

Structural Optimization Ambiguity and Simplicity Bias in Unsupervised Neural Grammar Induction

Jinwook Park and Kangil Kim*

AI Graduate School

Gwangju Institute of Science and Technology

jinwookpark@gm.gist.ac.kr, kangil.kim.01@gmail.com

Abstract

Neural parameterization has significantly advanced unsupervised grammar induction. However, training these models with a traditional likelihood loss for all possible parses exacerbates two issues: 1) *structural optimization ambiguity* that arbitrarily selects one among structurally ambiguous optimal grammars despite the specific preference of gold parses, and 2) *structural simplicity bias* that leads a model to underutilize rules to compose parse trees. These challenges subject unsupervised neural grammar induction (UNGI) to inevitable prediction errors, high variance, and the necessity for extensive grammars to achieve accurate predictions. This paper tackles these issues, offering a comprehensive analysis of their origins. As a solution, we introduce *sentence-wise parse-focusing* to reduce the parse pool per sentence for loss evaluation, using the structural bias from pre-trained parsers on the same dataset. In unsupervised parsing benchmark tests, our method significantly improves performance while effectively reducing variance and bias toward overly simplistic parses. Our research promotes learning more compact, accurate, and consistent explicit grammars, facilitating better interpretability.

1 Introduction

Grammar induction (GI) has been intensively studied due to its critical role in understanding, controlling, and utilizing the structural information of internal symbols called constituents in linguistic analysis. Unsupervised approaches to GI are continuously studied for their benefits of avoiding structured data as parses, which requires enormous human efforts and expertise to generate supervised data. In the literature of this field, neural parameterization for representing grammars (Kim et al., 2019a) has shown significant improvement over traditional probabilistic models (Klein and Manning,

2002, 2001; Clark, 2001). Furthermore, advanced models that incorporate conventional approaches, such as lexicalization (Zhu et al., 2020; Yang et al., 2021a) and symbol expansion (Yang et al., 2021b, 2022) have been proposed.

However, these Unsupervised Neural Grammar Inductions (UNGIs) lack analyses of the traditional challenges, particularly regarding the training loss. Some studies, such as Kim et al. (2019a), have highlighted the difficulty of optimization presented by the complex loss landscape in UGIs (Lari and Young, 1990; Carroll and Charniak, 1992; Klein and Manning, 2001). Subsequent efforts (Yang et al., 2021b, 2022) have alleviated these issues by only enhancing expressiveness through over-parameterization rather than the cause analysis.

In this paper, we aim to reveal two unhandled problems in UNGIs: 1) *structural optimization ambiguity*, which increases the number of optima of the same quality that prefer different parse structures, leading to prediction errors by random selection for branching and high variance, and 2) *structural simplicity bias*, which induces the grammar to use only a few rules in parse trees and resulting in simplistic parse trees even with a large capacity. Both problems are caused by training a neural grammar with sentence probability.

To address these issues, we propose a simple but effective method to inject focusing-bias that accounts for only a few selected parses for training loss in the inside algorithm. This approach reduces the ambiguity by selecting more distinct parses across sentences and restricts the simplicity bias by excluding of unnecessary simplistic parse trees by sentences. We also introduce a method of generating structural bias by parse selection from single or multiple parsers pre-trained on the same training data.

Our method outperforms state-of-the-art implicit and explicit grammar induction models in English parsing tasks on the Penn Treebank (PTB) (Marcus

*Corresponding author.

et al., 1993), as well as in ten other languages, including the Chinese Penn Treebank (CTB) (Xue et al., 2005) and the SPMRL datasets (Seddah et al., 2014). In-depth analysis shows a reduction in high variance and overly simplistic parses. We provide the investigation results of various structural biases, recommending heterogeneous multi-parsers.

Our contributions are summarized as follows:

- We raise and clarify *structural optimization ambiguity* and *structural simplicity bias*, which lead to high variance and overly simplistic parse problems in UNGIs.
- We propose *sentence-wise parse-focusing* that employs biases from pre-trained parsers using the same training data. This approach enables the stable learning of more compact and accurate explicit grammars.
- Through in-depth empirical analysis, we demonstrate the effectiveness of our approach in reducing the identified causes, investigating various focusing-biases, and achieving significant performance improvements compared to state-of-the-art UNGIs.

2 Background

2.1 Notations of Probabilistic Context-Free Grammar

In the following analysis, we use the notations below for a PCFG $G = (S, N, P, \Sigma, R, \Pi)$: the root symbol S , finite sets of nonterminals N , preterminals P , terminal words Σ , production rules R , and production rule probabilities Π . R consists of three types of rules:

$$\begin{aligned} S &\rightarrow A && \text{where } A \in N \\ A &\rightarrow \alpha \beta && \text{where } \alpha, \beta \in (N \cup P) \\ B &\rightarrow \omega && \text{where } B \in P \text{ and } \omega \in \Sigma \end{aligned}$$

Note that the left-hand side (parent) of the rules is distinguished between nonterminal and preterminal categories.

2.2 Negative Log Likelihood Loss

The loss of grammar induction is well-known to increase the probability mass of only observed parse trees (Chi and Geman, 1998) and to maximize rule probability for representing correct parses of sentences (Corazza and Satta, 2006). The loss of UGI is the same as SGI in that they both maximize sentence probability for a given sentence; however, it

is different in that UGI uses all possible parse trees to calculate sentence probability.

The likelihood loss for a given sentence is:

$$L(s) = -\log \sum_{\tau \in T(s)} \exp \sum_{r \in R(\tau)} f(r|\tau) \log p(r|G)$$

where the exponential term is the probability of a parse tree τ , and $f(r|\tau)$ denotes the frequency of rules used in the tree. In the case of SGI, $T(s)$ only contains the gold parse tree for s . However, for UGI, $T(s)$ contains all parse trees that can generate s .

In the equation 2.2, the loss reaches its minimum when the exponential term within the summation, which is always negative or equal to zero, approaches zero. This can be rewritten as $\sum f(A \rightarrow \alpha) \log P(A|\alpha) \simeq 0$.

3 Limits on Unsupervised Learning of Neural Grammar Induction

As explained earlier, UNGI Loss has the advantage of inducing grammar without the information of gold parse trees. However, UNGIs still inherit the limitations of probabilistic UGIs. In this section, we reveal two issues arising from these limitations in UNGIs: *structural optimization ambiguity* and *structural simplicity bias*, and discuss derived negative effects and their causes. For precise analysis, we focus on inducing explicit neural grammars based on PCFGs, mainly Neural PCFGs (N-PCFGs).

3.1 Structural Optimization Ambiguity (SOA)

Structural optimization ambiguity causes a huge variance of performance by arbitrary branching bias that depends on random seeds. In this section, we clarify the definition of SOA, and explain the arbitrary convergence caused by SOA. After that, we mathematically prove the existence of SOA under the single pre-terminal condition and extend this to the multiple pre-terminal condition using empirical evidences.

Definition of Structure Optimization Ambiguity by UGI Loss SOA is the ambiguity that disrupts the distinction between two optima that predict different parse trees for the same sentence. For clear analysis, we define the ambiguity as follows.

Definition 1. Given $G_1 = (S, N, P, \Sigma, R, \Pi_1)$ and $G_2 = (S, N, P, \Sigma, R, \Pi_2)$ that assign different rule probabilities from the same G , the grammar G is




Figure 1: The optimization of two different grammars, G_1 and G_2 , in the optimization landscape for a given sentence. Each grammar has different preferences for parse trees based on parse tree probability; however, they have the same sentence probability. The two initial points converge to the different optima close to their own.

structural-optimization-ambiguous if and only if $p(s|G_1) = p(s|G_2)$ for all sentences $s \in S$ in the training data, and there exists a sentence s whose best parses respectively derived through G_1 and G_2 are different.

As mentioned in Section 2.2, the sentence probability is the loss function of GI. This means that if two grammars have the same sentence probability, they are not distinguished in the optimization landscape. In other words, they are ambiguous in the perspective of optimization.

Arbitrary Convergence by Structural Optimization Ambiguity The discrepancy between structure and loss function leads to that each grammar randomly initialized favors different optimal points. In Figure 1, highlighted two optimal points are equally probable yet structurally divergent. At the blue optimum, the parse tree t_1 has a higher probability than the parse tree t_2 , whereas, at the red optimum, the situation is reversed, with t_2 having a higher probability than t_1 . Consequently, it becomes unclear which optimum will be chosen based solely on sentence probabilities from an optimization perspective. Moreover, it is intuitively understood that the reachable optimum may vary with the randomness, such as initial states.

Proof in Single Pre-terminal Condition We prove that SOA always exists under the single pre-terminal condition by using equivalence of sen-

tence probability between two grammars that have different probability distributions.

In the first step, we consider constraints on the grammars: 1) both grammars share the same configuration (S, N, P, Σ, R) , and 2) the rule probabilities of both grammars are identical except for a single type of flipped rule pair ($N_i \rightarrow N_j T$ and $N_i \rightarrow T N_j$). Under these constraints, suppose that the two grammars yield the same sentence probability for a given sentence; we can derive the following equation in general condition from the formula of the inside algorithm:

$$p_G(N_i \rightarrow N_j T) + p_G(N_i \rightarrow T N_j)\alpha = p_{G'}(N_i \rightarrow N_j T) + p_{G'}(N_i \rightarrow T N_j)\alpha, \quad (1)$$

Where,

$$\alpha = \frac{p_{G_1}(T \rightarrow w_1)I_{G_1,j}(2, r)}{I_{G_1,j}(1, r-1)p_{G_1}(T \rightarrow w_r)}$$


Then, we also consider that each grammar includes only one pre-terminal¹. Under the aforementioned constraints, $\alpha = 1$ is always satisfied; therefore, the above equation 1 can be simplified as follows:

$$p_G(N_i \rightarrow N_j T) + p_G(N_i \rightarrow T N_j) = p_{G'}(N_i \rightarrow N_j T) + p_{G'}(N_i \rightarrow T N_j) \quad (2)$$

Under the identical probability distribution except for the flipped rule pair, the sum of the probabilities for the flipped rule pair is always equal in a proper grammar (Nederhof and Satta, 2006). Thus, equation 2 is always satisfied, and the two grammars always yield the same probability for the same sentence. Furthermore, due to the strict inequality of the flipped rule pair, the two grammars can derive different parse trees. This signifies that the two grammars are structural-optimization-ambiguous. Appendix A.1 provides a more detailed proof of this.

Empirical Evidence in General Condition: Low Correlation of Loss to S-F1 Because proving the existence is significantly complex under general condition using multiple pre-terminals, we prove it indirectly by presenting empirical evidences. The

¹The single pre-terminal condition is also established when 1) rule probabilities are uniformly initialized or 2) all tokens in sentences cannot be distinguished (e.g. all tokens are $\langle unk \rangle$).



(a) Correlation of S-F1 score and likelihood (b) Amplification for the variance of S-F1 by symbol size

Figure 2: (a) Low correlation between S-F1 score and UGI likelihood for PTB test set. We evaluate trained FGG-TNPCFGs with 4500 nonterminal symbols and 9000 preterminal symbols using 32 random seeds over 10 epochs. (b) Amplification of variance by symbol size increasing. Circles are the same as in (a) and others are trained with the same settings as in (a) except for the number of symbols.

first evidence is a low correlation between sentence-level F1 (S-F1) scores² and negative log likelihood loss, which implies that the optimization of grammars is inconsistent, and can lead to generating structures different from those desired by gold parses.

In Figure 2a, the negative log likelihood and S-F1 scores are shown to be almost uncorrelated, as indicated by a low P-value and Pearson coefficient. This low correlation means that the distinction between the grammars that induce different structures is hard to determine using negative log likelihood, which causes the large variance of S-F1 scores due to randomness. Through this analysis, we demonstrate the existence of SOA under general conditions.

Empirical Evidence in General Condition: Variance Amplification by Symbol Size Another piece of evidence is the variance amplification of S-F1 scores with the increase in the number of symbols, as shown in Figure 2b. It appears that more nonterminals create more local optima of similar quality that derive different parses, leading to a greater variety of parses compared to a single gold parse. This demonstrates the variance amplification of SOA as the number of symbols increases under general conditions.

²The S-F1 score is one of the common metrics to evaluate parsing quality. It evaluates the similarity between predicted parses and gold parses. In this paper, we only use unlabeled S-F1 scores that do not consider the symbol labels the same as N-PCFG.

3.2 Structural Simplicity Bias (SSB)

The second problem, *structural simplicity bias*, causes a decrease in expressive power for parse trees by leading the grammar to have low structural diversity. In this section, we discuss what SSB is, the problems it causes, and why it occurs. Lastly, we present empirical evidence for how SSB is revealed during training.


What is SSB? SSB is the bias that induces a grammar to utilize as few rules as possible to compose parse trees. This bias is a natural and intended effect of the objective to maximize sentence probability. However, it leads to learning undesirable structures that are different from gold parses in UNGI.

Specific Conditions of UGI UGI is affected by different biases, unlike SGI, due to two properties: 1) the loss maximizes the probabilities of all possible trees for a given sentence. 2) all rules can be used in any position that has the same type³. Therefore, UGI prefers the structure that maximizes tree probability rather than optimizing the given parses as SGI does.

Rule Simplification for Expressing a Parse If all rules in the tree are the same, it is intuitively easy to maximize tree probability because the model only needs to optimize one rule. Therefore, maximization using tree probability prefers to train utilizing a small number of rules. This preference leads to a grammar that includes simple structures in parses. This means that the grammar has low expressive power for parse trees. Specifically, only rules of same type can share their position in the tree due to symbol types; therefore, parses prefer structures that utilize structural repetition. For these reasons, it prefers left- or right-binarized structures with many repetitions.

Space Inefficiency and Performance Degradation by the Potential Expressive Limit The grammar affected by SSB utilizes only a small part of its potential expressive power for parse trees, which causes two problems: 1) space inefficiency, as expressive power grows slowly compared to an increase in grammar size, and 2) low expressive power compared to gold parses, leading to decreased performance. The low rule diversity of

³The type of rules means the four types, $N \rightarrow NN|NP|PN|PP$, distinguished by the type of symbols: nonterminal, preterminal.



(a) Comparison between the gold parse trees and FGG-TNPCFGs. (b) Comparison between FGG-TNPCFGs and parse-focused N-PCFGs.

Figure 3: The average number of unique rules in each parse for each sentence length. We evaluate the average for models with 5 / 15 / 30 nonterminals trained with four different random seeds. These evaluations use the WSJ train set. (a) The gold uses non-binarized gold parses. (b) Parse-focused N-PCFGs use models trained by Structformer, NBL-PCFG, and FGG-TNPCFG.

grammar compared to the necessity of gold parses cannot represent the necessary tree structures for a given sentence. Induced parses that are not sufficiently represented cause low performance.

Empirical Evidence: Average Diversity in Each Parse Tree We demonstrate that existing models induce simplified grammars by analyzing the predicted parses. We provide evidence that SSB leads a model to learn simplistic diversity compared to gold parses, as shown in Figure 3a, which simplifies the induced tree structures and results in similar structures among trees of the same length. In other words, this can damage diversity. This decrease in diversity appears to lead to a decrease in performance.

In Figure 3b, the grammars of the same size with parse-focusing have greater rule diversity than FGG-TNPCFGs. This shows that 1) parse-focusing relaxes the SSB, and 2) the decrease in rule diversity is not caused by the upper bound for grammar capacity. We verify the loss of rule diversity in other languages in Appendix A.2 and come to the same conclusion. Additionally, we provide the qualitative analysis of overly simplistic parses in Appendix A.3.

4 Method

4.1 Preliminaries

Motivation SOA and SSB have not yet been reported as critical problems, even though they can also exist in conventional probabilistic GIs (Klein and Manning, 2001). The reason is that conven-

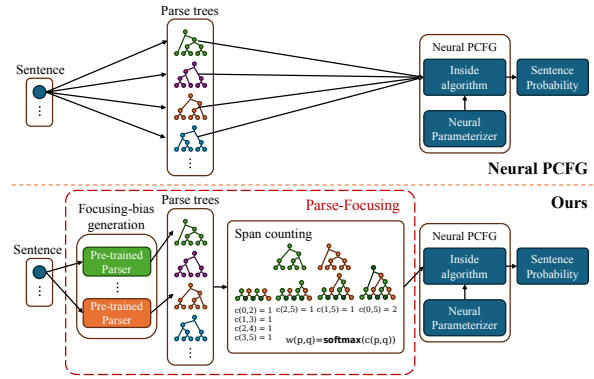


Figure 4: Overview of N-PCFG (above) and Ours (below). The red box shows the parse-focusing method.

tional methods can avoid equal frequency assignment to all parses, which is the basic condition that causes these problems. For example, the EM algorithm in UGIs reassigns different rule frequencies at every iteration (Petrov et al., 2006), and SGIs use only gold parse trees while implicitly assigning zero probability to all incorrect parses. Motivated by this, we propose *sentence-wise parse-focusing* to learn from a few parses for each sentence and *structural bias generation* to select the parses from unsupervised multi-parsers pre-trained on the same training data.

Overview of Parse-Focusing Figure 4 presents an example application of sentence-wise parse-focusing to N-PCFGs. Compared to N-PCFG, the distinguishing feature is selecting a few parses from unsupervised parsers pre-trained on the same training data and deriving the training loss from these parses. Details of the algorithm are provided in Appendix A.4.

4.2 Sentence-wise Parse-Focusing

Learning from Focused Parses To learn grammar from a few selected parses, we simply modify the inside algorithm to count constituent frequency only for observed parses. However, this strict bias injection completely excludes unselected parses, resulting in losing potential structures. To alleviate this limitation, we use soft weights derived from softmax over all constituent counts, which maintains the focus on the selected parses and also leaves small weights on potential structures. The process of applying the parse-focusing is detailed in Algorithm 1.

Given a finite set of sentences $s \in \mathcal{S}$, we represent a finite set of selected parse trees for each sentence s as $\tau_s \in \mathcal{T}_s$. We denote the frequency

for span (p, q) of word p to q for parse tree τ_s as $f((p, q)|\tau_s)$. Then, the soft weights are

$$w(p, q|s) = \frac{\exp \sum_{\tau_s \in \mathcal{T}_s} f((p, q)|\tau_s)}{\sum_{p, q} \exp \sum_{\tau_s \in \mathcal{T}_s} f((p, q)|\tau_s)}$$

Why Does It Alleviate High Variance and Overly Simplistic Parse Issues? As mentioned in Section 4.1, the base condition for causing SOA and SSB is that all parse trees are considered equally.

This equal consideration establishes conflicting pairs of flipped rules with equal frequency, which causes SOA. Sentence-wise parse-focusing avoids this equal frequency by using soft weights for spans. Additionally, sentence-wise parse-focusing has no specific dynamics to consider them equally. These reasons significantly decrease the possibility of establishing Equation 1 under general condition, which avoids the problems caused by SOA in training.

This consideration also makes a model maximize advantageous simplistic parse trees that are not related to gold parse trees, which causes SSB. These simplistic parse tree structures are shared across sentences. Sentence-wise parse-focusing avoids the maximization of the simplistic parse trees by assigning selected parse trees for given sentences. The selected parse trees may not involve the simplistic parse trees and are also different across sentences.

4.3 Focusing-Bias Generation

Bias Generation from Pre-Trained Unsupervised Parsers The performance of grammar relies on the quality of selected parses because focused parses control the loss for given sentences. Therefore, we can expect to further improve performance when we use high-quality parses. However, in unsupervised learning, finding parses close to gold parse trees without supervision-related information is a challenging issue. As an alternative solution, we use additional parsers pre-trained on the same data with unsupervised learning. Note that our method allows any type of parser as long as it provides parse span information.

Data-Specific Bias Induction via Heterogeneous Multi-Parsers A potential approach to generate good parses is to use parsers induced by different models, which we call heterogeneous parsers. According to Ishii and Miyao (2023), the branching bias consists of model-specific inherent bias

Homogeneous		Heterogeneous	
Parser Pair	IoU	Parser Pair	IoU
(SF, SF)	56.3±7.0	(SF, NBL)	32.8±9.8
(FGG, FGG)	47.0±7.3	(SF, FGG)	38.9±5.6
(NBL, NBL)	38.5±19.7	(NBL, FGG)	37.9±12.3

Table 1: IoU scores of homogeneous and heterogeneous parser combinations. SF means Structformer, FGG means FGG-TNPCFGs and NBL means NBL-PCFGs.

and data-specific inductive bias. They also show that different parsers have different model-specific inherent biases. Therefore, if we use heterogeneous parsers when counting constituents in parse-focusing, the opposing inherent biases between parsers are offset, while the same structures by data-specific bias are enhanced.

In Table 1, we compare the structural differences between parsers within pairs of heterogeneous parsers and within pairs of homogeneous parsers. Homogeneous parsers are trained using the same model but with different random seeds. To directly quantify structural heterogeneity, we use IoU scores on parse spans defined as:

$$\frac{\bigcap_{G_p \in \mathcal{G}_p} T_{G_p}}{\bigcup_{G_p \in \mathcal{G}_p} T_{G_p}}$$

The IoU scores of heterogeneous parsers are consistently lower than those of homogeneous parsers, which implies that these parsers have more different branching biases. Therefore, when using heterogeneous parsers, the different model-specific inherent biases are offset, making the shared part more likely to indicate data-specific bias.

5 Experiments

5.1 Settings

Datasets and Hyperparameters We use the Penn Treebank (PTB) (Marcus et al., 1994), an English constituency parsing task dataset, to evaluate parsing performance. Additionally, we use the CTB (Xue et al., 2005) and SPMRL dataset (Seddah et al., 2014) for multilingual evaluation.

We use the same hyperparameters as Yang et al. (2022); for the analysis of the number of symbols, we follow the ratio of 1:2 between nonterminals and preterminals. The details for the dataset and training are in Appendix A.6

Inference We primarily evaluate grammars with MBR decoding (Yang et al., 2021b). However, we

use CYK decoding when it is necessary to distinguish each rule in the grammar.

Base Model We use the model from Yang et al. (2022), which we refer to as Factor Graph Grammar-based TN-PCFGs (FGG-TNPCFGs), as our base model. Our model is implemented based on the open-sourced code from Yang et al. (2022)⁴. There are two reasons for choosing this model: 1) The strict assumption of PCFGs is preserved. In the case of Compound PCFGs (C-PCFGs) and Lexicalized PCFGs (Zhu et al., 2020; Yang et al., 2021a), because they are dependent on words or sentences, which makes the induced grammar complex to understand. We avoid this to verify whether better performance can be achieved without relaxed assumptions. 2) FGG-TNPCFGs rapidly train large-scale grammars through efficient use of memory and computation. This is favorable for obtaining diverse results in different conditions. Additionally, we note that while FGG-TNPCFGs provide efficient computation, the fundamental mathematical principle for grammar remains the same.

Selected Parsers for Focusing-Bias We adopt three parsers that use different algorithms; 1) Structformer (Shen et al., 2021) is a transformer-based and implicit grammar model. 2) NBL-PCFG (Yang et al., 2021a) is a lexicalized explicit grammar model. 3) FGG-TNPCFG (Yang et al., 2022) is the largest and state-of-the-art grammar model. Note that all of these parsers use the same hyperparameter settings as in their original papers, and they are trained only with PTB. The primary model trained by parse-focusing consists of these three parsers, which shows the best performance. Additionally, we use DIORA (Drozdov et al., 2019), which is a self-supervised model utilizing a recursive autoencoder architecture. However, it is not used as the main parser because it is trained with the additional dataset MultiNLI (Williams et al., 2018b) besides PTB.

5.2 Performance

Performance in English We evaluated the performance of state-of-the-art UNGI methods on English PTB data, as shown in Table 2. Our approach significantly outperforms the state-of-the-art model (FGG-TNPCFGs) with only 30 nonterminals and 60 preterminals. When using 4500 nonterminals, as with FGG-TNPCFGs, the performance further

grammar type	model	S-F1	
		mean	max
implicit grammar	Structformer	54.0	-
	Structformer [†]	52.3 \pm 2.3	54.2
	DIORA	55.7	56.8
	DIORA [†]	43.6 \pm 0.9	44.8
explicit grammar	N-PCFG	50.9 \pm 2.3	52.3
	N-PCFG w/ MBR	52.3 \pm 2.3	55.8
	C-PCFG	55.4 \pm 2.2	59.0
	C-PCFG w/ MBR	56.3 \pm 2.1	60.1
	NL-PCFG	55.3	-
	NBL-PCFG	60.4	-
	NBL-PCFG [†]	53.3 \pm 11.5	62.3
	TN-PCFG	57.7 \pm 4.2	55.6
	FGG-TNPCFG	64.1	-
	FGG-TNPCFG [†]	57.4 \pm 6.0	66.8
Ours (NT=30)		67.4 \pm 0.9	68.4
Ours (NT=4500)		69.6 \pm 0.6	70.3

Table 2: Performance of UNGIs on English trained on Penn TreeBank ([†]: reproduced result). All statistics are from 32 runs with different random seeds. Our model is trained with randomly and separately generated focusing-biases from Structformer[†], NBL-PCFG[†], and FGG-TNPCFG[†].

increases significantly. Note the difference between reproduced models and the variance among models. Instead of running only four runs, we performed 32 runs, resulting in lower performance and higher variance, especially in FGG-TNPCFG and NBL-PCFG. The significant performance difference between DIORA and DIORA[†] is due to differing evaluation criteria. While DIORA utilizes a test set of binarized trees by Stanford CoreNLP, DIORA[†] employs non-binarized trees. This approach is intended to assess performance accurately within the evaluation environment of N-PCFG.

Performance in Multilingual Table 3 shows the performance of explicit grammar-based UNGIs in multilingual parsing on CTB and SPMRL data. The results indicate that our model significantly outperforms other models across most languages. While a few languages, such as Hungarian, Polish, and Swedish, show second or third-best performance, our model ranks the highest overall. Note that our method has consistently and significantly lower variance compared to all other methods, which indicate that it is a more reliable estimator.

⁴<https://github.com/sustcsonglin/TN-PCFG>

Model	Basque	Chinese	English	French	German	Hebrew	Hungarian	Korean	Polish	Swedish	Mean rank
N-PCFG*	35.1 \pm 2.0	26.3 \pm 2.5	52.3 \pm 2.3	45.0 \pm 2.0	42.3 \pm 1.6	45.7 \pm 2.2	43.5 \pm 1.2	28.4 \pm 6.5	43.2 \pm 0.8	17.0 \pm 9.9	4.3
C-PCFG*	36.0 \pm 1.2	38.7 \pm 6.6	56.3 \pm 2.1	45.0 \pm 1.1	43.5 \pm 1.2	45.2 \pm 0.5	44.9\pm1.5	30.5 \pm 4.2	43.8 \pm 1.3	33.0 \pm 15.4	3.2
TN-PCFG*	36.0 \pm 3.0	39.2 \pm 5.0	57.7 \pm 4.2	39.1 \pm 4.1	47.1 \pm 1.7	39.2 \pm 10.7	43.1 \pm 1.1	35.4 \pm 2.8	48.6\pm3.1	40.0\pm4.8	3.1
FGG-TNPCFG \dagger	38.4 \pm 7.3	31.0 \pm 8.4	59.6 \pm 7.7	43.9 \pm 3.1	48.0 \pm 1.4	46.2 \pm 4.1	42.2 \pm 0.7	31.5 \pm 4.0	41.6 \pm 4.3	40.0\pm0.6	3.0
Ours	45.9\pm0.3	46.1\pm0.9	69.7\pm0.9	50.5\pm0.5	49.1\pm0.3	49.5\pm0.2	43.7 \pm 0.2	42.1\pm0.3	47.9 \pm 0.3	33.4 \pm 1.2	1.4

Table 3: Performance (S-F1 score) in multilingual parsing on CTB and SPMRL datasets. * indicates reported by Yang et al. (2021b), \dagger indicates results reproduced using open-sourced code provided by Yang et al. (2022).

Bias	S-F1	NLL
LBranching	8.9 \pm 0.0	116.5 \pm 2.3
RBranching	39.7 \pm 0.0	116.1 \pm 1.6
Random	34.9 \pm 1.2	113.5 \pm 1.0
Structformer	58.9 \pm 0.2	112.0 \pm 0.9
NBL-PCFG	64.7 \pm 0.6	111.5 \pm 0.7
FGG-TNPCFG	65.3 \pm 0.4	113.4 \pm 1.3
DIORA	50.0 \pm 1.4	111.6 \pm 0.3
FGG-TNPCFG \dagger	57.4 \pm 6.0	111.9 \pm 1.1

Table 4: S-F1 and NLL for parse-focused grammar by each focusing-bias. LBranching and RBranching refer to left-branching and right-branching, respectively. Random and pre-trained biases are generated once and used for all runs. FGG-TNPCFG \dagger denotes reproduced performance to compare variance.

5.3 In-depth Analysis

Performance Consistency about Randomness

We investigate the impact of our method on the high variance problem, which is the main limitation of UNGIs. To observe the variance change caused solely by parse-focusing, we also evaluate the variance using the generated sets of parse trees as the focusing-bias: randomly, left-only branched, and right-only branched. As shown in Table 4, all of the variances of S-F1 scores are significantly reduced compared to those of FGG-TNPCFG, while their negative log likelihoods exhibit a similar bias. This means that sentence-wise parse-focusing reduces variance without depending on the selected parses. The results indicate that applying parse-focusing reduces the variance problem and indirectly supports the idea that the cause is due to SOA under uncorrelated S-F1 and NLL. Moreover, they show that focusing-bias from pre-trained parsers can lead to higher performance than algorithmically generated parses.

Performance and Variance by Symbol Size In Section 3.2, we show that parse-focusing relaxes the low diversity of unique rules in conventional methods. Here, we demonstrate how the relaxation

Nonterminal	Preterminal	S-F1	
		Ours	FGG
1	2	44.3 \pm 0.1	39.7 \pm 0.0
5	10	60.4 \pm 1.9	41.6 \pm 9.6
15	30	65.5 \pm 0.7	43.2 \pm 0.8
30	60	67.4 \pm 0.9	51.2 \pm 3.1
250	500	69.7 \pm 0.9	54.3 \pm 3.9
4500	9000	69.6 \pm 0.7	57.4 \pm 6.0

Table 5: The S-F1 score for each number of symbols for FGG-TNPCFGs and induced grammars by parse-focusing.




Figure 5: Frequency distribution sorted in descending order for contained rules in the parsed training dataset. The colored fill area represents the frequency difference between FGG-TNPCFGs and our method with 5 nonterminals. The small box zooms in on the top 10 rules for differences.

from parse-focusing affects the increase in performance according to changes in symbol size. In Table 5, although our method boosts performance in the larger model, it significantly enhances performance compared to FGG-TNPCFGs in the smaller grammar model. These results imply that our method can stably induce more compact grammars by reducing SSB while achieving more accurate models using the same model capacity.

Rule Utilization We show that parse-focusing reduces the extreme reliance on a small number of rules in FGG-TNPCFG. Figure 5 shows the sorted frequency of all rules in a grammar. After applying our method, the frequency of the most frequent




Figure 6: The difference between the mean of S-F1 scores for pre-trained parsers and the parse-focused model trained by those parsers. The blue bars indicate homogeneous parsers, and the orange bars indicate heterogeneous parsers. The raw S-F1 scores are shown in Figure 10.

three rules decreases from 46.1% to 38.7%, and the reduced frequency is redistributed to other less frequently used rules. These results imply the effective enhancement of rule utilization.

Practical Multi-Parser Combinations We show the significance of using heterogeneous multi-parsers that have different model-specific inherent biases. Figure 6 shows that the performance increase in heterogeneous parsers is generally more considerable compared to homogeneous ones. These parsers exhibit more significant variances in performance, which supports the idea that the inherent biases of models are offset. Consequently, the results implies that heterogeneity is a practically helpful strategy.

6 Related Works

TN-PCFG/FGGs Yang et al. (2021b) and Yang et al. (2022) verify that overparameterization leads to better grammars with many symbols. Conversely, we highlight the inefficiency in learning UGIs caused by simplicity and ambiguity issues. Solving this inefficiency enables learning even a small grammar with significantly high performance and low variance.

S-DIORA Drozdov et al. (2020) shows that structural ambiguity between subtrees can affect locally greedy training, and that this problem can be resolved by single tree encoding through beam search. This work supports the impact of the ambiguity issue, although our research addresses the issue of learning explicit grammars in UGIs differently.

Ensemble Distillation Shayegh et al. (2023) introduce a variant of the CYK algorithm that gains

refined parse trees by combining trees from the different pre-trained parsers. Our approach differs in searching for structure through training trees in a relaxed condition using soft weights from the count of spans. It utilizes the model’s search capability. In addition, it is different that we resolve the optimization problem for the loss function.

Tree-shape Uncertainty The tree-shape uncertainty referred to in Ishii and Miyao (2023) is the existence of similar grammars that assign different tree structures for a given sentence. In our work, the issue is analyzed from an optimization perspective, highlighting how undifferentiated optima can cause unstable learning.

What Do Implicit Grammar Models Learn? Williams et al. (2018a) investigate how consistently latent tree learning models learn in the presence of randomness. In our work, we focus on identifying the causes of inconsistency and its negative effects, which are primarily due to likelihood loss over all parses in UGI.

7 Conclusion

In this paper, we tackle the challenges of high variance and overly simplistic parse in unsupervised neural grammar induction (UNGI). We identify *structural optimization-ambiguity* and *structural simplicity bias* as primary causes, arising from UGI loss on all possible parses of all sentence parses. We introduce a parse-focusing approach that narrows down the parse set per sentence and incorporates biases from parses generated by pre-trained parsers on the same dataset. This method has been shown to effectively mitigate these issues, significantly surpassing state-of-the-art methods in both English and multilingual parsing benchmarks. Our contributions facilitate the consistent learning of more compact, explicit, accurate, and therefore interpretable grammar models in unsupervised learning environment.

Acknowledgements

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2022R1A2C2012054, Development of AI for Canonicalized Expression of Trained Hypotheses by Resolving Ambiguity in Various Relation Levels of Representation Learning).

Limitations

Our work is limited to offline learning due to additional computational costs for pre-training parsers and for the hyper-parameter search of better combinations. In the future, more light and adaptive structural bias generation is required.

References

- Glenn Carroll and Eugene Charniak. 1992. Two experiments on learning probabilistic dependency grammars from corpora. In *AAAI Workshop on Statistically-Based NLP Techniques*.
- Zhiyi Chi and Stuart Geman. 1998. Estimation of probabilistic context-free grammars. *Computational linguistics*, 24(2):299–305.
- Alexander Clark. 2001. [Unsupervised induction of stochastic context-free grammars using distributional clustering](#). In *Proceedings of the 2001 Workshop on Computational Natural Language Learning - Volume 7, ConLL '01, USA*. Association for Computational Linguistics.
- Anna Corazza and Giorgio Satta. 2006. Cross-entropy and estimation of probabilistic context-free grammars. In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 335–342.
- Andrew Drozdov, Subendhu Rongali, Yi-Pei Chen, Tim O’Gorman, Mohit Iyyer, and Andrew McCallum. 2020. [Unsupervised parsing with S-DIORA: Single tree encoding for deep inside-outside recursive autoencoders](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4832–4845, Online. Association for Computational Linguistics.
- Andrew Drozdov, Pat Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. [Unsupervised latent tree induction with deep inside-outside recursive autoencoders](#).
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. [Recurrent neural network grammars](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 199–209, San Diego, California. Association for Computational Linguistics.
- Taiga Ishii and Yusuke Miyao. 2023. [Tree-shape uncertainty for analyzing the inherent branching bias of unsupervised parsing models](#). In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 532–547, Singapore. Association for Computational Linguistics.
- Yoon Kim, Chris Dyer, and Alexander M Rush. 2019a. Compound probabilistic context-free grammars for grammar induction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2369–2385.
- Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019b. [Unsupervised recurrent neural network grammars](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1105–1117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dan Klein and Christopher D Manning. 2001. Natural language grammar induction using a constituent-context model. *Advances in neural information processing systems*, 14.
- Dan Klein and Christopher D Manning. 2002. A generative constituent-context model for improved grammar induction. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 128–135.
- Karim Lari and Steve J Young. 1990. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer speech & language*, 4(1):35–56.
- Mitch Marcus, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. 1994. The penn treebank: Annotating predicate argument structure. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Mark-Jan Nederhof and Giorgio Satta. 2006. [Estimation of consistent probabilistic context-free grammars](#). In *Proceedings of the Main Conference on Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics, HLT-NAACL '06*, page 343–350, USA. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Slav Petrov, Leon Barrett, Romain Thibaux, and Dan Klein. 2006. Learning accurate, compact, and interpretable tree annotation. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 433–440.

- Djamé Seddah, Sandra Kübler, and Reut Tsarfaty. 2014. [Introducing the SPMRL 2014 shared task on parsing morphologically-rich languages](#). In *Proceedings of the First Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 103–109, Dublin, Ireland. Dublin City University.
- Behzad Shayegh, Yanshuai Cao, Xiaodan Zhu, Jackie C. K. Cheung, and Lili Mou. 2023. [Ensemble distillation for unsupervised constituency parsing](#).
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. 2018. [Neural language modeling by jointly learning syntax and lexicon](#).
- Yikang Shen, Shawn Tan, Alessandro Sordani, and Aaron Courville. 2019. [Ordered neurons: Integrating tree structures into recurrent neural networks](#).
- Yikang Shen, Yi Tay, Che Zheng, Dara Bahri, Donald Metzler, and Aaron Courville. 2021. [Structformer: Joint unsupervised induction of dependency and constituency structure from masked language modeling](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7196–7209.
- Yau-Shian Wang, Hung-Yi Lee, and Yun-Nung Chen. 2019. [Tree transformer: Integrating tree structures into self-attention](#).
- Adina Williams, Andrew Drozdov*, and Samuel R. Bowman. 2018a. [Do latent tree learning models identify meaningful structure in sentences?](#) *Transactions of the Association for Computational Linguistics*, 6:253–267.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018b. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Nianwen Xue, Fei Xia, Fu-Dong Chiou, and Martha Palmer. 2005. [The penn chinese treebank: Phrase structure annotation of a large corpus](#). *Natural Language Engineering*, 11:207–238.
- Songlin Yang, Wei Liu, and Kewei Tu. 2022. [Dynamic programming in rank space: Scaling structured inference with low-rank HMMs and PCFGs](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4797–4809, Seattle, United States. Association for Computational Linguistics.
- Songlin Yang, Yanpeng Zhao, and Kewei Tu. 2021a. [Neural bi-lexicalized PCFG induction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2688–2699, Online. Association for Computational Linguistics.
- Songlin Yang, Yanpeng Zhao, and Kewei Tu. 2021b. [PCFGs can do better: Inducing probabilistic context-free grammars with many symbols](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1487–1498, Online. Association for Computational Linguistics.
- Hao Zhu, Yonatan Bisk, and Graham Neubig. 2020. [The return of lexical dependencies: Neural lexicalized pcfgs](#). *Transactions of the Association for Computational Linguistics*, 8:647–661.

A Appendices

A.1 Detail Proof of Structural Optimization Ambiguity with the Inside Algorithm

We incorporate the three stipulated conditions into the conventional inside algorithm as per Equation 3: 1) distinguishing between nonterminals and preterminals, 2) allowing all symbols to occupy identical positions in parse trees, and 3) encompassing all structurally possible binary trees for given sentences. Therefore, we can reformulate the inside algorithm as shown in Equation 4.

We examine a special case where a grammar contains only one preterminal symbol, as detailed in Equation 5, to investigate the potential for structural optimization ambiguity. We recursively derive the conditions leading to ambiguity.

Before we begin the proof, let us consider two grammars, G_1 and G_2 . The relationship between them is defined as follows:

$$\begin{aligned} p_{G_1}(N_i \rightarrow N_j T) &\neq p_{G_2}(N_i \rightarrow N_j T) \\ p_{G_1}(N_i \rightarrow T N_j) &\neq p_{G_2}(N_i \rightarrow T N_j) \\ \forall r \in R^c, p_{G_1}(r) &= p_{G_2}(r) \end{aligned}$$

Here, N_i, N_j represent nonterminal symbols, T denotes a preterminal symbol, $p_{G_1}(N_i \rightarrow N_j T)$ represents the probability of the rule $N_i \rightarrow N_j T$ assigned by G_1 , and R^c is the complement set of rules that do not involve $N_i \rightarrow N_j T$ and $N_i \rightarrow T N_j$.

Initially, we can directly determine the probability of a span with width $w = 2$:

$$I_{G,i}(p, p+1) = p_G(N_i \rightarrow TT) p_G(T \rightarrow w_p) p_G(T \rightarrow w_{p+1}) \quad (9)$$

Conventional inside algorithm

$$I_i(1, l) = \sum_i^n \sum_j^n \sum_k^n p(N_i \rightarrow N_j N_k) I_j(1, p) I_k(p + 1, l) \quad (3)$$

Inside algorithm in Unsupervised Neural Grammar Induction

$$\begin{aligned} I_i(1, l) = & \sum_j^n \sum_k^m p(N_i \rightarrow N_j T_k) I_j(1, l - 1) p(T_k \rightarrow w_l) + p(N_i \rightarrow T_k N_j) p(T_k \rightarrow w_1) I_j(2, l) \\ & + \sum_j^n \sum_k^n \sum_{a=2}^{n-2} p(N_i \rightarrow N_j N_k) I_j(1, a) I_k(a + 1, l) \end{aligned} \quad (4)$$

Inside algorithm in UNGI with one preterminal symbol

$$\begin{aligned} I_i(1, l) = & \sum_j^n p(N_i \rightarrow N_j T) I_j(1, l - 1) p(T \rightarrow w_l) + p(N_i \rightarrow T N_j) p(T \rightarrow w_1) I_j(2, l) \\ & + \sum_j^n \sum_k^n \sum_{a=2}^{n-2} p(N_i \rightarrow N_j N_k) I_j(1, a) I_k(a + 1, l) \end{aligned} \quad (5)$$

Inside probability for span width 3

$$I_{G,i}(1, l) = \sum_j^n p_G(N_i \rightarrow N_j T) I_{G,j}(1, l - 1) p_G(T \rightarrow w_l) + p_G(N_i \rightarrow T N_j) p_G(T \rightarrow w_1) I_{G,j}(2, l) \quad (6)$$

Equivalence between two different grammars

$$\begin{aligned} & p_{G_1}(N_i \rightarrow N_j T) I_{G_1,j}(1, r - 1) p_{G_1}(T \rightarrow w_r) + p_{G_1}(N_i \rightarrow T N_j) p_{G_1}(T \rightarrow w_1) I_{G_1,j}(2, r) \\ & = p_{G_2}(N_i \rightarrow N_j T) I_{G_2,j}(1, r - 1) p_{G_2}(T \rightarrow w_r) + p_{G_2}(N_i \rightarrow T N_j) p_{G_2}(T \rightarrow w_1) I_{G_2,j}(2, r) \end{aligned} \quad (7)$$

$$\begin{aligned} & p_G(N_i \rightarrow N_j T) + p_G(N_i \rightarrow T N_j) \alpha \\ & = p_{G'}(N_i \rightarrow N_j T) + p_{G'}(N_i \rightarrow T N_j) \alpha \end{aligned} \quad (8)$$

Table 6: The equation of inside algorithm and derived equation for proof.

For both grammars G_1 and G_2 , the equality $I_{G_1,i}(p, p+1) = I_{G_2,i}(p, p+1)$ consistently holds for all p because the probability of a span with width $w = 2$ does not involve the rules $N_i \rightarrow N_jT$ and $N_i \rightarrow TN_j$.

For spans of width $w = 3$, we refer to Equation 6. We set the inside probabilities for grammars G_1 and G_2 to be equal, such that $I_{G_1,i}(p, p+2) = I_{G_2,i}(p, p+2)$. Noting that only the probabilities of $N_i \rightarrow N_jT$ and $N_i \rightarrow TN_j$ differ, we see that the other terms cancel out, as shown in Equation 7. Therefore, Equation 7 can be simplified to Equation 8, given that $I_{G_1,i}(p, p+1) = I_{G_2,i}(p, p+1)$. Consequently:

$$\alpha = \frac{p_{G_1}(T \rightarrow w_1)I_{G_1,j}(2, r)}{I_{G_1,j}(1, r-1)p_{G_1}(T \rightarrow w_r)} \quad (10)$$

It must be noted that $p_{G_1}(T \rightarrow w_1)I_{G_1,j}(2, r) = I_{G_1,j}(1, r-1)p_{G_1}(T \rightarrow w_r)$ since the preterminal symbol T is unified, leading to $\alpha = 1$.

This process can be applied recursively to spans of any width. As a result, $p_{G_1}(N_i \rightarrow N_jT) + p_{G_1}(N_i \rightarrow TN_j) = p_{G_2}(N_i \rightarrow N_jT) + p_{G_2}(N_i \rightarrow TN_j)$ if and only if $I_{G_1,i}(p, q) = I_{G_2,i}(p, q)$. In simpler terms, the sentence probabilities calculated by G_1 and G_2 will be identical as long as the sum of probabilities $p_{G_1}(N_i \rightarrow N_jT) + p_{G_1}(N_i \rightarrow TN_j)$ matches that of G_2 .

A.2 Structural Simplicity Bias in Multilingual

We demonstrate that structural simplicity bias is observed in all language datasets. For this, we train and evaluate FGG-TNPCFGs with 5, 15, and 30 nonterminals using CTB and SPMRL datasets. In Figure 7, it is shown that in all languages, the average number of unique rules decreases as the number of symbols decreases, and verifying that structural simplicity bias is a general issue occurring in all languages, not just limited to PTB.

A.3 Qualitative analysis of tree structure

In Figure 8, we qualitatively analyze the structural differences between parses induced by different grammars (Gold, FGG-TNPCFGs and parse-focused N-PCFG) with 5 nonterminal symbols for the same sentence. We verify that the FGG-TNPCFGs utilize a small number of rule types to induce parses, which leads a different structure compared to the gold parse.

A.4 Algorithm of Proposed Method

Algorithm 1 presents the pseudocode for our entire methodology. We calculate the weights to be applied in the Inside Algorithm using softmax, based on the frequency of each span within a sentence as provided by the parse trees from pre-trained parsers.

Algorithm 1 Parse-Focusing Inside algorithm

```

1: function PARSE-FOCUSING( $s \in S, \mathcal{T}_p$ )
2:    $\mathcal{T}_p(s) \leftarrow \mathcal{T}_p$  ▷ Subset of parse trees
3:    $C(s) \leftarrow \text{COUNTSPANS}(\tau_{p_i}, |s|)$ 
4:    $W(s) \leftarrow \frac{\exp f(c(p, q) | \tau_{G_p}(s_i))}{\sum_{p,q} \exp f(c(p, q) | \tau_{G_p}(s_i))}$ 
5:   for  $w \leftarrow 2$  to  $|s|$  do
6:     for  $n \leftarrow 0$  to  $|s|$  do
7:        $I_i(n, n+w) \leftarrow$ 
8:          $\sum_{j,k} \sum_m p(i, j, k) I_j(n, m) I_k(m+1, n+w)$ 
9:        $I_i(n, n+w) \leftarrow w_s(n, n+w) I_i(n, n+w)$ 
9:   return  $p(s)$ 
10: function COUNTSPANS( $\tau_{p_i}, |s|$ )
11:   Initialize  $c(i, j) = 0$ , for all  $0 \leq i, j \leq |s|$ 
12:   for  $i \leftarrow 0$  to  $|s|$  do
13:     for  $j \leftarrow i+2$  to  $|s|$  do
14:       for each  $\tau_{p_i} \in \mathcal{T}_p$  do
15:         if span  $(i, j)$  in  $\tau_{p_i}$  then
16:            $c(i, j) \leftarrow c(i, j) + 1$ 
17:   return  $C(s)$ 

```

A.5 Quality of Focusing-Bias by the number of Multiple Parsers

In Figure 9, the common structures across the selected parsers are more similar to gold parse trees. This gives the intuition that assigning strong weights to parts that commonly appear across multiple parsers will lead to better performance.

A.6 Experiment Details

Dataset Detail We follow the standard setup and preprocessing for the PTB⁵, utilizing sections 02-21 as the training set, section 22 as the validation set, and section 23 as the test set excluding punctuation and trivial constituents. The most frequent 10,000 words were selected as the vocabulary, and the remaining words were treated as <unk>. To process the data, we followed the pipeline used by models employing base models and parsers as outlined in (Kim et al., 2019a; Yang et al., 2022; Shen et al., 2021; Yang et al., 2021a; Drozdov et al., 2019). We also use CTB⁶ and Basque, French, Ger-

⁵The license of PTB is LDC User Agreement for Non-Members. <https://catalog.ldc.upenn.edu/LDC99T42>

⁶The license of CTB is LDC User Agreement for Non-Members. <https://catalog.ldc.upenn.edu/LDC2005T01>.

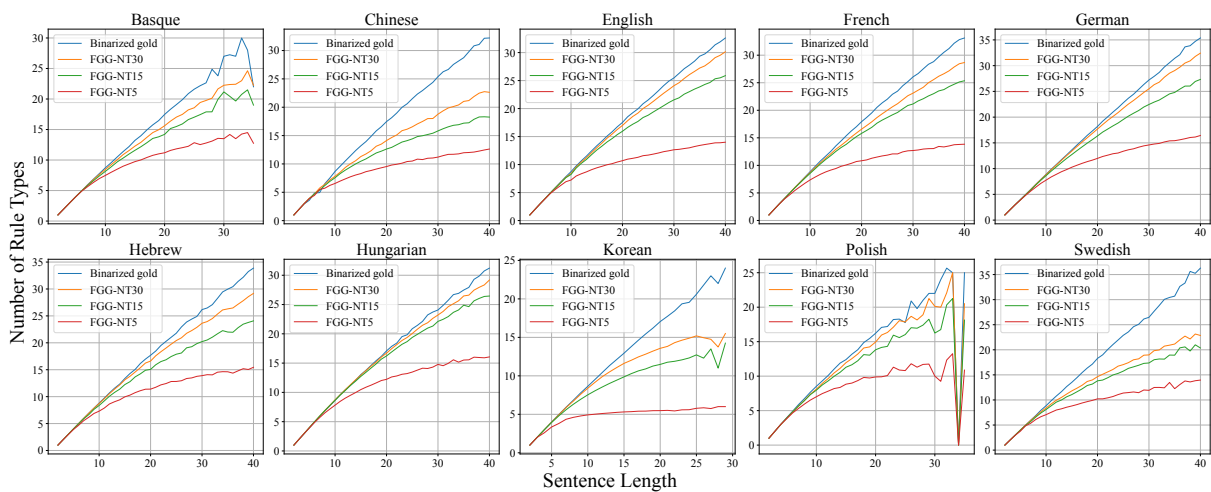


Figure 7: Average of the number of unique rules in each tree for CTB and SPMRL datasets. Binarized gold is the

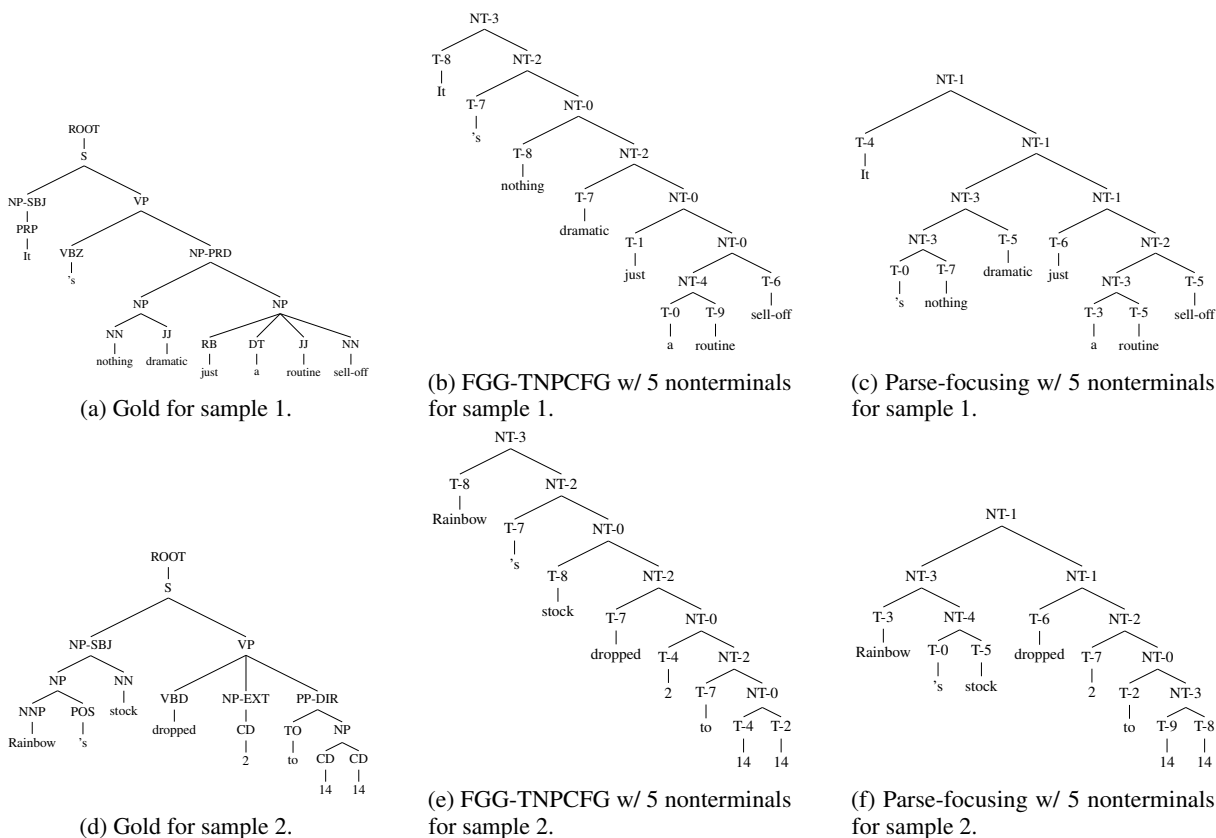


Figure 8: Induced parse trees for two sentences, "It's nothing dramatic just a routine sell-off." and "Rainbow's stock dropped 2 to 14 14."




Figure 9: Frequency (10^6 scale) of common parse spans in gold parse trees by the number of combined parsers. Using more parsers shows significantly more spans in gold parse trees without supervision.

man, Hebrew, Hungarian, Korean, Polish, Swedish dataset in SPMRL⁷.

Implementation Detail To implement our methodology on top of the base model FGG-TNPCFGs, we utilized PyTorch version 2.2 (Paszke et al., 2019). For smooth processing and analysis of tree structures, we employed NLTK⁸.

Training Detail The hyperparameter details are identical to those found in Yang et al. (2022). The ratio of nonterminal to preterminal symbols is primarily set at 1:2, and to analyze performance based on the number of symbols, experiments were conducted using a minimum of 1 to a maximum of 4500 nonterminal symbols. Training was carried out for up to 10 epochs, with early stopping applied based on validation likelihood. For optimization, the Adam optimizer was utilized with a learning rate of $2e^{-3}$, $\beta_1 = 0.75$, and $\beta_2 = 0.999$. The model’s hyperparameters were set to a rank size of 1000, a symbol embedding size of 256, and a word embedding size of 200. The same hyperparameters were used for all languages without any specific tuning. All experiments were performed using an NVIDIA RTX 2080ti, and 32 experiments were conducted to verify variance for both the original FGG-TNPCFGs and our method, as shown in Table 2, while the rest of the experiments were carried out four times each. Each run takes approximately 3 hours.

⁷The license of SPMRL is Creative Commons Attribution 4.0 International License. <https://www.spmrl.org/spmrl2013-sharedtask.html>

⁸<https://www.nltk.org/>

A.7 Focusing-Bias of Single Parser

To validate the learning of various types of focusing-bias in the N-PCFG, we assessed the Intersection over Union (IoU) scores between the pre-trained single parsers and our parse-focused N-PCFGs, as shown in Table 7. The results in Table 7 indicate that our parse-focusing N-PCFGs can replicate the tree structures generated by the pre-trained parsers. They show the highest congruence with FGG-TNPCFGs and the least congruence with Structformer. Despite this, the Structformer, which has the lowest IoU score, still surpasses the highest IoU score listed in Table 1. This indirectly indicates the extent of similarity between the parse trees derived from the pre-trained parsers and those obtained from the parse-focused N-PCFG.

Bias	S-F1		IoU
	Pre-trained	Injected	
SF	52.32 \pm 2.25	56.35 \pm 2.03	57.18 \pm 2.81
NBL	52.11 \pm 15.32	53.04 \pm 16.03	61.37 \pm 2.39
FGGs	59.65 \pm 7.73	58.67 \pm 7.15	74.49 \pm 6.67

Table 7: S-F1 score of Pre-trained parsers, Parse-focused N-PCFG, and IoU score between their results.

A.8 Numerical Result for Figure 10

In Table 8, the S-F1 and IoU scores for combinations of homogeneous parsers and heterogeneous parsers. However, heterogeneous parsers show significantly lower IoU scores across all combinations compared to homogeneous parsers. A lower IoU score signifies that the two grammars make different decisions for more spans, allowing the model to handle a wider variety of structural information in parse-focusing. The model can select the most advantageous structural information from the given options, opening up opportunities for exploration, potentially leading to higher performance.

A.9 Recent Studies

With the dramatic growth of neural networks, various approaches have been developed to integrate the concepts of neural networks and grammars. (Shen et al., 2018, 2019) utilized tree structures within their models, hierarchically learning implicit structural information from sentences using tree structures. Meanwhile, Wang et al. (2019); Shen et al. (2021) aimed to leverage structural information within transformers by inferring structural information from sentences and using it to

Parser Pair	S-F1	IoU
Homogeneous (same method, different seeds)		
(SF,SF)	59.8 \pm 0.6	50.6
(NBL,NBL)	65.0 \pm 0.2	57.1
(FGG,FGG)	68.6 \pm 0.5	40.8
(DIORA,DIORA)	53.5 \pm 0.7	43.9
Heterogeneous (different method, different seeds)		
(SF,NBL)	65.0 \pm 1.7	40.3
(SF,FGG)	65.7 \pm 0.3	39.6
(SF,DIORA)	59.1 \pm 0.4	36.2
(NBL,FGG)	69.2 \pm 0.2	40.6
(NBL,DIORA)	66.1 \pm 1.5	32.9
(FGG,DIORA)	65.1 \pm 0.9	34.3
(SF,NBL,DIORA)	64.3 \pm 1.1	21.0
(SF,FGG,NBL)	69.7 \pm 0.9	24.7
(SF,FGG,DIORA)	66.3 \pm 0.7	21.7
(FGG,NBL,DIORA)	68.5 \pm 1.6	20.7

Table 8: IoU scores of homogeneous and heterogeneous parser combinations. S-F1 is the test performance of ours with the parser combination on PTB. FGG denotes FGG-TNPCFG.




Figure 10: S-F1 scores of the combination of homogeneous and heterogeneous parsers.

mask the attention matrix to control relationships between words. Notably, Shen et al. (2021) proposed a method for simultaneously learning both constituency and dependency parsing.

There have been studies applying structural concepts, either by mimicking grammars using recurrent neural networks as in Dyer et al. (2016); Kim et al. (2019b) or by combining the inside algorithm with the concept of autoencoders as in Drozdov et al. (2019).

In particular, Kim et al. (2019a) proposed neural parameterization, assigning rule probabilities using neural networks to grammar component embeddings. Kim et al. (2019a) utilized explicit grammatical structures and the inside algorithm to induce grammars, thus enabling traditional grammat-

ical approaches. This led to the emergence of lexicalized N-PCFGs (Zhu et al., 2020; Yang et al., 2021a), which leverage lexicons. However, the inside algorithm still demanded high computational complexity, which limited the expansion of grammar sizes.

Yang et al. (2021b, 2022) addressed the high computational complexity of the inside algorithm by utilizing tensor decomposition and FGGs, which allowed for the learning of larger grammars. This approach, which solves issues through the number of parameters, aligns with the methods of Petrov et al. (2006).