charniak, 1992; Charniak, 1993), though Pereira and Schabes (1992) reported some success on partially bracketed data. Clark (2001) and Klein and Manning (2002) were some of the first successful statistical approaches to grammar induction. In particular, the constituent-context model (CCM) of Klein and Manning (2002), which explicitly models both constituents and disjuncts, was the basis for much subsequent work (Klein and Manning, 2004; Huang et al., 2012; Golland et al., 2012). Other works have explored imposing inductive biases through Bayesian priors (Johnson et al., 2007; Liang et al., 2007; Wang and Blunsom, 2013), modified objectives (Smith and Eisner, 2004), and additional constraints on recursion depth (Noji et al., 2016; Jin et al., 2018).

While the framework of specifying the structure of a grammar and learning the parameters is common, other methods exist. Bod (2006) consider a nonparametric-style approach to unsupervised parsing by using random subsets of training subtrees to parse new sentences. Seginer (2007) utilize an incremental algorithm to unsupervised parsing which makes local decisions to create con-

stituents based on a complex set of heuristics. Ponvert et al. (2011) induce parse trees through cascaded applications of finite state models.

More recently, neural network-based approaches to grammar induction have shown promising results on inducing parse trees directly from words. In particular, Shen et al. (2018, 2019) learn tree structures through gated mechanisms within hidden layers of neural language models, while Drozdov et al. (2019) combine recursive autoencoders with the inside-outside algorithm. Kim et al. (2019) train unsupervised recurrent neural network grammars with a structured inference network to induce latent trees.

## 7 Conclusion

This work explores grammar induction with compound PCFGs, which modulate rule probabilities with per-sentence continuous latent vectors. The latent vector induces marginal dependencies beyond the traditional first-order context-free assumptions within a tree-based generative process, leading to improved performance. The collapsed amortized variational inference approach is general and can be used for generative models which admit tractable inference through partial conditioning. Learning deep generative models which exhibit such conditional Markov properties is an interesting direction for future work.

**Acknowledgments** We thank Phil Blunsom for initial discussions, Yonatan Belinkov and Shay Cohen for helpful feedback, and Andrew Drozdov for the DIORA dataset. YK is supported by a Google Fellowship. AMR acknowledges the support of NSF 1704834, 1845664, AWS, and Oracle.

## References

- Sanjeev Arora, Nadav Cohen, and Elad Hazan. 2018. On the Optimization of Deep Networks: Implicit Acceleration by Overparameterization. In *Proceedings of ICML*.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. In *Proceedings of NIPS*.
- James K. Baker. 1979. Trainable Grammars for Speech Recognition. In *Proceedings of the Spring Conference of the Acoustical Society of America*.
- Yoshua Bengio, Jerome Louradour, Ronan Collobert, and Jason Weston. 2009. Curriculum Learning. In *Proceedings of ICML*.
- James O. Berger. 1985. *Statistical Decision Theory and Bayesian Analysis*. Springer.
- Rens Bod. 2006. An All-Subtrees Approach to Unsupervised Parsing. In *Proceedings of ACL*.
- Glenn Carroll and Eugene Charniak. 1992. Two Experiments on Learning Probabilistic Dependency Grammars from Corpora. In *AAAI Workshop on Statistically-Based NLP Techniques*.
- Eugene Charniak. 1993. *Statistical Language Learning*. MIT Press.
- Danqi Chen and Christopher D. Manning. 2014. A Fast and Accurate Dependency Parser using Neural Networks. In *Proceedings of EMNLP*.
- Mingda Chen, Qingming Tang, Sam Wiseman, and Kevin Gimpel. 2019a. Controllable Paraphrase Generation with a Syntactic Exemplar. In *Proceedings of ACL*.
- Mingda Chen, Qingming Tang, Sam Wiseman, and Kevin Gimpel. 2019b. A Multi-task Approach for Disentangling Syntax and Semantics in Sentence Representations. In *Proceedings of NAACL*.
- Alexander Clark. 2001. Unsupervised Induction of Stochastic Context Free Grammars Using Distributional Clustering. In *Proceedings of CoNLL*.
- Shay B. Cohen. 2016. *Bayesian Analysis in Natural Language Processing*. Morgan and Claypool.
- Shay B. Cohen, Kevin Gimpel, and Noah A Smith. 2009. Logistic Normal Priors for Unsupervised Probabilistic Grammar Induction. In *Proceedings of NIPS*.
- Shay B. Cohen and Noah A Smith. 2009. Shared Logistic Normal Distributions for Soft Parameter Tying in Unsupervised Grammar Induction. In *Proceedings of NAACL*.
- Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. 1977. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38.
- Andrew Drozdov, Patrick Verga, Mohit Yadev, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised Latent Tree Induction with Deep Inside-Outside Recursive Auto-Encoders. In *Proceedings of NAACL*.
- Simon S. Du, Xiyu Zhai, Barnabas Póczos, and Aarti Singh. 2019. Gradient Descent Provably Optimizes Over-parameterized Neural Networks. In *Proceedings of ICLR*.
- Greg Durrett and Dan Klein. 2015. Neural CRF Parsing. In *Proceedings of ACL*.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent Neural Network Grammars. In *Proceedings of NAACL*.
- Dave Golland, John DeNero, and Jakob Uszkoreit. 2012. A Feature-Rich Constituent Context Model for Grammar Induction. In *Proceedings of ACL*.
- Junxian He, Graham Neubig, and Taylor Berg-Kirkpatrick. 2018. Unsupervised Learning of Syntactic Structure with Invertible Neural Projections. In *Proceedings of EMNLP*.
- Phu Mon Htut, Kyunghyun Cho, and Samuel R. Bowman. 2018. Grammar Induction with Neural Language Models: An Unusual Replication. In *Proceedings of EMNLP*.
- Yun Huang, Min Zhang, and Chew Lim Tan. 2012. Improved Constituent Context Model with Features. In *Proceedings of PACLIC*.
- Lifeng Jin, Finale Doshi-Velez, Timothy Miller, William Schuler, and Lane Schwartz. 2018. Unsupervised Grammar Induction with Depth-bounded PCFG. In *Proceedings of TACL*.
- Mark Johnson. 1998. PCFG Models of Linguistic Tree Representations. *Computational Linguistics*, 24:613–632.
- Mark Johnson, Thomas L. Griffiths, and Sharon Goldwater. 2007. Bayesian Inference for PCFGs via Markov chain Monte Carlo. In *Proceedings of NAACL*.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the Limits of Language Modeling. *arXiv:1602.02410*.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-Aware Neural Language Models. In *Proceedings of AAAI*.
- Yoon Kim, Alexander M. Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019. Unsupervised Recurrent Neural Network Grammars. In *Proceedings of NAACL*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of ICLR*.

- Diederik P. Kingma, Tim Salimans, and Max Welling. 2016. Improving Variational Inference with Autoregressive Flow. *arXiv:1606.04934*.
- Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *Proceedings of ICLR*.
- Nikita Kitaev and Dan Klein. 2018. Constituency Parsing with a Self-Attentive Encoder. In *Proceedings of ACL*.
- Dan Klein and Christopher Manning. 2002. A Generative Constituent-Context Model for Improved Grammar Induction. In *Proceedings of ACL*.
- Dan Klein and Christopher Manning. 2004. Corpus-based Induction of Syntactic Structure: Models of Dependency and Constituency. In *Proceedings of ACL*.
- Dan Klein and Christopher D. Manning. 2003. Accurate Unlexicalized Parsing. In *Proceedings of ACL*.
- Adhiguna Kuncoro, Chris Dyer, John Hale, Dani Yogatama, Stephen Clark, and Phil Blunsom. 2018. LSTMs Can Learn Syntax-Sensitive Dependencies Well, But Modeling Structure Makes Them Better. In *Proceedings of ACL*.
- Kenichi Kurihara and Taisuke Sato. 2006. Variational Bayesian Grammar Induction for Natural Language. In *Proceedings of International Colloquium on Grammatical Inference*.
- Karim Lari and Steve Young. 1990. The Estimation of Stochastic Context-Free Grammars Using the Inside-Outside Algorithm. *Computer Speech and Language*, 4:35–56.
- Bowen Li, Lili Mou, and Frank Keller. 2019. An Imitation Learning Approach to Unsupervised Parsing. In *Proceedings of ACL*.
- Percy Liang, Slav Petrov, Michael I. Jordan, and Dan Klein. 2007. The Infinite PCFG using Hierarchical Dirichlet Processes. In *Proceedings of EMNLP*.
- Chu-Cheng Lin, Waleed Ammar, Chris Dyer, , and Lori Levin. 2015. Unsupervised POS Induction with Word Embeddings. In *Proceedings of NAACL*.
- Yang Liu, Matt Gardner, and Mirella Lapata. 2018. Structured Alignment Networks for Matching Sentences. In *Proceedings of EMNLP*.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a Large Annotated Corpus of English: The Penn Treebank. *Computational Linguistics*, 19:313–330.
- Rebecca Marvin and Tal Linzen. 2018. Targeted Syntactic Evaluation of Language Models. In *Proceedings of EMNLP*.
- Bernard Merialdo. 1994. Tagging English Text with a Probabilistic Model. *Computational Linguistics*, 20(2):155–171.
- Hiroshi Noji, Yusuke Miyao, and Mark Johnson. 2016. Using Left-corner Parsing to Encode Universal Structural Constraints in Grammar Induction. In *Proceedings of EMNLP*.
- Ankur P. Parikh, Shay B. Cohen, and Eric P. Xing. 2014. Spectral Unsupervised Parsing with Additive Tree Metrics. In *Proceedings of ACL*.
- Hao Peng, Ankur P. Parikh, Manaal Faruqui, Bhuwan Dhingra, and Dipanjan Das. 2019. Text Generation with Exemplar-based Adaptive Decoding. In *Proceedings of NAACL*.
- Fernando Pereira and Yves Schabes. 1992. Inside-Outside Reestimation from Partially Bracketed Corpora. In *Proceedings of ACL*.
- Elis Ponvert, Jason Baldridge, and Katrin Erk. 2011. Simplified Unsupervised Grammar Induction from Raw Text with Cascaded Finite State Methods. In *Proceedings of ACL*.
- Ofir Press and Lior Wolf. 2016. Using the Output Embedding to Improve Language Models. In *Proceedings of EACL*.
- Roi Reichart and Ari Rappoport. 2010. Improved Fully Unsupervised Parsing with Zoomed Learning. In *Proceedings of EMNLP*.
- Danilo J. Rezende and Shakir Mohamed. 2015. Variational Inference with Normalizing Flows. In *Proceedings of ICML*.
- Danilo J. Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of ICML*.
- Herbert Robbins. 1951. Asymptotically Subminimax Solutions of Compound Statistical Decision Problems. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 131–149. Berkeley: University of California Press.
- Herbert Robbins. 1956. An Empirical Bayes Approach to Statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, pages 157–163. Berkeley: University of California Press.
- Cícero Nogueira dos Santos and Bianca Zadrozny. 2014. Learning Character-level Representations for Part-of-Speech Tagging. In *Proceedings of ICML*.
- Yoav Seginer. 2007. Fast Unsupervised Incremental Parsing. In *Proceedings of ACL*.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. 2018. Neural Language Modeling by Jointly Learning Syntax and Lexicon. In *Proceedings of ICLR*.

- Yikang Shen, Shawn Tan, Alessandro Sordani, and Aaron Courville. 2019. Ordered Neurons: Integrating Tree Structures into Recurrent Neural Networks. In *Proceedings of ICLR*.
- Noah A. Smith and Jason Eisner. 2004. Annealing Techniques for Unsupervised Statistical Language Learning. In *Proceedings of ACL*.
- Benjamin Snyder, Tahira Naseem, and Regina Barzilay. 2009. Unsupervised Multilingual Grammar Induction. In *Proceedings of ACL*.
- Valentin I. Spitkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2012. Three Dependency-and-Boundary Models for Grammar Induction. In *Proceedings of EMNLP-CoNLL*.
- Valentin I. Spitkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2013. Breaking Out of Local Optima with Count Transforms and Model Recombination: A Study in Grammar Induction. In *Proceedings of EMNLP*.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017. A Minimal Span-Based Neural Constituency Parser. In *Proceedings of ACL*.
- Karl Stratos. 2019. Mutual Information Maximization for Simple and Accurate Part-of-Speech Induction. In *Proceedings of NAACL*.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. In *Proceedings of ACL*.
- Ke Tran, Yonatan Bisk, Ashish Vaswani, Daniel Marcu, and Kevin Knight. 2016. Unsupervised Neural Hidden Markov Models. In *Proceedings of the Workshop on Structured Prediction for NLP*.
- Pengyu Wang and Phil Blunsom. 2013. Collapsed Variational Bayesian Inference for PCFGs. In *Proceedings of CoNLL*.
- Wenhui Wang and Baobao Chang. 2016. Graph-based Dependency Parsing with Bidirectional LSTM. In *Proceedings of ACL*.
- Ethan Wilcox, Peng Qian, Richard Futrell, Miguel Ballesteros, and Roger Levy. 2019. Structural Supervision Improves Learning of Non-Local Grammatical Dependencies. In *Proceedings of NAACL*.
- Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2018. Learning Neural Templates for Text Generation. In *Proceedings of EMNLP*.
- Ji Xu, Daniel Hsu, and Arian Maleki. 2018. Benefits of Over-Parameterization with EM. In *Proceedings of NeurIPS*.
- Naiwen Xue, Fei Xia, Fu dong Chiou, and Marta Palmer. 2005. The Penn Chinese Treebank: Phrase Structure Annotation of a Large Corpus. *Natural Language Engineering*, 11:207–238.
- Cun-Hui Zhang. 2003. Compound Decision Theory and Empirical Bayes Methods. *The Annals of Statistics*, 31:379–390.
- Xiaodan Zhu, Parinaz Sobhani, and Hongyu Guo. 2015. Long Short-Term Memory Over Tree Structures. In *Proceedings of ICML*.

## A Appendix

### A.1 Model Parameterization

**Neural PCFG** We associate an input embedding  $\mathbf{w}_N$  for each symbol  $N$  on the left side of a rule (i.e.  $N \in \{S\} \cup \mathcal{N} \cup \mathcal{P}$ ) and run a neural network over  $\mathbf{w}_N$  to obtain the rule probabilities. Concretely, each rule type  $\pi_r$  is parameterized as follows,

$$\begin{aligned}\pi_{S \rightarrow A} &= \frac{\exp(\mathbf{u}_A^\top f_1(\mathbf{w}_S))}{\sum_{A' \in \mathcal{N}} \exp(\mathbf{u}_{A'}^\top f_1(\mathbf{w}_S))}, \\ \pi_{A \rightarrow BC} &= \frac{\exp(\mathbf{u}_{BC}^\top \mathbf{w}_A)}{\sum_{B'C' \in \mathcal{M}} \exp(\mathbf{u}_{B'C'}^\top \mathbf{w}_A)}, \\ \pi_{T \rightarrow w} &= \frac{\exp(\mathbf{u}_w^\top f_2(\mathbf{w}_T))}{\sum_{w' \in \Sigma} \exp(\mathbf{u}_{w'}^\top f_2(\mathbf{w}_T))},\end{aligned}$$

where  $\mathcal{M}$  is the product space  $(\mathcal{N} \cup \mathcal{P}) \times (\mathcal{N} \cup \mathcal{P})$ , and  $f_1, f_2$  are MLPs with two residual layers,

$$\begin{aligned}f_i(\mathbf{x}) &= g_{i,1}(g_{i,2}(\mathbf{W}_i \mathbf{x})), \\ g_{i,j}(\mathbf{y}) &= \text{ReLU}(\mathbf{V}_{i,j} \text{ReLU}(\mathbf{U}_{i,j} \mathbf{y})) + \mathbf{y}.\end{aligned}$$

The bias terms for the above expressions (including for the rule probabilities) are omitted for notational brevity. In Figure 1 we use the following to refer to rule probabilities of different rule types,

$$\begin{aligned}\pi_S &= \{\pi_r \mid r \in L(S)\}, \\ \pi_{\mathcal{N}} &= \{\pi_r \mid r \in L(A), A \in \mathcal{N}\}, \\ \pi_{\mathcal{P}} &= \{\pi_r \mid r \in L(T), T \in \mathcal{P}\}, \\ \boldsymbol{\pi} &= \pi_S \cup \pi_{\mathcal{N}} \cup \pi_{\mathcal{P}},\end{aligned}$$

where  $L(A)$  denotes the set of rules with  $A$  on the left hand side.

**Compound PCFG** The compound PCFG rule probabilities  $\pi_{\mathbf{z}}$  given a latent vector  $\mathbf{z}$ ,

$$\begin{aligned}\pi_{\mathbf{z}, S \rightarrow A} &= \frac{\exp(\mathbf{u}_A^\top f_1([\mathbf{w}_S; \mathbf{z}]))}{\sum_{A' \in \mathcal{N}} \exp(\mathbf{u}_{A'}^\top f_1([\mathbf{w}_S; \mathbf{z}]))}, \\ \pi_{\mathbf{z}, A \rightarrow BC} &= \frac{\exp(\mathbf{u}_{BC}^\top [\mathbf{w}_A; \mathbf{z}])}{\sum_{B'C' \in \mathcal{M}} \exp(\mathbf{u}_{B'C'}^\top [\mathbf{w}_A; \mathbf{z}])}, \\ \pi_{\mathbf{z}, T \rightarrow w} &= \frac{\exp(\mathbf{u}_w^\top f_2([\mathbf{w}_T; \mathbf{z}]))}{\sum_{w' \in \Sigma} \exp(\mathbf{u}_{w'}^\top f_2([\mathbf{w}_T; \mathbf{z}]))}.\end{aligned}$$

Again the bias terms are omitted for brevity, and  $f_1, f_2$  are as before where the first layer’s input dimensions are appropriately changed to account for concatenation with  $\mathbf{z}$ .

### A.2 Corpus/Sentence $F_1$ by Sentence Length

For completeness we show the corpus-level and sentence-level  $F_1$  broken down by sentence length in Table 6, averaged across 4 different runs of each model.

### A.3 Experiments with RNNs

For experiments on supervising RNNs with induced trees, we use the parameterization and hyperparameters from Kim et al. (2019), which uses a 2-layer 650-dimensional stack LSTM (with dropout of 0.5) and a 650-dimensional tree LSTM (Tai et al., 2015; Zhu et al., 2015) as the composition function.

Concretely, the generative story is as follows: first, the stack representation is used to predict the next action (SHIFT or REDUCE) via an affine transformation followed by a sigmoid. If SHIFT is chosen, we obtain a distribution over the vocabulary via another affine transformation over the stack representation followed by a softmax. Then we sample the next word from this distribution and shift the generated word onto the stack using the stack LSTM. If REDUCE is chosen, we pop the last two elements off the stack and use the tree LSTM to obtain a new representation. This new representation is shifted onto the stack via the stack LSTM. Note that this RNNG parameterization is slightly different than the original from Dyer et al. (2016), which does not ignore constituent labels and utilizes a bidirectional LSTM as the composition function instead of a tree LSTM. As our RNNG parameterization only works with binary trees, we binarize the gold trees with right binarization for the RNNG trained on gold trees (trees from the unsupervised methods explored in this paper are already binary). The RNNG also trains a discriminative parser alongside the generative model for evaluation with importance sampling. We use a CRF parser whose span score parameterization is similar similar to recent works (Wang and Chang, 2016; Stern et al., 2017; Kitaev and Klein, 2018): position embeddings are added to word embeddings, and a bidirectional LSTM with 256 hidden dimensions is run over the input representations to obtain the forward and backward hidden states. The score  $s_{ij} \in \mathbb{R}$  for a constituent spanning the  $i$ -th and  $j$ -th word is given by,

$$s_{ij} = \text{MLP}([\vec{\mathbf{h}}_{j+1} - \vec{\mathbf{h}}_i; \overleftarrow{\mathbf{h}}_{i-1} - \overleftarrow{\mathbf{h}}_j]),$$

where the MLP has a single hidden layer with

	Sentence-level $F_1$				
	WSJ-10	WSJ-20	WSJ-30	WSJ-40	WSJ-Full
Left Branching	17.4	12.9	9.9	8.6	8.7
Right Branching	58.5	49.8	44.4	41.6	39.5
Random Trees	31.8	25.2	21.5	19.7	19.2
PRPN (tuned)	58.4	54.3	50.9	48.5	47.3
ON (tuned)	63.9	57.5	53.2	50.5	48.1
Neural PCFG	64.6	58.1	54.6	52.6	50.8
Compound PCFG	70.5	63.4	58.9	56.6	55.2
Oracle	82.1	84.1	84.2	84.3	84.3

	Corpus-level $F_1$				
	WSJ-10	WSJ-20	WSJ-30	WSJ-40	WSJ-Full
Left Branching	16.5	11.7	8.5	7.2	6.0
Right Branching	58.9	48.3	42.5	39.4	36.1
Random Trees	31.9	23.9	20.0	18.1	16.4
PRPN (tuned)	59.3	53.6	49.7	46.9	44.5
ON (tuned)	64.7	56.3	51.5	48.3	45.6
Neural PCFG	63.5	56.8	53.1	51.0	48.7
Compound PCFG	70.6	62.0	57.1	54.6	52.4
Oracle	83.5	85.2	84.9	84.9	84.7

Table 6: Average unlabeled  $F_1$  for the various models broken down by sentence length on the PTB test set. For example WSJ-10 refers to  $F_1$  calculated on the subset of the test set where the maximum sentence length is at most 10. Scores are averaged across 4 runs of the model with different random seeds. Oracle is the performance of binarized gold trees (with right branching binarization). Top shows sentence-level  $F_1$  and bottom shows corpus-level  $F_1$ .

ReLU nonlinearity followed by layer normalization (Ba et al., 2016).

For experiments on fine-tuning the RNNG with the unsupervised RNNG, we take the discriminative parser (which is also pretrained alongside the RNNG on induced trees) to be the structured inference network for optimizing the evidence lower bound. We refer the reader to Kim et al. (2019) and their open source implementation<sup>21</sup> for additional details. We also observe that as noted by Kim et al. (2019), a URNNG trained from scratch on this version of PTB without punctuation failed to outperform a right-branching baseline.

The LSTM language model baseline is the same size as the stack LSTM (i.e. 2 layers, 650 hidden units, dropout of 0.5), and is therefore equivalent to an RNNG with completely right branching trees. For all models we share input/output word embeddings (Press and Wolf, 2016). Perplexity estimation for the RNNGs and the compound PCFG uses 1000 importance-weighted samples.

For grammaticality judgment, we modify the publicly available dataset from Marvin and Linzen (2018)<sup>22</sup> to only keep sentence pairs that did not have any unknown words with respect to our PTB

vocabulary of 10K words. This results in 33K sentence pairs for evaluation.

#### A.4 Nonterminal/Preterminal Alignments

Figure 3 shows the part-of-speech alignments and Table 7 shows the nonterminal label alignments for the compound PCFG/neural PCFG.

#### A.5 Subtree Analysis

Table 8 lists more examples of constituents within each subtree as the top principal component is varied. Due to data sparsity, the subtree analysis is performed on the full dataset. See section 5 for more details. See section 5 for more details.

<sup>21</sup><https://github.com/harvardnlp/urnng>

<sup>22</sup>[https://github.com/BeckyMarvin/LM\\_syneval](https://github.com/BeckyMarvin/LM_syneval)

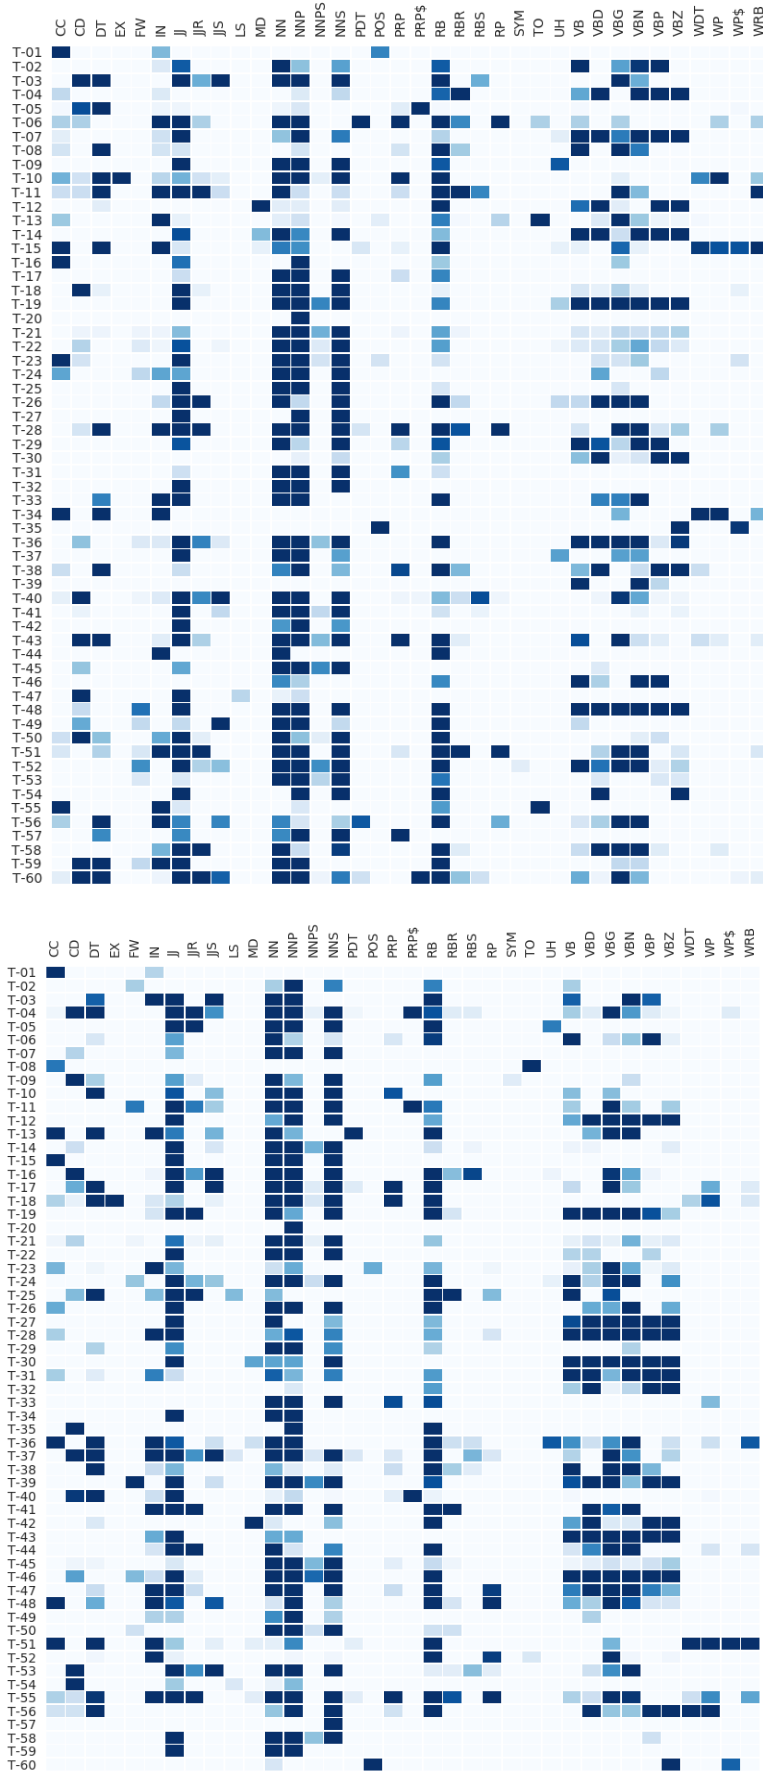


Figure 3: Preterminal alignment to part-of-speech tags for the compound PCFG (top) and the neural PCFG (bottom).

Label	S	SBAR	NP	VP	PP	ADJP	ADVP	Other	Freq.	Acc.
NT-01	0.0%	0.0%	81.8%	1.1%	0.0%	5.9%	0.0%	11.2%	2.9%	13.8%
NT-02	2.2%	0.9%	90.8%	1.7%	0.9%	0.0%	1.3%	2.2%	1.1%	44.0%
NT-03	1.0%	0.0%	2.3%	96.8%	0.0%	0.0%	0.0%	0.0%	1.8%	37.1%
NT-04	0.3%	2.2%	0.5%	2.0%	93.9%	0.2%	0.6%	0.3%	11.0%	64.9%
NT-05	0.2%	0.0%	36.4%	56.9%	0.0%	0.0%	0.2%	6.2%	3.1%	57.1%
NT-06	0.0%	0.0%	99.1%	0.0%	0.1%	0.0%	0.2%	0.6%	5.2%	89.0%
NT-07	0.0%	0.0%	99.7%	0.0%	0.3%	0.0%	0.0%	0.0%	1.3%	59.3%
NT-08	0.5%	2.2%	23.3%	35.6%	11.3%	23.6%	1.7%	1.7%	2.0%	44.3%
NT-09	6.3%	5.6%	40.2%	4.3%	32.6%	1.2%	7.0%	2.8%	2.6%	52.1%
NT-10	0.1%	0.1%	1.4%	58.8%	38.6%	0.0%	0.8%	0.1%	3.0%	50.5%
NT-11	0.9%	0.0%	96.5%	0.9%	0.9%	0.0%	0.0%	0.9%	1.1%	42.9%
NT-12	0.5%	0.2%	94.4%	2.4%	0.2%	0.1%	0.2%	2.0%	8.9%	74.9%
NT-13	1.6%	0.1%	0.2%	97.7%	0.2%	0.1%	0.1%	0.1%	6.2%	46.0%
NT-14	0.0%	0.0%	0.0%	98.6%	0.0%	0.0%	0.0%	1.4%	0.9%	54.1%
NT-15	0.0%	0.0%	99.7%	0.0%	0.3%	0.0%	0.0%	0.0%	2.0%	76.9%
NT-16	0.0%	0.0%	0.0%	100.0%	0.0%	0.0%	0.0%	0.0%	0.3%	29.9%
NT-17	96.4%	2.9%	0.0%	0.7%	0.0%	0.0%	0.0%	0.0%	1.2%	24.4%
NT-18	0.3%	0.0%	88.7%	2.8%	0.3%	0.0%	0.0%	7.9%	3.0%	28.3%
NT-19	3.9%	1.0%	86.6%	2.4%	2.6%	0.4%	1.3%	1.8%	4.5%	53.4%
NT-20	0.0%	0.0%	99.0%	0.0%	0.0%	0.3%	0.2%	0.5%	7.4%	17.5%
NT-21	94.4%	1.7%	2.0%	1.4%	0.3%	0.1%	0.0%	0.1%	6.2%	34.7%
NT-22	0.1%	0.0%	98.4%	1.1%	0.1%	0.0%	0.2%	0.2%	3.5%	77.6%
NT-23	0.4%	0.9%	14.0%	53.1%	8.2%	18.5%	4.3%	0.7%	2.4%	49.1%
NT-24	0.0%	0.2%	1.5%	98.3%	0.0%	0.0%	0.0%	0.0%	2.3%	47.3%
NT-25	0.3%	0.0%	1.4%	98.3%	0.0%	0.0%	0.0%	0.0%	2.2%	34.6%
NT-26	0.4%	60.7%	18.4%	3.0%	15.4%	0.4%	0.4%	1.3%	2.1%	23.4%
NT-27	0.0%	0.0%	48.7%	0.5%	0.7%	13.1%	3.2%	33.8%	2.0%	59.7%
NT-28	88.2%	0.3%	3.8%	0.9%	0.1%	0.0%	0.0%	6.9%	6.7%	76.5%
NT-29	0.0%	1.7%	95.8%	1.0%	0.7%	0.0%	0.0%	0.7%	1.0%	62.8%
NT-30	1.6%	94.5%	0.6%	1.2%	1.2%	0.0%	0.4%	0.4%	2.1%	49.4%
NT-01	0.0%	0.0%	0.0%	99.2%	0.0%	0.0%	0.0%	0.8%	2.6%	41.1%
NT-02	0.0%	0.3%	0.3%	99.2%	0.0%	0.0%	0.0%	0.3%	5.3%	15.4%
NT-03	88.2%	0.3%	3.6%	1.0%	0.1%	0.0%	0.0%	6.9%	7.2%	71.4%
NT-04	0.0%	0.0%	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.5%	2.4%
NT-05	0.0%	0.0%	0.0%	96.6%	0.0%	0.0%	0.0%	3.4%	5.0%	1.2%
NT-06	0.0%	0.4%	0.4%	98.8%	0.0%	0.0%	0.0%	0.4%	1.2%	43.7%
NT-07	0.2%	0.0%	95.3%	0.9%	0.0%	1.6%	0.1%	1.9%	2.8%	60.6%
NT-08	1.0%	0.4%	95.3%	2.3%	0.4%	0.2%	0.3%	0.2%	9.4%	63.0%
NT-09	0.6%	0.0%	87.4%	1.9%	0.0%	0.0%	0.0%	10.1%	1.0%	33.8%
NT-10	78.3%	17.9%	3.0%	0.5%	0.0%	0.0%	0.0%	0.3%	1.9%	42.0%
NT-11	0.3%	0.0%	99.0%	0.3%	0.0%	0.3%	0.0%	0.0%	0.9%	70.3%
NT-12	0.0%	8.8%	76.5%	2.9%	5.9%	0.0%	0.0%	5.9%	2.0%	3.6%
NT-13	0.5%	2.0%	1.0%	96.6%	0.0%	0.0%	0.0%	0.0%	1.7%	50.7%
NT-14	0.0%	0.0%	99.1%	0.0%	0.0%	0.6%	0.0%	0.4%	7.7%	14.8%
NT-15	2.9%	0.5%	0.4%	95.5%	0.4%	0.0%	0.0%	0.2%	4.4%	45.2%
NT-16	0.4%	0.4%	17.9%	5.6%	64.1%	0.4%	6.8%	4.4%	1.4%	38.1%
NT-17	0.1%	0.0%	98.2%	0.5%	0.1%	0.1%	0.1%	0.9%	9.6%	85.4%
NT-18	0.1%	0.0%	95.7%	1.6%	0.0%	0.1%	0.2%	2.3%	4.7%	56.2%
NT-19	0.0%	0.0%	98.9%	0.0%	0.4%	0.0%	0.0%	0.7%	1.3%	72.6%
NT-20	2.0%	22.7%	3.0%	4.8%	63.9%	0.6%	2.3%	0.6%	6.8%	59.0%
NT-21	0.0%	0.0%	14.3%	42.9%	0.0%	0.0%	42.9%	0.0%	2.2%	0.7%
NT-22	1.4%	0.0%	11.0%	86.3%	0.0%	0.0%	0.0%	1.4%	1.0%	15.2%
NT-23	0.1%	0.0%	58.3%	0.8%	0.4%	5.0%	1.7%	33.7%	2.8%	62.7%
NT-24	0.0%	0.0%	100.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.6%	70.2%
NT-25	2.2%	0.0%	76.1%	4.3%	0.0%	2.2%	0.0%	15.2%	0.4%	23.5%
NT-26	0.0%	0.0%	2.3%	94.2%	3.5%	0.0%	0.0%	0.0%	0.8%	24.0%
NT-27	96.6%	0.2%	1.5%	1.1%	0.3%	0.2%	0.0%	0.2%	4.3%	32.2%
NT-28	1.2%	3.7%	1.5%	5.8%	85.7%	0.9%	0.9%	0.3%	7.6%	64.9%
NT-29	3.0%	82.0%	1.5%	13.5%	0.0%	0.0%	0.0%	0.0%	0.6%	45.4%
NT-30	0.0%	0.0%	1.0%	60.2%	19.4%	1.9%	4.9%	12.6%	2.1%	10.4%
Gold	15.0%	4.8%	38.5%	21.7%	14.6%	1.7%	0.8%	2.9%		

Table 7: Analysis of label alignment for nonterminals in the compound PCFG (top) and the neural PCFG (bottom). Label alignment is the proportion of correctly-predicted constituents that correspond to a particular gold label. We also show the predicted constituent frequency and accuracy (i.e. precision) on the right. Bottom line shows the frequency in the gold trees.



	(NT-13 (T-12 $w_1$ ) (NT-25 (T-39 $w_2$ ) (T-58 $w_3$ )))	
would be irresponsible could be delayed can be held can be proven could be used		has been growing 've been neglected had been made had been canceled have been wary
	(NT-04 (T-13 $w_1$ ) (NT-12 (T-60 $w_2$ ) (NT-18 (T-60 $w_3$ ) (T-21 $w_4$ )))	
of federally subsidized loans of criminal racketeering charges for individual retirement accounts without prior congressional approval between the two concerns		in fairly thin trading in quiet expiration trading in big technology stocks from small price discrepancies by futures-related program buying
	(NT-04 (T-13 $w_1$ ) (NT-12 (T-05 $w_2$ ) (NT-01 (T-18 $w_3$ ) (T-25 $w_4$ )))	
by the supreme court of the bankruptcy code to the bankruptcy court in a foreign court for the supreme court		in a stock-index arbitrage as a hedging tool of the bond market leaving the stock market after the new york
	(NT-12 (NT-20 (NT-20 (T-05 $w_1$ ) (T-40 $w_2$ )) (T-40 $w_3$ ) (T-22 $w_4$ ))	
a syrian troop pullout a conventional soviet attack the house-passed capital-gains provision the official creditors committee a syrian troop withdrawal		the frankfurt stock exchange the late sell programs a great buying opportunity the most active stocks a major brokerage firm
	(NT-21 (NT-22 (NT-20 (T-05 $w_1$ ) (T-40 $w_2$ )) (T-22 $w_3$ )) (NT-13 (T-30 $w_4$ ) (T-58 $w_5$ )))	
the frankfurt market was mixed the u.s. unit edged lower a news release was prepared the stock market closed wednesday the stock market remains fragile		the gramm-rudman targets are met a private meeting is scheduled the key assumption is valid the budget scorekeeping is completed the tax bill is enacted
	(NT-03 (T-07 $w_1$ ) (NT-19 (NT-20 (NT-20 (T-05 $w_2$ ) (T-40 $w_3$ )) (T-40 $w_4$ )) (T-22 $w_5$ )))	
have a high default risk have a lower default risk has a strong practical aspect have a good strong credit have one big marketing edge		rejected a reagan administration plan approved a short-term spending bill has an emergency relief program writes the hud spending bill adopted the underlying transportation measure
	(NT-13 (T-12 $w_1$ ) (NT-25 (T-39 $w_2$ ) (NT-23 (T-58 $w_3$ ) (NT-04 (T-13 $w_4$ ) (T-43 $w_5$ ))))	
has been operating in paris has been taken in colombia has been vacant since july have been dismal for years has been improving since then		will be used for expansion might be room for flexibility may be built in britain will be supported by advertising could be used as weapons
	(NT-04 (T-13 $w_1$ ) (NT-12 (NT-06 (NT-20 (T-05 $w_2$ ) (T-40 $w_3$ )) (T-22 $w_4$ )) (NT-04 (T-13 $w_5$ ) (NT-12 (T-18 $w_6$ ) (T-53 $w_7$ ))))	
for a health center in south carolina by a federal jury in new york of the appeals court in new york of the further thaw in u.s.-soviet relations of the service corps of retired executives		with an opposite trade in stock-index futures from the recent volatility in financial markets of another steep plunge in stock prices over the past decade as pension funds by a modest recovery in share prices
	(NT-10 (T-55 $w_1$ ) (NT-05 (T-02 $w_2$ ) (NT-19 (NT-06 (T-05 $w_3$ ) (T-41 $w_4$ )) (NT-04 (T-13 $w_5$ ) (NT-12 (T-60 $w_6$ ) (T-21 $w_7$ ))))	
to integrate the products into their operations to offset the problems at radio shack to purchase one share of common stock to tighten their hold on their business to use the microprocessor in future products		to defend the company in such proceedings to dismiss an indictment against her claiming to death some N of his troops to drop their inquiry into his activities to block the maneuver on procedural grounds
	(NT-13 (T-12 $w_1$ ) (NT-25 (T-39 $w_2$ ) (NT-23 (T-58 $w_3$ ) (NT-04 (T-13 $w_4$ ) (NT-12 (NT-20 (T-05 $w_5$ ) (T-40 $w_6$ )) (T-22 $w_7$ ))))	
has been mentioned as a takeover candidate has been stuck in a trading range had left announced to the trading mob only become active during the closing minutes will get settled in the short term		would be run by the joint chiefs would be made into a separate bill would be included in the final bill would be costly given the financial arrangement would be restricted by a new bill
	(NT-10 (T-55 $w_1$ ) (NT-05 (T-02 $w_1$ ) (NT-19 (NT-06 (T-05 $w_2$ ) (T-41 $w_3$ )) (NT-04 (T-13 $w_4$ ) (NT-12 (T-60 $w_5$ ) (NT-18 (T-18 $w_6$ ) (T-53 $w_7$ ))))	
to supply that country with other defense systems to transfer its skill at designing military equipment to improve the availability of quality legal service to unveil a family of high-end personal computers to arrange an acceleration of planned tariff cuts		to enjoy a loyalty among junk bond investors to transfer their business to other clearing firms to soften the blow of declining stock prices to keep a lid on short-term interest rates to urge the fed toward lower interest rates
	(NT-21 (NT-22 (T-60 $w_1$ ) (NT-18 (T-60 $w_2$ ) (T-21 $w_3$ ))) (NT-13 (T-07 $w_4$ ) (NT-02 (NT-27 (T-47 $w_5$ ) (T-50 $w_6$ )) (NT-10 (T-55 $w_7$ ) (NT-05 (T-47 $w_8$ ) (T-50 $w_9$ ))))	
unconsolidated pretax profit increased N % to N billion its total revenue rose N % to N billion total operating revenue grew N % to N billion its group sales rose N % to N billion total operating expenses increased N % to N billion		amex short interest climbed N % to N shares its pretax profit rose N % to N million its pretax profit rose N % to N billion fiscal first-half sales slipped N % to N million total operating expenses increased N % to N billion

Table 8: For each subtree (shown at the top of each set of examples), we perform PCA on the variational posterior mean vectors that are associated with that particular subtree and take the top principal component. We then list the top 5 constituents that had the lowest (left) and highest (right) principal component values.