

Tree-Averaging Algorithms for Ensemble-Based Unsupervised Discontinuous Constituency Parsing

Behzad Shayegh Yuqiao Wen Lili Mou*

Dept. Computing Science & Alberta Machine Intelligence Institute, University of Alberta

*Canada CIFAR AI Chair

the.shayegh@gmail.com yq.when@gmail.com doublepower.mou@gmail.com

Abstract

We address unsupervised discontinuous constituency parsing, where we observe a high variance in the performance of the only previous model in the literature. We propose to build an ensemble of different runs of the existing discontinuous parser by averaging the predicted trees, to stabilize and boost performance. To begin with, we provide comprehensive computational complexity analysis (in terms of \mathbf{P} and \mathbf{NP} -complete) for tree averaging under different setups of binarity and continuity. We then develop an efficient exact algorithm to tackle the task, which runs in a reasonable time for all samples in our experiments. Results on three datasets show our method outperforms all baselines in all metrics; we also provide in-depth analyses of our approach.¹

1 Introduction

Unsupervised parsing has been attracting the interest of researchers over decades (Klein and Manning, 2002; Klein, 2005; Snyder et al., 2009; Shen et al., 2019; Cao et al., 2020). Compared with supervised methods, unsupervised parsing has its own importance: (1) It reduces the reliance on linguistically annotated data, and is beneficial for low-resource languages and domains (Kann et al., 2019). (2) Discovering language structures in an unsupervised way helps to verify theories in linguistics (Goldsmith, 2001) and cognitive science (Bod, 2009). (3) Unsupervised parsing methods are applicable to other types of streaming data, such as motion-sensor signals (Peng et al., 2011).

A constituency tree, a hierarchy of words and phrases shown in Figure 1, is an important parse structure (Carnie, 2007). Researchers have proposed various approaches to address unsupervised constituency parsing, such as latent-variable methods (Clark, 2001; Petrov and Klein, 2007;

¹Code available at <https://github.com/MANGA-UOFA/TAA4EUDCP>




Figure 1: (a) A continuous parse structure in English. (b) An arguably discontinuous parse structure in English. (c) A discontinuous parse structure in German. Interesting structures (binarity and fan-out) are illustrated.

Kim et al., 2019a) and rule-based systems (Cao et al., 2020; Li and Lu, 2023). In our previous work (Shayegh et al., 2024), we reveal that different unsupervised parsers have low correlations with each other, and further propose an ensemble approach based on dynamic programming to boost performance.

One limitation of most existing studies is that they only address continuous constituency parsing, that is, a constituent can only be a continuous span of words (Figure 1a). However, a constituent may be discontinuous (McCawley, 1982).² In Figure 1b, for example, it is linguistically arguable that *wake* and *up* form a constituent in the sentence “Wake your friend up.” Such discontinuous constituents are more common in certain languages (such as German shown in Figure 1c) than in English (Skut et al., 1997), but they have been less tackled in the literature. Very recently, Yang et al. (2023) proposed an unsupervised parsing method based on a mildly context-sensitive grammar (Joshi, 1985) that allows discontinuous parse structures, known as discontinuous trees (Tomita, 1990). As an early attempt, their performance appears to be low, and we find that their approach is noisy and exhibits low

²In the literature, a continuous parse structure is also known to be *projective*, and a discontinuous parse structure is *non-projective* (Versley, 2014).

correlation in different runs with random seeds.

In this work, we propose an ensemble method to discontinuous constituency parsing by tree averaging, following our previous work (Shayegh et al., 2024), in which we address the ensemble of continuous constituency parsing by CYK-like dynamic programming (Younger, 1967). However, the previous approach is not directly applicable to this paper because of the discontinuous structures in our setup. Consequently, the seemingly small change in the setting (continuous vs. discontinuous) leads to a completely different landscape of problems, requiring advanced algorithmic analysis and development.

Specifically, we first analyze the computational complexity of tree averaging under different setups, such as binarity and continuity. We show certain problems are in the **P** category, whereas general tree averaging is **NP**-complete. Then, we develop an algorithm utilizing the meet-in-the-middle technique (Horowitz and Sahni, 1974) with effective pruning strategies, making our search practical despite its exponential time complexity.

Our experiments on German and Dutch show that our approach largely outperforms previous work in terms of continuous, discontinuous, and overall F_1 scores. To the best of our knowledge, our algorithm is the first that can handle non-binary constituents in the setting of unsupervised constituency parsing.

To sum up, our main contributions conclude: (1) proposing an ensemble approach to unsupervised discontinuous constituency parsing, (2) theoretically analyzing the computational complexity of various tree-averaging settings, and (3) conducting experiments on benchmark datasets to verify the effectiveness of our approach.

2 Approach

2.1 Unsupervised Discontinuous Constituency Parsing

In linguistics, a *constituent* is one or more words that act as a semantic unit in a hierarchical tree structure (Carnie, 2007). Discontinuous constituents are intriguing, where a constituent is split into two or more components by other words (Figure 1; McCawley, 1982; Tomita, 1990).

Unsupervised discontinuous constituency parsing aims to induce—without using linguistically annotated data for training—a parse structure that may contain discontinuous constituents. As men-

tioned in §1, unsupervised parsing is an important research topic, as it potentially helps low-resource languages, the development of linguistic and cognitive theory, as well as the processing of non-textual streaming data.

Yang et al. (2023) propose the only known unsupervised discontinuous parsing approach, based on linear context-free rewriting systems (LCFRS; Vijay-Shanker et al., 1987), a type of mildly context-sensitive grammars (Joshi, 1985), which can model certain discontinuous constituents. They focus on binary LCFRS and limit the *fan-out* of their grammar to be at most 2, known as LCFRS-2 (Stanojević and Steedman, 2020), meaning that each constituent can contain up to two nonadjacent components³ (illustrated in Figure 1). Maier et al. (2012) observe such structures cover most cases in common treebanks. Essentially, the LCFRS-2 grammar used in Yang et al. (2023) is a 6-tuple $\mathcal{G} = (S, \mathcal{N}_1, \mathcal{N}_2, \mathcal{P}, \Sigma, \mathcal{R})$, where S is the start symbol, \mathcal{N}_1 a finite set of non-terminal symbols with fan-out being 1, \mathcal{N}_2 fan-out being 2, \mathcal{P} preterminals, and Σ terminals. \mathcal{R} is a finite set of rules, following one of the forms:

$$\begin{array}{ll}
 R_1 : S(x) \rightarrow A(x) & A \in \mathcal{N}_1 \\
 R_2 : A(xy) \rightarrow U(x)U'(y) & A \in \mathcal{N}_1 \\
 R_3 : A(xyz) \rightarrow U(y)B(x, z) & A \in \mathcal{N}_1 \\
 R_4 : A(x, y) \rightarrow U(x)U'(y) & A \in \mathcal{N}_2 \\
 R_5 : A(xy, z) \rightarrow U(x)B(y, z) & A \in \mathcal{N}_2 \\
 R_6 : A(xy, z) \rightarrow U(y)B(x, z) & A \in \mathcal{N}_2 \\
 R_7 : A(x, yz) \rightarrow U(y)B(x, z) & A \in \mathcal{N}_2 \\
 R_8 : A(x, yz) \rightarrow U(z)B(x, y) & A \in \mathcal{N}_2 \\
 R_9 : T(w) \rightarrow w & T \in \mathcal{P}, w \in \Sigma
 \end{array}$$

for $B \in \mathcal{N}_2$, and $U, U' \in \mathcal{N}_1 \cup \mathcal{P}$. Note that x in a rewriting rule $A(x)$ is not an input string. Instead, it is a placeholder suggesting that $A(x)$ has a fan-out 1. Likewise, $A(x, y)$ has a fan-out 2, with x and y being the two nonadjacent components.

Such a grammar may handle certain types of discontinuous constituents. Take R_3 as an example. The left-hand side $A(xyz)$ implies the placeholder string xyz is a constituent that can be split into three adjacent components: x , y , and z in order. The notation $B(x, z)$ in the rule asserts that x and z form a constituent, although they are not adjacent in the original string xyz .

We follow Yang et al. (2023) and train a

³A *component* refers to a span of one or more consecutive words in a sentence.




Figure 2: F_1 scores on continuous and discontinuous constituents in the NEGRA test set (Skut et al., 1997), where each point is a random run of TN-LCFRS (Yang et al., 2023).

probabilistic LCFRS-2, parametrized by a tensor decomposition-based neural network (TN-LCFRS). The objective is to maximize the likelihood of sentence reconstruction by marginalizing the grammar rules. Compared with general context-sensitive grammar, LCFRS-2 balances the modeling capacity and polynomial-time inference efficiency; thus, it is also used in supervised discontinuous parsing (Maier, 2010; Cranenburgh et al., 2016). We refer interested readers to Yang et al. (2023) for the details of TN-LCFRS training and inference.

In our work, we observe that TN-LCFRS exhibits high variance in performance, especially for discontinuous constituents (Figure 2). In addition, there is a negative correlation between performance on continuous and discontinuous constituents among different runs: some are better at continuous constituents, while others are better at discontinuous ones.

In our previous work (Shayegh et al., 2024), we show that an ensemble model can utilize different expertise of existing continuous parsers and smooth out their noise. A natural question now is: *Can we build an ensemble of discontinuous constituency parsers?*

In the rest of this section, we perform in-depth theoretical analysis of the problem, and show that with seemingly harmless alteration of setups, the problem may belong to either **P** or **NP**-complete complexity categories. We further develop a meet-in-the-middle search algorithm with efficient pruning to solve our problem efficiently in practice.

2.2 Averaging over Constituency Trees

We adopt the tree-averaging notion in Shayegh et al. (2024), which suggests the ensemble output should be the tree with the highest average F_1 -score against the ensemble components (referred to as *individuals*):

$$T^* = \operatorname{argmax}_{T \in \mathcal{T}} \sum_{k=1}^K F_1(T, T_k) \quad (1)$$

where \mathcal{T} is the search space given an input sentence, and T_k is the parse tree predicted by the k th individual. The F_1 score is commonly used for evaluating constituency parsers, and is also our ensemble objective, given by $F_1(T_{\text{prd}}, T_{\text{ref}}) = \frac{2|C(T_{\text{prd}}) \cap C(T_{\text{ref}})|}{|C(T_{\text{prd}})| + |C(T_{\text{ref}})|}$, where T_{prd} and T_{ref} denote the predicted and reference trees, respectively, while $C(T)$ represents the set of constituents in a tree T .

It is noted that existing unsupervised parsers can only produce binary trees in both continuous and discontinuous settings (Shen et al., 2018a; Kim et al., 2019a; Yang et al., 2023). The binary property asserts that, given a length- n sentence, $|C(T_k)| = 2n - 1$ for every individual k , thus simplifying Eqn. (1) to

$$T^* = \operatorname{argmax}_{T \in \mathcal{T}} \frac{\sum_{c \in C(T)} h(c)}{|C(T)| + 2n - 1} \quad (2)$$

Here, $h(c)$ counts the occurrences of a constituent c in the trees predicted by the individuals. We call $h(c)$ the *hit count* of c .

We point out that the output of our approach does not have to be a binary tree, and we will empirically analyze the output binarity in §3.2.

We would like to examine the computational complexity for tree averaging, such as **P** and **NP** categories. To begin with, we consider the decision problem (i.e., whether there exists a tree satisfying some conditions) corresponding to the search problem (i.e., finding the best tree), which is standard in complexity analysis (Arora and Barak, 2009).

Problem 1 (Averaging binary trees with bounded fan-out). *Consider a number z and constituency trees T_1, \dots, T_K with the same leaf nodes, where the trees are binary and have a fan-out of at most F . Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 1. *Problem 1 belongs to **P**.*

Proof sketch. In our previous work (Shayegh et al., 2024), we present a polynomial-time dynamic programming (DP) algorithm to average continuous binary trees, where the outputs are restricted to binary trees as well; here, continuous trees can be seen as having max fan-out 1. For non-binary outputs, the DP table can be augmented by an additional axis whose size is bounded by the sentence

length n . To handle discontinuity, we may augment the DP table with additional axes based on the maximum fan-out F . Overall, the time complexity of the DP algorithm is $\mathcal{O}(n^{4F+1})$; thus, Problem 1 is in **P**. See Appendix A.1 for the detailed proof.⁴ \square

A follow-up question then is whether there is a polynomial-time algorithm when we relax the assumptions of input being binary and having bounded fan-out. Having non-binary inputs is intriguing since the above DP relies on Eqn. (2), which only holds for binary inputs.

Problem 2 (Averaging trees with bounded fan-out). *Consider a number z and constituency trees T_1, \dots, T_K with the same leave nodes, where the fan-out is at most F but the trees may be non-binary. Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 2. *Problem 2 belongs to **P**.*

Proof sketch. If we fix the number of nodes τ of the output tree, the above DP can be reused to solve this problem. We may enumerate all possible values of τ , which must satisfy $n < \tau < 2n$. Thus, Problem 2 can be solved in $\mathcal{O}(n^{4F+2})$ time and belongs to **P**. See Appendix A.2 for the proof. \square

However, if the fan-out is unbounded, the difficulty is that the above DP table grows exponentially with respect to F . Since the problem is polynomial-time verifiable, it surely belongs to **NP**, but whether it belongs to **P**, **NP**-complete, or both remains unknown for binary inputs.

Open Problem 1 (Averaging binary trees). *What is the exact complexity category (**P** or **NP**-complete) of averaging binary trees with unbounded fan-out?*

For averaging general trees (non-binary inputs with unbounded fan-out), we can show that it belongs to **NP**-complete.

Problem 3 (Averaging trees). *Consider a number z and constituency trees T_1, \dots, T_K with the same leave nodes, where the fan-out is unbounded and the trees may be non-binary. Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 3. *Problem 3 belongs to **NP**-complete.*

⁴Corro (2023) develops a CYK-like algorithm for span-based nested named-entity recognition that takes discontinuous and non-binary structures into account. However, their algorithm limits its search space to constituents that contain at most one smaller multi-word constituent inside, making it not applicable to our scenario.

Binary individuals? \ Bounded fan-out?	Bounded	Unbounded
Binary	P	Unknown
Non-binary	P	NP -complete

Table 1: Summary of the complexity categories of tree-averaging problems.

Proof sketch. It is easy to show that, given a certificate, Problem 3 is polynomial-time verifiable. For the completeness, we reduce the max clique problem, a known **NP**-complete problem, to Problem 3. See Appendix A.3 for the detailed proof. \square

The above theoretical analysis provides a deep understanding of building tree ensembles, summarized in Table 1. In our experiments, the task falls into Problem 1 because the only existing unsupervised discontinuous parser (Yang et al., 2023) produces binary trees with fan-out at most 2. The problem belongs to **P**.

That being said, a practical approach should consider not only algorithmic complexity, but also the specific properties of the task at hand. A DP algorithm that one may develop for our task has a complexity of $\mathcal{O}(n^9)$, as discussed in the proof of Theorem 1, which appears impractical despite its polynomial time complexity.

In the rest of this section, we develop a more general search algorithm that works at the level of Open Problem 1 but also solves Problem 1 more efficiently in practice than high-order polynomial DP.

2.3 Our Search Algorithm

Our work concerns building an ensemble of binary, bounded-fan-out, discontinuous trees generated by the unsupervised parser in Yang et al. (2023).

We will develop a general search algorithm (regardless of fan-out) with strong pruning that only needs to consider a few candidate constituents, bringing down the $\mathcal{O}(2^{2^n})$ complexity of exhaustive search⁵ to $\mathcal{O}(2^n)$. We will further halve the exponent using a meet-in-the-middle technique, resulting in an overall complexity of $\mathcal{O}(2^{\frac{n}{2}}n^2)$. As a result, our algorithm remains exact search while being efficient for all samples in the dataset; when a sentence is short, our algorithm is even much faster than the DP with $\mathcal{O}(n^9)$ complexity.

⁵For a length- n sentence, there are $2^n - 1$ possible constituents because an arbitrary non-empty set of words can be a constituent. To build a constituency tree, the exhaustive search needs to look into any combination of constituents.

To solve the tree-averaging problem, we first convert it into an equivalent graph problem for clarity. Specifically, we construct an undirected graph G with vertices being all possible constituents (although they can be largely pruned), weighted by corresponding hit counts, as defined after Eqn. (2). We put an edge between two vertices if and only if their corresponding constituents can coexist within a constituency tree, that is to say, they are either disjoint or inclusive of one another. Then, we formulate the below graph-search problem, which can solve the original tree-search problem for ensemble discontinuous parsing.

Problem 4 (Normalized Max Weighted Clique). *Consider a weighted undirected graph $G = \langle V, E \rangle$ and an objective function $f(Q; \alpha_1, \alpha_2) = \frac{\sum_{v \in Q} w(v) + \alpha_1}{|Q| + \alpha_2}$, where $\alpha_1, \alpha_2 \in \mathbb{R}$, $Q \subseteq V$, and $w(v)$ is the weight of v . What is the clique Q that maximizes $f(Q; \alpha_1, \alpha_2)$?*

Lemma 1. *Given G as described above, the clique Q^* maximizing $f(Q; \alpha_1, \alpha_2)$ corresponds to a constituency tree, if $\alpha_1 \leq K\alpha_2$ where K is the number of individuals. (Proof in Appendix A.4.)*

Theorem 4. *The average constituency tree T^* in Eqn. (2) is polynomial-time reducible from Q^* in Problem 4 with $\alpha_1 = 0$ and $\alpha_2 = 2n - 1$, if G is built as above.*

Proof. Suppose the constituents corresponding to a clique Q in the graph G form a constituency tree, denoted by T_Q . We have $\sum_{v \in Q} w(v) = \sum_{c \in T_Q} h(c)$ by the construction of G . Thus, the objective function of Problem 4 is

$$f(Q; 0, 2n - 1) = \frac{\sum_{v \in Q} w(v) + 0}{|Q| + 2n - 1} \quad (3)$$

$$= \frac{\sum_{c \in C(T_Q)} h(c)}{|C(T_Q)| + 2n - 1} \quad (4)$$

which is the same as the argmax objective in Eqn. (2). It is easy to see that each constituency tree corresponds to a clique in G . Therefore, the corresponding constituency tree to Q^* , guaranteed by Lemma 1, maximizes Eqn. (2). In other words, we have $T^* = T_{Q^*}$. \square

Problem 4 generalizes a standard max clique problem (Arora and Barak, 2009), which may be solved in $\mathcal{O}(2^{|V|})$ time complexity by exhaustive search. The meet-in-the-middle technique (Horowitz and Sahni, 1974) can be used to address the max clique problem, reducing the

complexity from $\mathcal{O}(2^{|V|})$ to $\mathcal{O}(2^{\frac{|V|}{2}} |V|^2)$. In Appendix B, we develop a variant of the meet-in-the-middle algorithm to solve Problem 4.

2.4 Candidate Constituents Pruning

The efficiency of our algorithm for Problem 4 depends on the number of vertices in the graph. In our construction, the vertices in G correspond to possible constituents, which we call *candidates*.

A naïve approach may consider all $\mathcal{O}(2^n)$ possible constituents, which are non-empty combinations of words in a length- n sentence. Thus, our meet-in-the-middle algorithm has an overall complexity of $\mathcal{O}(2^{2^{n-1}+2n})$ for tree averaging.

In this part, we theoretically derive lower and upper bounds for a candidate’s hit count. If a constituent has a lower hit count than the lower bound, it is guaranteed not to appear in the average tree. On the other hand, a constituent must appear in the average tree, if it has a higher hit count than the upper bound. We may exclude both cases from our search process, and directly add the must-appear candidates to the solution in a *post hoc* fashion.

Let P be a set of constituents that are known to be in the average tree T^* . Here, P may even be an empty set. We may derive a hit-count lower bound for other constituents (not in P), as stated in the following theorem.

Theorem 5 (Lower bound). *For every constituent $c \in C(T^*) \setminus P$, where $P \subseteq C(T^*)$, we have*

$$h(c) > \min_{|P| \leq j \leq 2n-2} \frac{\sum_{i=1}^{j-|P|} \lambda_i^+ + \sum_{c' \in P} h(c')}{j + 2n - 1} \quad (5)$$

where λ_i^+ is the i th smallest positive hit count. (Proof in Appendix A.5.)

The theorem suggests that constituents having a hit count lower than the above threshold can be removed from the search. If we set $P = \emptyset$, we have $h(c) > 0$, immediately pruning all zero-hit constituents. In other words, we may only consider the constituents appearing in at least one individual, which largely reduces the graph size from $\mathcal{O}(2^n)$ to $\mathcal{O}(nK)$ for K individuals. In fact, the graph can be further pruned with the below theorem.

Theorem 6 (Upper bound). *Let c be a constituent with a hit count of K , where K is the number of individuals. (a) The constituent c is compatible—i.e., may occur in the same constituency tree—with every constituent in the average tree. (b) The constituent c appears in the average tree. (Proof in Appendix A.6.)*

Theorem 6b suggests that the search process may exclude the constituents that occur in all individuals, denoted by $P = \{c : h(c) = K\}$. In this case, we may solve a reduced version of Problem 4 with the pruned graph and $\alpha_1 = \sum_{c' \in P} h(c')$, $\alpha_2 = |P| + 2n - 1$. The α_1 and α_2 in Theorem 4 can be modified accordingly.⁶ When we add P back to the solution, the connectivity of P is guaranteed by Theorem 6a, forming the solution (guaranteed to be a clique) for the original Problem 4.

To analyze the worst-case complexity, we notice that single words are constituents that must occur in all individuals. At least, we may set $P = \{c : c \text{ is a single word}\}$, and Theorem 5 yields a lower bound of

$$h(c) > \min_{n \leq j \leq 2n-2} \frac{0 + nK}{j + 2n - 1} > \frac{nK}{4n} = \frac{K}{4} \quad (6)$$

Then, we can find an upper bound for the number of candidates kept for the search process. Since the sum of all hit counts is always $(2n - 1)K$, the number of candidates with a higher hit count than the threshold in Eqn. (6) is bounded by

$$|V| \leq \left\lfloor \frac{(2n - 1)K}{K/4} \right\rfloor < 8n \quad (7)$$

This shows that our pruning mechanism, profoundly, reduces the graph size from $\mathcal{O}(2^n)$ to $\mathcal{O}(n)$, which in turn reduces the overall time complexity of our tree averaging from $\mathcal{O}(2^{2n-1+2n})$ to $\mathcal{O}(2^{\frac{n}{2}}n^2)$ in the worst-case scenario.

In practice, the pruning is even more effective when the individuals either largely agree or largely disagree with each other. In the former case, we will have more all-hit constituents excluded from the search and consequently the lower bound increases, whereas in the latter case, many constituents will fall short of the lower bound. Empirically, our algorithm is efficient for all samples in our experiments.

3 Experiments

3.1 Settings

Datasets. We evaluated our method on Dutch and German datasets, where discontinuous constituents are relatively common. In particular, we used the LASSY treebank (Van Noord et al., 2013) for Dutch. For German, we trained our individuals on the union of NEGRA (Skut et al., 1997) and TIGER (Brants et al., 2002) treebanks, while test-

⁶The proof is parallel to that of the original version of Theorem 4, and we leave it as an exercise for readers.

ing and reporting the performance on their test sets respectively. Our settings strictly followed Yang et al. (2023) for fair comparison.

Evaluation Metrics. F_1 scores are commonly used for evaluating constituency parsing performance (Klein, 2005; Shen et al., 2018a, 2019; Kim et al., 2019a,b). In our work of unsupervised discontinuous constituency parsing, we report corpus-level F_1 scores of all constituents, continuous constituents, and discontinuous constituents, denoted by F_1^{overall} , F_1^{cont} , and F_1^{disco} , respectively. We followed the previous work (Yang et al., 2023), which discards punctuation and excludes trivial constituents (the whole sentence and single words) when calculating the F_1 scores.

Setups of Our Ensemble. As mentioned in §2.1, we used different runs of the tensor decomposition-based neural LCFRS (TN-LCFRS; Yang et al., 2023) as our individual models. Our experiment was based on the released code⁷ with default training hyperparameters. However, we find the TN-LCFRS is highly unstable, with a very high variance and lower overall performance than Yang et al. (2023). As a remedy, we trained the model 10 times and selected the top five based on validation F_1^{overall} , and this yields performance close to Yang et al. (2023). In addition, we observe that random weight initializations lead to near-zero correlation of the predicted discontinuous constituents, but our ensemble method expects the individuals to at least agree with each other to some extent. Therefore, our different runs started with the same weight initialization⁸ but randomly shuffled the order of training samples to achieve stochasticity.

Our ensemble method does not have hyperparameters. However, F_1^{overall} and F_1^{disco} may not correlate well. Since our ensemble objective is solely based on F_1^{overall} , we may enhance F_1^{disco} by weighting the individuals with validation F_1^{disco} scores in hopes of achieving high performance in all aspects. Note that weighting individuals does not hurt our theoretical analysis and algorithm, because weighting is equivalent to duplicating individuals (since all the weights are rational numbers) and can be implemented by modifying hit counts without actual duplication.

Baselines. Our ensemble individual, TN-

⁷<https://github.com/sustcsonglin/TN-LCFRS>

⁸Our pilot study shows the proposed ensemble method is not sensitive to the initialization, as long as it is shared among different runs. This is also supported by the evidence that our approach performs well consistently on three datasets.

Method(# preterminal symbols)	NEGRA			TIGER			LASSY		
	F_1^{overall}	F_1^{cont}	F_1^{disco}	F_1^{overall}	F_1^{cont}	F_1^{disco}	F_1^{overall}	F_1^{cont}	F_1^{disco}
Baselines (four runs each) [†]									
1 Left branching	7.8	–	0.0	7.9	–	0.0	7.2	–	0.0
2 Right branching	12.9	–	0.0	14.5	–	0.0	24.1	–	0.0
3 N-PCFG ⁽⁴⁵⁾	40.8±0.5	–	0.0	39.5±0.4	–	0.0	40.1±3.9	–	0.0
4 C-PCFG ⁽⁴⁵⁾	39.1±1.9	–	0.0	38.8±1.3	–	0.0	37.9±3.4	–	0.0
5 TN-PCFG ⁽⁴⁵⁰⁰⁾	45.4±0.5	–	0.0	44.7±0.6	–	0.0	44.3±6.4	–	0.0
6 N-LCFRS ⁽⁴⁵⁾	33.7±2.8	–	2.0±0.8	32.7±1.8	–	1.2±0.8	36.9±1.5	–	0.9±0.8
7 C-LCFRS ⁽⁴⁵⁾	36.7±1.5	–	2.7±1.4	35.2±1.2	–	1.7±1.1	36.9±3.7	–	2.2±1.0
8 TN-LCFRS ⁽⁴⁵⁰⁰⁾	46.1±1.1	–	8.0±1.1	45.4±0.9	–	6.1±0.8	45.6±2.3	–	8.9±1.5
Individuals: TN-LCFRS ⁽⁴⁵⁰⁰⁾									
9 Five runs	46.4±0.5	49.8±1.3	6.0±4.0	45.8±1.3	49.9±1.1	4.0±3.2	46.7±2.0	50.9±1.7	6.2±1.9
10 F_1^{overall} -best run	<u>46.9</u>	50.2	1.3	<u>47.2</u>	51.1	5.9	<u>48.2</u>	52.4	5.8
11 F_1^{cont} -best run	46.7	<u>51.3</u>	7.3	47.2	<u>51.1</u>	5.9	48.2	<u>52.4</u>	5.8
12 F_1^{disco} -best run	46.0	48.3	<u>10.4</u>	45.4	48.8	<u>6.6</u>	48.0	52.1	<u>8.6</u>
13 Binary ensemble	47.6*	50.1	9.9	47.8**	51.5	6.5	50.9**	54.6**	9.7**
14 Non-binary ensemble	49.1**	51.5	10.6	48.7**	52.4**	6.6	51.4**	55.0**	10.2*

Table 2: Main results. [†]Quoted from Yang et al. (2023). * p -value < 0.05 in an Improved Nonrandomized Sign test (Starks, 1979) against the best ensemble individual in each metric, indicated by underline. ** p -value < 0.01.

LCFRS (Yang et al., 2023), is naturally included as a baseline. In addition, we consider different variants of continuous PCFG parsers and discontinuous LCFRS parsers, based on vanilla neural networks, compound prior (Kim et al., 2019a), and tensor decomposition-based neural networks (Yang et al., 2021), denoted by [N|C|TN]-[PCFG|LCFRS].

3.2 Results and Analyses

Main Results. Table 2 presents the main results on three datasets. We replicated TN-LCFRS with five runs as our ensemble individuals; our results are similar to Yang et al. (2023), showing the success of our replication (Rows 9 vs. 8).

In our study, we observe TN-LCFRS behaves inconsistently in different runs: some runs are good at continuous constituents (Row 11), while others are good at discontinuous ones (Row 12); both may disagree with the best run according to F_1^{overall} (Row 10).

Our ensemble approach (Rows 14) achieves F_1^{cont} and F_1^{disco} scores comparable to, or higher than, all individuals. This eventually leads to a much higher F_1^{overall} score, with a p -value < 0.01 in an Improved Nonrandomized Sign (INS) test (Starks, 1979) with the binomial comparison over sentence-level performance.⁹ Results are consistent on all three datasets.

Binary vs. Non-Binary Ensemble. To the best of our knowledge, we are the first to address

⁹We chose INS test because it can properly handle a large number of ties (neutral cases), when they do not necessarily suggest the equivalency of the models but the inadequacy of test samples (e.g., no discontinuous constituents in a sentence).

Model	Overall		Cont		Disco	
	TP	FP	TP	FP	TP	FP
Binary ensemble	6392	9318	6290	7934	102	1384
Non-binary ensemble	6321	8858	6221	7605	100	1253
Difference	71	460	69	329	2	131

Table 3: Number of true positive (TP) and false positive (FP) constituents on LASSY.

non-binary unsupervised constituency parsing. A natural question is whether and how non-binarity in the output improves the performance, given that all individuals are binary. To this end, we conducted an experiment by restricting the output to binary trees.¹⁰ We compare the results with the non-binary setting in Rows 13–14, Table 2, showing the consistent superiority of our non-binary ensemble in all metrics and datasets. To further understand the effect of non-binarity, we report in Table 3 true positive and false positive counts on the LASSY dataset.¹¹ Results suggest that the non-binary setting largely eliminates false positive constituents, as the model may opt to predict fewer constituents than that in the binary setting.

Number of Ensemble Individuals. We varied the number of individuals in the ensemble to analyze its effect. In Figure 3a, we picked a random subset of individuals from a pool of 10 and av-

¹⁰For binary outputs, we may still perform the meet-in-the-middle search for non-binary trees, with the hit count as the objective, and binarize them by post-processing. Note that Theorem 5 for pruning does not hold in this case, but we may still safely ignore zero-hit constituents, making the algorithm affordable (although it is not as efficient as the non-binary setting where Theorem 5 can be applied).

¹¹We chose LASSY as the testbed for all our analyses due to the limit of space and time, because it is relatively more stable in terms of F_1^{disco} than other datasets.

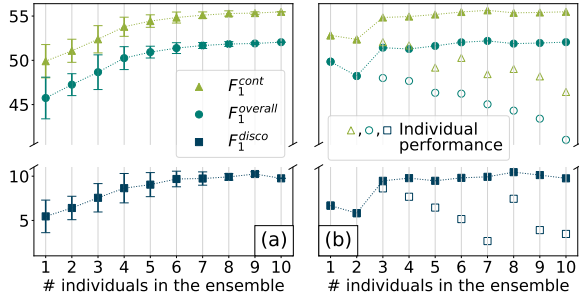


Figure 3: Effect of the number of ensemble individuals on LASSY. (a) Averaged over 30 trials with error bars indicating standard deviations. (b) Best-to-worst incremental ensemble.




Figure 4: Effectiveness of pruning on LASSY for (a) different sentence lengths, and (b) different numbers of ensemble individuals. Note that the dashed orange line does not fit the range of y -axis in (b).

eraged the results over 30 trials. Overall, more individuals lead to higher performance and lower variance in all F_1 scores, although the performance may be saturated if there are already a large number of individuals.

In addition, we present an analysis of the best-to-worst incremental ensembles in Figure 3b, where we add individuals to the ensemble from the best to worst based on $F_1^{overall}$. As seen, adding weak individuals to the ensemble does not hurt, if not help, the performance in our experiments, demonstrating that our ensemble approach is robust to the quality of individuals.

Efficiency Analysis. We analyze the effect of our pruning mechanism and the meet-in-the-middle algorithm. Pruning serves as preprocessing for our search, and we show the number of remaining candidate constituents in Figure 4. The empirical results confirm Theorem 5 bounding the number of candidates linear in both sentence length and the number of individuals, as well as Theorem 6 further bounding it by a constant with respect to the number of individuals.

We further analyze wall clock run time in Figure 5, as the actual execution time may be different from algorithmic complexity. In §A.1, we provide




Figure 5: Wall clock run time of tree-averaging algorithms on LASSY for different sentence lengths, using an Intel(R) Core(TM) i9-9940X (@3.30GHz) CPU.

a dynamic programming (DP) algorithm with time complexity of $\mathcal{O}(n^9)$ for a length- n sentence. As shown in the figure, the DP fails to serve as a practical algorithm due to the high-order polynomial. By contrast, our exponential-time search algorithm, even without the meet-in-the-middle technique, is able to run in a reasonable time for many samples (especially short ones) because of the strong pruning mechanism. Our meet-in-the-middle technique further speeds up the search, making our algorithm efficient and faster than DP for all the data samples.

It is emphasized that this experiment also empirically verifies the correctness of all our algorithms, as they perform exact inference and we have obtained exact results by using different algorithms.

Additional Results. We provide additional results in the appendix. **C:** Performance by constituency types; and **D:** Case study.

4 Related Work

Constituency parsing carries a long history in natural language processing research (Charniak, 2000; Klein, 2005; Kallmeyer and Maier, 2010; Li et al., 2019). Different setups of the task have been introduced and explored, including supervised and unsupervised, continuous and discontinuous constituency parsing (Shen et al., 2018a,b, 2019; Corro, 2020; Fernández-González and Gómez-Rodríguez, 2020, 2021, 2023; Chen and Komachi, 2023; Yang et al., 2023; Yang and Tu, 2023). In the unsupervised setup, researchers typically define a joint distribution over the parse structure and an observable variable, e.g., the sentence itself, and maximize the observable variable’s likelihood through marginalization (Kim et al., 2019a,b). To the best of our knowledge, previous unsupervised parsing studies are all restricted to binary structures to squeeze the marginalization space, and we are the first to address non-binary unsupervised parsing and show

non-binarity is beneficial to parsing performance.

Ensemble methods strategically combine multiple models to improve performance, rooted in the bagging concept where different data portions are used to train multiple models (Breiman, 1996a; Hastie et al., 2009). To build an ensemble, straightforward methods include averaging and voting (Breiman, 1996a,b). For outputs with internal structures, minimum Bayes risk decoding (MBR; Bickel and Doksum, 2015) can be used to build an ensemble, where the vote is the negative risk in MBR. However, existing MBR approaches are mostly *selective*, where the output is selected from a candidate set (Kumar and Byrne, 2004; Titov and Henderson, 2006; Shi et al., 2022). Petrov and Klein (2007) formulate MBR for supervised constituency parsing and propose to search for the global optimum decoding when the risk allows for dynamic programming. Smith and Smith (2007) extend the idea to non-projective dependency parsing. In our previous work (Shayegh et al., 2024), we formally introduce *generative* MBR and develop an algorithm that searches for a binary continuous tree. This paper extends our previous work and searches in the space of discontinuous constituency trees, leading to significant algorithmic design and theoretical analysis.

5 Conclusion

In this work, we address ensemble-based unsupervised discontinuous constituency parsing by tree averaging, where we provide comprehensive complexity analysis and develop an efficient search algorithm to obtain the average tree. Our experiments on Dutch and German demonstrate the effectiveness of our ensemble method. To the best of our knowledge, we are also the first to address, and show the importance of, non-binary structures in unsupervised constituency parsing.

6 Limitations

Our work demonstrates both theoretical depth and empirical effectiveness, but may also have limitations.

First, our work is focused on syntactic parsing, and we have provided a series of theoretical analyses and algorithmic designs for averaging constituency trees under different setups. Following the trajectory, our ensemble approach may be extended to other structures beyond parsing. We are happy to explore this direction as future work. For

example, a concurrent study of ours develops algorithms for ensemble-based text generation (Wen et al., 2024).

Second, our model is only tested on Dutch and German datasets. This is partially because of the established setups in previous work (Yang et al., 2023) and the lack of annotated treebanks. Notice that English is usually excluded from the study of unsupervised discontinuous parsing, because English discontinuous structures are too rare for any model to discover. A potential future direction is multilingual linguistic structure discovery, perhaps, with ensembles of different languages.

Third, our theoretical analysis leaves an open problem about averaging binary trees. However, pointing to open problems is usually considered a contribution (instead of a weakness) in theoretical computer science. Our theoretical analysis is also crucial to the understanding of our algorithms, because we now know that our proposed method works at the level of Open Problem 1 and can be easily extended to the level of Problem 3 following the technique proposed in Appendix A.2.

Acknowledgments

We would like to thank all reviewers and chairs for their valuable and constructive comments. The research is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Amii Fellow Program, the Canada CIFAR AI Chair Program, an Alberta Innovates Program, a donation from DeepMind, and the Digital Research Alliance of Canada (alliancecan.ca). We also thank Yara Kamkar for providing advice on the algorithms.

References

- Sanjeev Arora and Boaz Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press.
- Peter J Bickel and Kjell A Doksum. 2015. *Mathematical Statistics: Basic Ideas and Selected Topics*. CRC Press.
- Rens Bod. 2009. From exemplar to grammar: A probabilistic analogy-based model of language learning. *Cognitive Science*, 33(5):752–793.
- Sabine Brants, Stefanie Dipper, Silvia Hansen, Wolfgang Lezius, and George Smith. 2002. The TIGER treebank. In *Proceedings of the Workshop on Treebanks and Linguistic Theories*, pages 24–41.
- Leo Breiman. 1996a. Bagging predictors. *Machine Learning*, 24:123–140.
- Leo Breiman. 1996b. Heuristics of instability and stabilization in model selection. *The Annals of Statistics*, 24(6):2350–2383.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. Unsupervised parsing via constituency tests. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 4798–4808.
- Andrew Carnie. 2007. *Syntax: A Generative Introduction*, 2 edition. Wiley Blackwell.
- Eugene Charniak. 2000. A maximum-entropy-inspired parser. In *the North American Chapter of the Annual Meeting of the Association for Computational Linguistics*, pages 132–139.
- Zhousi Chen and Mamoru Komachi. 2023. Discontinuous combinatory constituency parsing. *Transactions of the Association for Computational Linguistics*, pages 267–283.
- Alexander Clark. 2001. Unsupervised induction of stochastic context-free grammars using distributional clustering. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics Workshop on Computational Natural Language Learning*.
- Caio Corro. 2020. Span-based discontinuous constituency parsing: a family of exact chart-based algorithms with time complexities from $O(n^6)$ down to $O(n^3)$. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 2753–2764.
- Caio Corro. 2023. A dynamic programming algorithm for span-based nested named-entity recognition in $o(n^2)$. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 10712–10724.
- Andreas Cranenburgh, Remko Scha, and Rens Bod. 2016. Data-oriented parsing with discontinuous constituents and function tags. *Journal of Language Modelling*, 4(1):57–111.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2020. Discontinuous constituent parsing with pointer networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7724–7731.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2021. Reducing discontinuous to continuous parsing with pointer network reordering. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 10570–10578.
- Daniel Fernández-González and Carlos Gómez-Rodríguez. 2023. Discontinuous grammar as a foreign language. *Neurocomputing*, 524:43–58.
- John Goldsmith. 2001. Unsupervised learning of the morphology of a natural language. *Computational Linguistics*, 27(2):153–198.
- Trevor Hastie, Robert Tibshirani, Jerome Friedman, Trevor Hastie, Robert Tibshirani, and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer.
- Ellis Horowitz and Sartaj Sahni. 1974. Computing partitions with applications to the knapsack problem. *Journal of the Association for Computing Machinery*, 21(2):277–292.
- Aravind K Joshi. 1985. Tree adjoining grammars: How much context-sensitivity is required to provide reasonable structural descriptions? In *Natural Language Parsing*, pages 206–250.
- Laura Kallmeyer and Wolfgang Maier. 2010. Data-driven parsing with probabilistic linear context-free rewriting systems. In *Proceedings of the International Conference on Computational Linguistics*, pages 537–545.
- Katharina Kann, Anhad Mohanane, Samuel R. Bowman, and Kyunghyun Cho. 2019. Neural unsupervised parsing beyond English. In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource Natural Language Processing*, pages 209–218.
- Yoon Kim, Chris Dyer, and Alexander Rush. 2019a. Compound probabilistic context-free grammars for grammar induction. In *Proceedings of the 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 North American Chapter of the Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1105–1117.
- Dan Klein. 2005. *The Unsupervised Learning of Natural Language Structure*. Stanford University.

- Dan Klein and Christopher D. Manning. 2002. [A generative constituent-context model for improved grammar induction](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, page 128–135.
- Shankar Kumar and William Byrne. 2004. [Minimum Bayes-risk decoding for statistical machine translation](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 169–176.
- Bowen Li, Lili Mou, and Frank Keller. 2019. [An imitation learning approach to unsupervised parsing](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 3485–3492.
- Jiaxi Li and Wei Lu. 2023. [Contextual distortion reveals constituency: Masked language models are implicit parsers](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 5208–5222.
- Wolfgang Maier. 2010. [Direct parsing of discontinuous constituents in German](#). In *Proceedings of the NAACL HLT Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 58–66.
- Wolfgang Maier, Miriam Kaeshammer, and Laura Kallmeyer. 2012. [PLCFRS parsing revisited: Restricting the fan-out to two](#). In *Proceedings of the International Workshop on Tree Adjoining Grammars and Related Formalisms*, pages 126–134.
- James D. McCawley. 1982. [Parentheticals and discontinuous constituent structure](#). *Linguistic Inquiry*, 13(1):91–106.
- Huan-Kai Peng, Pang Wu, Jiang Zhu, and Joy Ying Zhang. 2011. [Helix: Unsupervised grammar induction for structured activity recognition](#). In *Proceedings of the International Conference on Data Mining*, pages 1194–1199.
- Slav Petrov and Dan Klein. 2007. [Improved inference for unlexicalized parsing](#). In *Proceedings of Human Language Technologies: The Conference of the North American Chapter of the Association for Computational Linguistics*, pages 404–411.
- Behzad Shayegh, Yanshuai Cao, Xiaodan Zhu, Jackie CK Cheung, and Lili Mou. 2024. [Ensemble distillation for unsupervised constituency parsing](#). In *International Conference on Learning Representations*.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. 2018a. [Neural language modeling by jointly learning syntax and lexicon](#). In *International Conference on Representation Learning*.
- Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordoni, Aaron Courville, and Yoshua Bengio. 2018b. [Straight to the tree: Constituency parsing with neural syntactic distance](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 1171–1180.
- Yikang Shen, Shawn Tan, Alessandro Sordoni, and Aaron Courville. 2019. [Ordered neurons: Integrating tree structures into recurrent neural networks](#). In *International Conference on Representation Learning*.
- Freda Shi, Daniel Fried, Marjan Ghazvininejad, Luke Zettlemoyer, and Sida I. Wang. 2022. [Natural language to code translation with execution](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 3533–3546.
- Wojciech Skut, Brigitte Krenn, Thorsten Brants, and Hans Uszkoreit. 1997. [An annotation scheme for free word order languages](#). In *Proceedings of the Applied Natural Language Processing*, page 88–95.
- David A. Smith and Noah A. Smith. 2007. [Probabilistic models of nonprojective dependency trees](#). In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 132–140.
- Benjamin Snyder, Tahira Naseem, and Regina Barzilay. 2009. [Unsupervised multilingual grammar induction](#). In *Proceedings of the Joint Conference of the Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, pages 73–81.
- Miloš Stanojević and Mark Steedman. 2020. [Span-based LCFRS-2 parsing](#). In *Proceedings of the International Conference on Parsing Technologies and the International Workshop on Parsing Technologies Shared Task on Parsing into Enhanced Universal Dependencies*, pages 111–121.
- Thomas H Starks. 1979. [An improved sign test for experiments in which neutral responses are possible](#). *Technometrics*, 21(4):525–530.
- Ivan Titov and James Henderson. 2006. [Loss minimization in parse reranking](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 560–567.
- Masaru Tomita. 1990. *Current Issues in Parsing Technology*. Kluwer Academic Publishers.
- Gertjan Van Noord, Gosse Bouma, Frank Van Eynde, Daniel De Kok, Jelmer Van der Linde, Ineke Schuurman, Erik Tjong Kim Sang, and Vincent Vandeghinste. 2013. [Large scale syntactic annotation of written dutch: Lassy](#). In *Essential Speech and Language Technology for Dutch*, pages 147–164.
- Yannick Versley. 2014. [Experiments with easy-first nonprojective constituent parsing](#). In *Proceedings of the Joint Workshop on Statistical Parsing of Morphologically Rich Languages and Syntactic Analysis of Non-Canonical Languages*, pages 39–53.

- K. Vijay-Shanker, David J. Weir, and Aravind K. Joshi. 1987. [Characterizing structural descriptions produced by various grammatical formalisms](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, page 104–111.
- Yuqiao Wen, Behzad Shayegh, Chenyang Huang, Yan-shuai Cao, and Lili Mou. 2024. [EBBS: An ensemble with bi-level beam search for zero-shot machine translation](#). *arXiv preprint arXiv:2403.00144*.
- Songlin Yang, Roger Levy, and Yoon Kim. 2023. [Unsupervised discontinuous constituency parsing with mildly context-sensitive grammars](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 5747–5766.
- Songlin Yang and Kewei Tu. 2023. [Don’t parse, choose spans! continuous and discontinuous constituency parsing via autoregressive span selection](#). In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, pages 8420–8433.
- Songlin Yang, Yanpeng Zhao, and Kewei Tu. 2021. [PCFGs can do better: Inducing probabilistic context-free grammars with many symbols](#). In *Proceedings of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1487–1498.
- Daniel H. Younger. 1967. [Recognition and parsing of context-free languages in time \$n^3\$](#) . *Information and Control*, 10(2):189–208.

A Proofs

A.1 Proof of Theorem 1

Problem 1 (Averaging binary trees with bounded fan-out). *Consider a number z and constituency trees T_1, \dots, T_K with the same leave nodes, where the trees are binary and have a fan-out of at most F . Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 1. *Problem 1 belongs to **P**.*

Proof. Following our previous work (Shayegh et al., 2024), we may design a DP table with two axes being the start and end of every component of a constituent, thus $2F$ -axes for fan-out F . For outputs being non-binary, our DP algorithm requires one additional axis to indicate the number of nodes in a subtree.¹²

We define a DP variable $H(c, \tau)$ as the best total hit count for a τ -node constituency subtree over a constituent c , which may be discontinuous. A constituent with fan-out f has f components, each being a span of consecutive words. Therefore, c can be represented by $2F$ numbers indicating the beginnings and ends of the components, as the fan-out is bounded by F . For τ , it is upper-bounded by $2n - 1$ for a length- n sentence. Overall, the DP table has a size of $\mathcal{O}(n^{2F+1})$. Specially, we further define $H(c, 0) = 0$ for every c .

For initialization, we consider every single-word constituent c and set $H(c, 1) = h(c)$, where $h(c)$ is the hit count of c in T_1, \dots, T_K . We also set $H(c, i) = -\infty$ for every $i > 1$.

For recursion, we divide a constituent c into smaller sub-constituents based on alternation points $\mathbf{j} = (0, j_1, j_2, \dots, j_{2F-1}, |c|)$, where $|c|$ is the number of words in c , shown in Figure 6. We further join sub-constituents based on the parity (even or odd) of their indexes.

$$o_j = \bigoplus_{i=1,3,\dots,2F-1} c[j_{i-1} : j_i] \quad (8)$$

$$e_j = \bigoplus_{i=2,4,\dots,2F} c[j_{i-1} : j_i] \quad (9)$$

where $c[b : e]$ denotes a sub-constituent containing b th to e th words in c , and \bigoplus symbol is the joint of sub-constituents. We also distribute the capacity of the number of nodes into e_j and o_j branches by setting s as the o_j 's share and assigning the rest to e_j .

To find the best alternation points and shares, we iterate over all possible values for them:

$$\mathbf{j}_\tau^*, s_\tau^* = \operatorname{argmax}_{\mathbf{j}, 0 \leq s \leq \tau} [H(o_j, s) + H(e_j, \tau - s)] \quad (10)$$

where $\mathbf{j} = (0, j_1, \dots, j_{2F-1}, |c|)$ must satisfy $j_i \leq j_{i+1} \leq |c|$, $j_1 \neq |c|$, and if $j_i \neq |c|$, then $j_i < j_{i+1}$. We discard the cases that o_j or e_j have a fan-out greater than F , because they are guaranteed not to appear in the output (Theorem 5, whose proof does not rely on this theorem).

Our tree-building search processing first assumes a node is binary and then decides whether to join its parent and children to achieve non-binarity. Therefore, we calculate the best hit count when c itself is excluded, denoted by H_{excl} , or included, denoted by H_{incl} . We have

$$H_{\text{excl}}(c, \tau) = H(o_{\mathbf{j}_\tau^*}, s_\tau^*) + H(e_{\mathbf{j}_\tau^*}, \tau - s_\tau^*) \quad (11)$$

$$H_{\text{incl}}(c, \tau) = H_{\text{excl}}(c, \tau - 1) + h(c) \quad (12)$$

Finally, the recursion is

$$H(c, \tau) = \max\{H_{\text{excl}}(c, \tau), H_{\text{incl}}(c, \tau)\} \quad (13)$$

The best score of $\sum_{k=1}^K F_1(T, T_k)$ for the sentence S is

$$\operatorname{argmax}_{n < \tau < 2n} \frac{H(S, \tau)}{\tau + 2n - 1} \quad (14)$$

¹²If the output is restricted to binary trees, we do not need to track the number of nodes in a subtree, because it is always $2n - 1$ for n leaves.




Figure 6: An example of the recursion step. The yellow hachures indicate the discontinuity in the constituent c .

according to Eqn. (2). Problem 1, as a decision problem, can then be solved trivially by checking whether the score in (14) is greater than or equal to z .

Overall, the time complexity for a recursion step is $\mathcal{O}(n^{2F})$ to find \mathbf{j}_τ^* and s_τ^* . With an $\mathcal{O}(n^{2F+1})$ -sized DP table, the complexity of the entire DP algorithm is $\mathcal{O}(n^{4F+1})$, which is polynomial. Therefore, Problem 1 belongs to **P**. \square

The above proof concerns the decision problem (whether the score reaches or exceeds a threshold). To search for the average constituency tree, we may backtrack the best corresponding constituency subtrees during the recursion.

Notice that restricting output to be binary tree simplifies the algorithm by always setting $s_\tau^* = 2|o_j| - 1$, $H(c, \tau) = H_{\text{incl}}(c, \tau)$, and only considering the computation of $H(c, \tau)$ if $\tau = 2|c| - 1$. It reduces the complexity of a recursion step to $\mathcal{O}(n^{2F-1})$ and that of the DP-table size to $\mathcal{O}(n^{2F})$, resulting in an overall complexity of $\mathcal{O}(n^{4F-1})$.

The soundness of the DP can be cross-validated by our meet-in-the-middle search (§2.3 and §B). As both are exact algorithms, they should output exactly the same results. This is what we observed in our experiments, providing strong empirical evidence that both of our algorithms are sound.

A.2 Proof of Theorem 2

Problem 2 (Averaging trees with bounded fan-out). *Consider a number z and constituency trees T_1, \dots, T_K with the same leaf nodes, where the fan-out is at most F but the trees may be non-binary. Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 2. *Problem 2 belongs to **P**.*

Proof. In §A.1, we provide a DP algorithm and by backtracking we can easily obtain, for any τ , the constituency tree T that maximizes $\sum_{c \in C(T)} h(c)$ such that $|C(T)| = \tau$ (i.e., the tree having τ nodes). The DP algorithm, in fact, works for any scoring function h over constituents. Therefore, we have a polynomial-time solver for the following problem:

Fixed-Size Maximization. *Given a scoring function h over constituents of a length- n sentence and a natural number τ such that $n < \tau < 2n$, what is the tree T that maximizes $\sum_{c \in C(T)} h(c)$ with $|C(T)| = \tau$?*

For Problem 2, we may enumerate all possible values of τ , i.e., $n < \tau < 2n$. Given a fixed τ , we have

$$T^{(\tau)} = \operatorname{argmax}_{T:|T|=\tau} \sum_{k=1}^K F_1(T, T_k) \quad (15)$$

$$= \operatorname{argmax}_{T:|T|=\tau} \sum_{k=1}^K \frac{|C(T) \cap C(T_k)|}{|C(T)| + |C(T_k)|} \quad (16)$$

$$= \operatorname{argmax}_{T:|T|=\tau} \sum_{c \in C(T)} \underbrace{\sum_{k=1}^K \frac{\mathbb{1}[c \in C(T_k)]}{\tau + |C(T_k)|}}_{\hat{h}(c)} \quad (17)$$

$$= \operatorname{argmax}_{T:|T|=\tau} \sum_{c \in C(T)} \hat{h}(c) \quad (18)$$

In Eqn. (18), we define a generalized scoring function \hat{h} , which can be thought of as a weighted hit count. In other words, the overall- F_1 maximization in Eqn. (15) can be reduced to the above fixed-size maximization problem where the number of nodes is given, which can be solved in polynomial time. We may repeatedly solve the problem for $(n - 2)$ -many values of τ and find the best answer among them:

$$\tau^* = \operatorname{argmax}_{n < \tau < 2n} \sum_{k=1}^K F_1(T^{(\tau)}, T_k) \quad (19)$$

$$T^* = T^{(\tau^*)} \quad (20)$$

This does not push the complexity beyond polynomial. Finally, we can answer Problem 2 by checking whether $\sum_{k=1}^K F_1(T^*, T_k) \geq z$. \square

A.3 Proof of Theorem 3

Problem 3 (Averaging trees). *Consider a number z and constituency trees T_1, \dots, T_K with the same leave nodes, where the fan-out is unbounded and the trees may be non-binary. Is there a constituency tree T such that $\sum_{k=1}^K F_1(T, T_k) \geq z$?*

Theorem 3. *Problem 3 belongs to NP-complete.*

Proof. To show a problem is **NP**-complete, we need to first show it is an **NP** problem, i.e., polynomial-time solvable with a non-deterministic Turing machine. Then, we need to show its completeness, i.e., any **NP** problem, or a known **NP**-complete problem, can be reduced to this problem in polynomial time.

[Being **NP**] Being non-deterministic polynomial-time solvable is equivalent to the ability to be verified with a *certificate* in polynomial time (Arora and Barak, 2009). In our case, verifying the score of any candidate tree (which may serve as the certificate) can be done in polynomial time, proving the **NP** part.

[Being complete] To show the completeness, it suffices to reduce a known **NP**-complete problem to the problem at hand. In particular, we would reduce the max clique problem, known to be **NP**-complete (Arora and Barak, 2009), to our problem.

Max Clique. *Consider a number z and a graph G . Is there a clique in G having at least z vertices?*

For reduction, we use a solver for Problem 3 to solve the max clique problem with $G = \langle V, E \rangle$ and z being given.

Without loss of generality, we may assume that G does not contain any vertex that is connected to all other vertices for the rest of this proof. This is because such a vertex is guaranteed to appear in any maximal clique and can be pruned. Then the original problem is equivalent to determining whether the pruned graph has a clique with $(z - a)$ -many vertices, where a is the number of pruned vertices. Notice that such pruning can be accomplished in polynomial time, thus not changing the complexity category.




Figure 7: An example of reduction.

We construct an instantiation of Problem 3 as follows. For every $u_i \in V$, construct a set S_i initialized by a special symbol t_i . For every disconnected vertex pair (u_i, u_j) such that $i < j$, place a special symbol t_i^j in both S_i and S_j , shown in Figure 7. We can view each t_i and t_i^j as words in a sentence. Each S_i represents a possible constituent containing the words corresponding to its elements.

For each S_i , construct a constituency tree T_i by all the single words, the whole-sentence constituent, and S_i . Note that S_i cannot be a single word because there is no vertex in G connected to all other vertices. Moreover, S_i cannot be the whole-sentence constituent either, because S_i does not contain the word $t_{i'}$ for $i' \neq i$. Therefore, $|C(T_i)| = n + 2$ where n is the number of words.

We next show that Problem 3 given $\{T_i\}_{i:u_i \in V}$ and the value $\tilde{z} = \frac{2z+2|V|(n+1)}{z+2n+3}$ is equivalent to the max clique problem given the graph G and the number z .

For every $u_i, u_j \in V$ with $i < j$, u_i and u_j are connected if and only if $S_i \cap S_j = \emptyset$; otherwise t_i^j is in both S_i and S_j . In addition, $S_i \not\subseteq S_j$ for $i \neq j$ because there is a special symbol t_i in each S_i . In other words, S_i and S_j are compatible constituents if and only if u_i and u_j are connected in G . As a result, every constituency tree—constructed by single words, the whole-sentence constituent, and a subset of $\{S_i\}_{i:u_i \in V}$ —corresponds to a clique in G . We refer to such a tree as a *clique tree*.

We can further find the correspondence of the scores between Problem 3 and the max clique problem. For every clique tree T , we have

$$\sum_{i:u_i \in V} F_1(T, T_i) = \sum_{i:u_i \in V} \frac{2(\mathbf{1}[C(T_i) \subseteq C(T)] + n + 1)}{|C(T)| + n + 2} \quad (21)$$

$$= \frac{2(|C(T)| - n - 1) + 2|V|(n + 1)}{|C(T)| + n + 2} \quad (22)$$

where $|C(T)| - n - 1$ is the same as the number of vertices in the corresponding clique. This shows a clique tree T with a score of $\tilde{z} = \frac{2z+2|V|(n+1)}{z+2n+3}$ corresponds to a clique with z vertices.

Now, let us consider using a solver for Problem 3 to solve the max clique problem. This is discussed by cases.

[Case 1] If the answer to Problem 3 is “no,” it means that there does not exist a tree (either a clique tree or a general tree) T satisfying $\sum_{i:u_i \in V} F_1(T, T_i) \geq \tilde{z}$. Thus, there does not exist a clique Q satisfying $|Q| \geq z$.

[Case 2] If the answer to Problem 3 is “yes,” there must exist a tree T satisfying the condition $\sum_{i:u_i \in V} F_1(T, T_i) \geq \tilde{z}$. Notice that such a tree is not guaranteed to be a clique tree, but in this case, we can show there must also exist a clique tree satisfying the condition.

In general, a tree is a clique tree if and only if it does not include any zero-hit constituents (zero hit means not appearing in any T_i). This is because, other than single words and the whole-sentence constituent, a clique tree selects a subset of $\{S_i\}_{i:u_i \in V}$, whereas each T_i only selects one element of this

set. Let $C_0(T)$ denote zero-hit constituents in T . We construct T' by removing $C_0(T)$ from T . We have

$$\tilde{z} \leq \sum_{i:u_i \in V} F_1(T, T_i) \quad (23)$$

$$= \sum_{i:u_i \in V} \frac{2 \sum_{c \in C(T)} \mathbb{1}[c \in C(T_i)]}{|C(T)| + |C(T_i)|} \quad (24)$$

$$\leq \sum_{i:u_i \in V} \frac{2 \sum_{c \in C(T) \setminus C_0(T)} \mathbb{1}[c \in C(T_i)]}{|C(T)| - |C_0(T)| + |C(T_i)|} \quad (25)$$

$$= \sum_{i:u_i \in V} \frac{2 \sum_{c \in C(T')} \mathbb{1}[c \in C(T_i)]}{|C(T')| + |C(T_i)|} \quad (26)$$

$$= \sum_{i:u_i \in V} F_1(T', T_i) \quad (27)$$

Here, Inequality (23) is due to the assumption of Case 2 that the answer to Problem 3 is “yes.” Inequality (25) is because $|C_0(T)| \geq 0$ and $h(c) = 0$ for $c \in C_0(T)$, potentially increasing the denominator while the numerator is the same.

Since T' is a clique tree that satisfies $\sum_{i:u_i \in V} F_1(T', T_i) \geq \tilde{z}$, we find the solution to the max clique problem is also “yes.”

The given reduction can be done in polynomial time, as it iterates over all vertices and vertex pairs. This concludes that Problem 3 belongs to **NP**-complete. \square

A.4 Proof of Lemma 1

Lemma 1. *Given a graph G as described in §2.3, the clique Q^* maximizing $f(Q; \alpha_1, \alpha_2)$ corresponds to a constituency tree, if $\alpha_1 \leq K\alpha_2$ where K is the number of individuals.*

Proof. A selection of constituents corresponds to a constituency tree over a sentence if and only if (1) it includes all the single words and the whole-sentence constituent, referred to as *trivial* constituents, and (2) all of its constituents are compatible with each other (already satisfied by the clique definition).

Trivial constituents are compatible with any possible constituent, and thus their corresponding vertices are connected to all vertices in G . Moreover, the hit count of such a constituent equals K , the number of individual trees, because trivial constituents appear in all constituency trees.

We assume by way of contradiction that a trivial constituent does not appear in the constituency tree corresponding to Q^* . This is equivalent to having a vertex u —connected to every vertex in G with $w(u) = K$, i.e., having a weight of K —not appearing in Q^* ; that is,

$$f(Q^* \cup \{u\}; \alpha_1, \alpha_2) < f(Q^*; \alpha_1, \alpha_2) \quad (28)$$

With the definition of f , we have

$$\frac{\sum_{v \in Q^*} w(v) + w(u) + \alpha_1}{|Q^*| + 1 + \alpha_2} < \frac{\sum_{v \in Q^*} w(v) + \alpha_1}{|Q^*| + \alpha_2} \quad (29)$$

$$\frac{\sum_{v \in Q^*} w(v) + w(u) + \alpha_1}{\sum_{v \in Q^*} w(v) + \alpha_1} < \frac{|Q^*| + 1 + \alpha_2}{|Q^*| + \alpha_2} \quad (30)$$

$$1 + \frac{w(u)}{\sum_{v \in Q^*} w(v) + \alpha_1} < 1 + \frac{1}{|Q^*| + \alpha_2} \quad (31)$$

$$w(u) < \frac{\sum_{v \in Q^*} w(v) + \alpha_1}{|Q^*| + \alpha_2} \quad (32)$$

$$w(u) < \frac{K|Q^*| + K\alpha_2}{|Q^*| + \alpha_2} = K \quad (33)$$

Here, the last inequality is due to $\alpha_1 < K\alpha_2$ and $w(v) \leq K$ for every v . But this inequality contradicts our assumption that $w(u) = K$. \square

A.5 Proof of Theorem 5

Theorem 5 (Lower bound). *For every constituent $c \in C(T^*) \setminus P$, where $P \subseteq C(T^*)$, we have*

$$h(c) > \min_{|P| \leq j \leq 2n-2} \frac{\sum_{i=1}^{j-|P|} \lambda_i^+ + \sum_{c' \in P} h(c')}{j + 2n - 1} \quad (5)$$

where λ_i^+ is the i th smallest positive hit count.

Proof. Consider any constituent $c_0 \in C(T^*) \setminus P$. We construct T' by removing c_0 from T^* , i.e., $C(T^*) = C(T') \cup \{c_0\}$. Based on Eqn. (2), we have

$$\frac{\sum_{c \in C(T^*)} h(c)}{|C(T^*)| + 2n - 1} > \frac{\sum_{c \in C(T')} h(c)}{|C(T')| + 2n - 1} \quad (34)$$

$$\frac{\sum_{c \in C(T^*)} h(c)}{\sum_{c \in C(T')} h(c)} > \frac{|C(T^*)| + 2n - 1}{|C(T')| + 2n - 1} \quad (35)$$

$$1 + \frac{h(c_0)}{\sum_{c \in C(T')} h(c)} > 1 + \frac{1}{|C(T')| + 2n - 1} \quad (36)$$

$$h(c_0) > \frac{\sum_{c \in C(T')} h(c)}{|C(T')| + 2n - 1} \quad (37)$$

Since the right-hand side is non-negative, we must have $h(c_0) > 0$. In other words, zero-hit constituents do not appear in $C(T^*) \setminus P$. Thus, m -many constituents in $C(T^*) \setminus P$ have a total of hit count greater than or equal to $\sum_{i=1}^m \lambda_i^{(+)}$, because $\lambda_i^{(+)}$ is the i th smallest positive hit count.

On the other hand, we have $c_0 \notin P$, implying that $P \subseteq C(T')$. Since a constituent $c \in C(T')$ is either in P or not, we can lower-bound the total hit count of T' by

$$\sum_{c \in C(T')} h(c) \geq \sum_{c' \in P} h(c') + \sum_{i=1}^{|C(T') \setminus P|} \lambda_i^+ \quad (38)$$

Putting (38) into (37), we have

$$h(c_0) > \frac{\sum_{i=1}^{|C(T')| - |P|} \lambda_i^+ + \sum_{c' \in P} h(c')}{|C(T')| + 2n - 1} \quad (39)$$

Moreover, $P \subseteq C(T')$ also implies

$$|P| \leq |C(T')| = |C(T^*)| - 1 \leq 2n - 2 \quad (40)$$

which derives (39) to a T' -independent lower bound for $h(c_0)$ if we consider all the values $|C(T')|$ can take, given by

$$h(c_0) > \min_{|P| \leq j \leq 2n-2} \frac{\sum_{i=1}^{j-|P|} \lambda_i^+ + \sum_{c' \in P} h(c')}{j + 2n - 1} \quad (41)$$

concluding the proof. \square

The lower bound can be applied for pruning the search. We may immediately prune zero-hit constituents by setting $P = \emptyset$, and more importantly, it yields a much tighter lower bound together with Theorem 6.

A.6 Proof of Theorem 6

Theorem 6 (Upper bound). *Let c be a constituent with a hit count of K , where K is the number of individuals. (a) The constituent c is compatible—i.e., may occur in the same constituency tree—with every constituent in the average tree. (b) The constituent c appears in the average tree.*

Proof. [Part (a)] A hit count of K indicates that the constituent appears in all the individuals. Therefore, c is compatible with every constituent in every individual. On the other hand, Theorem 5 shows that every constituent in the average tree has a positive hit count (appears in at least one individual) and thus is compatible with c .

[Part (b)] We assume by way of contradiction that a constituent c_0 with a hit count of K does not appear in the average tree T^* . As shown in (a), c_0 is compatible with every constituent in T^* ; hence, there exists a constituency tree \hat{T} such that $C(\hat{T}) = C(T^*) \cup \{c_0\}$. Based on Eqn. (2), we have

$$\frac{\sum_{c \in C(\hat{T})} h(c)}{|C(\hat{T})| + 2n - 1} < \frac{\sum_{c \in C(T^*)} h(c)}{|C(T^*)| + 2n - 1} \quad (42)$$

$$\frac{\sum_{c \in C(\hat{T})} h(c)}{\sum_{c \in C(T^*)} h(c)} < \frac{|C(\hat{T})| + 2n - 1}{|C(T^*)| + 2n - 1} \quad (43)$$

$$1 + \frac{h(c_0)}{\sum_{c \in C(T^*)} h(c)} < 1 + \frac{1}{|C(T^*)| + 2n - 1} \quad (44)$$

$$h(c_0) < \frac{\sum_{c \in C(T^*)} h(c)}{|C(T^*)| + 2n - 1} \quad (45)$$

$$h(c_0) < \frac{K|C(T^*)|}{|C(T^*)| + 2n - 1} \leq K \quad (46)$$

Here, the last inequality is due to $h(c) \leq K$ for every c . But this contradicts our assumption that $h(c_0) = K$. \square

B Our Meet-in-the-Middle Algorithm

The meet-in-the-middle technique (Horowitz and Sahni, 1974) has been used to tackle the max clique problem. This method halves the exponent of the time complexity of exhaustive search, cutting it down from $\mathcal{O}(2^{|V|})$ to $\mathcal{O}(2^{\lfloor \frac{|V|}{2} \rfloor} |V|^2)$ given a graph with $|V|$ vertices. However, the exact same algorithm is not sufficient for our needs, because Problem 4 is a generalization of the max clique problem and requires additional consideration. In this section, we develop an extended version of the algorithm that solves Problem 4 within the same time complexity.

Problem 4 (Normalized Max Weighted Clique). *Consider a weighted undirected graph $G = \langle V, E \rangle$ and an objective function $f(Q; \alpha_1, \alpha_2) = \frac{\sum_{v \in Q} w(v) + \alpha_1}{|Q| + \alpha_2}$, where $\alpha_1, \alpha_2 \in \mathbb{R}$, $Q \subseteq V$, and $w(v)$ is the weight of v . What is the clique Q that maximizes $f(Q; \alpha_1, \alpha_2)$?*

To find the solution to Problem 4, we utilize the meet-in-the-middle technique, which splits the vertex set V into two parts V_1 and V_2 , whose sizes are as equal as possible.

For every $V'_1 \subseteq V_1$ and every $0 \leq j \leq |V'_1|$, we define

$$\text{best}(V'_1, j) = \underset{\substack{V''_1: V''_1 \subseteq V'_1, |V''_1|=j, \\ V''_1 \text{ is a clique}}}{\text{argmax}} \sum_{v \in V''_1} w(v) \quad (47)$$

which is the best j -vertex clique in V'_1 , and can be computed recursively.

For initialization, we have $\text{best}(V'_1, 0) = \emptyset$ for every $V'_1 \subseteq V_1$ and $\text{best}(\{u\}, 1) = \{u\}$ for every $u \in V_1$.

The recursion is to compute the best set for any multi-vertex V'_1 assuming we have the best set for every subset of V_1 , whose size is smaller than V'_1 . To achieve this, we pick any $u \in V'_1$, and the $\text{best}(V'_1, j)$ may or may not include u , which is discussed by cases.

If $u \notin \text{best}(V'_1, j)$, we have

$$\text{best}(V'_1, j) = \text{best}(V'_1 \setminus \{u\}, j) \quad (48)$$

If otherwise $u \in \text{best}(V'_1, j)$, then the other $j - 1$ vertices must be the neighbors of u in V'_1 , because




Figure 8: Recall scores of 10 most frequent constituent types on LASSY: (a) among all constituents, and (b) among discontinuous constituents. The percentages under the types indicate their occurrences in the dataset. Error bars are the standard deviation for five individuals.

$\text{best}(\cdot, \cdot)$ must be a clique. Denoting by $A_{V_1'}^{(u)}$ all the vertices in V_1' connected to u , we have

$$\text{best}(V_1', j) = \{u\} \cup \text{best}(A_{V_1'}^{(u)}, j - 1) \quad (49)$$

Combining the two cases, we have $\text{best}(V_1', j)$ be either Eqn. (48) or Eqn. (49), depending on which yields a higher weight.

For the V_2 part, we would like to evaluate every clique $V_2' \subseteq V_2$ and its corresponding best match in V_1 to obtain $Q^* \subseteq V = V_1 \cup V_2$ that maximizes our scoring function $f(Q; \alpha_1, \alpha_2)$.

We enumerate all the subsets of V_2 . For every clique $V_2' \subseteq V_2$, we denote by $A_{V_1}^{(V_2')}$ the vertices in V_1 that are connected to every vertex in V_2' . We compute

$$\text{val}(V_2', j) = \sum_{v \in V_2' \cup \text{best}(A_{V_1}^{(V_2')}, j)} w(v) \quad (50)$$

being the weight of the best clique, which in V_1 has j vertices and in V_2 involves V_2' only.

Since the $V = V_1 \cup V_2$, we can find the vertices of Q^* in V_1 and V_2 separately:

$$Q_2^*, j^* = \underset{V_2', j}{\text{argmax}} \frac{\text{val}(V_2', j) + \alpha_1}{|V_2'| + j + \alpha_2} \quad (51)$$

$$Q_1^* = \text{best}(A_{V_1}^{(Q_2^*)}, j^*) \quad (52)$$

The final answer is $Q^* = Q_1^* \cup Q_2^*$.

We analyze the time complexity of our algorithm. For the V_1 part, our algorithm enumerates $\mathcal{O}(2^{\frac{|V_1|}{2}})$ -many subsets V_1' and $\mathcal{O}(|V|)$ -many j values; in each iteration, $\mathcal{O}(|V|)$ -many vertices are visited for $A_{V_1'}^{(u)}$. For the V_2 part, we enumerate $\mathcal{O}(2^{\frac{|V_2|}{2}})$ -many subsets of V_2 , each verified for being a clique in $\mathcal{O}(|V|^2)$ time; after that, $A_{V_1}^{(V_2')}$ is retrieved with complexity $\mathcal{O}(|V_1| \cdot |V_2'|) = \mathcal{O}(|V|^2)$, which can be absorbed in the \mathcal{O} -notation. Finally, retrieving Q^* happens in $\mathcal{O}(2^{\frac{|V_1|}{2}} |V|)$ time, and the overall time complexity of our algorithm is $\mathcal{O}(2^{\frac{|V_1|}{2}} |V|^2)$.

C Performance by Constituency Types

We would like to analyze our ensemble method's performance from a linguistic point of view. We provide a performance breakdown by constituency types, e.g., noun phrases (NP) and verb phrases (VP). Note that our unsupervised models offer unlabeled constituency structures, and thus we can only compute recall scores for different constituency types.

In Figures 8a and 8b, we report the recall scores for all constituents and discontinuous constituents, respectively. We can see for most of the types, the ensemble can retain the best individual's performance in that type, suggesting that the ensemble is able to leverage the diverse strengths of different individuals.


D Case Studies

In Figure 9, we present a case study on the LASSY dataset to show how the ensemble can take advantage of different individuals. In the figure, a groundtruth constituent is annotated by its type (e.g., CONJ representing conjunction). However, our unsupervised parsing is untyped, and we use “o” to denote a constituent. A skipped line “-|-” indicates discontinuity.


Consider the last three words “*moeten worden opgesteld*” (must be established).¹³ We see the ensemble output follows the majority votes (three out of five, namely Individuals 2, 4, and 5) and detects the correct structure of this phrase. Moreover, the ensemble is able to detect the discontinuous constituent “*in sommige gevallen ... niet toereikend zijn*” (in some cases ... not sufficient be), suggested by Individuals 4 and 5. On the other hand, we also notice that every individual produces incorrect discontinuous constituents, but our ensemble is able to smooth out such noise and achieve high F_1 scores.

¹³We provide English interpretations by ChatGPT-4 with the prompt: “Translate the following Dutch sentence. Give word-by-word literal translation.”

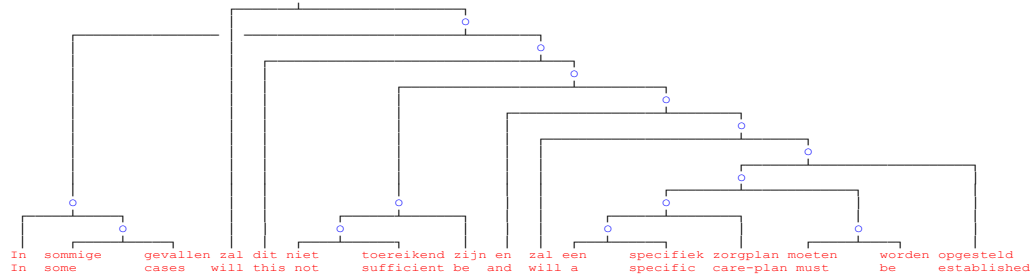
Ground truth:



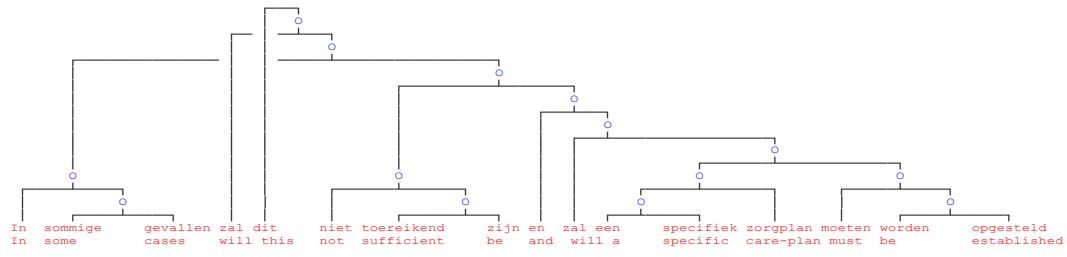
Our ensemble: $F_1^{\text{overall}} = 70.0\%$, $F_1^{\text{cont}} = 66.7\%$, $F_1^{\text{disco}} = 100\%$




Individual 1: $F_1^{\text{overall}} = 63.6\%$, $F_1^{\text{cont}} = 66.7\%$, $F_1^{\text{disco}} = 50.0\%$



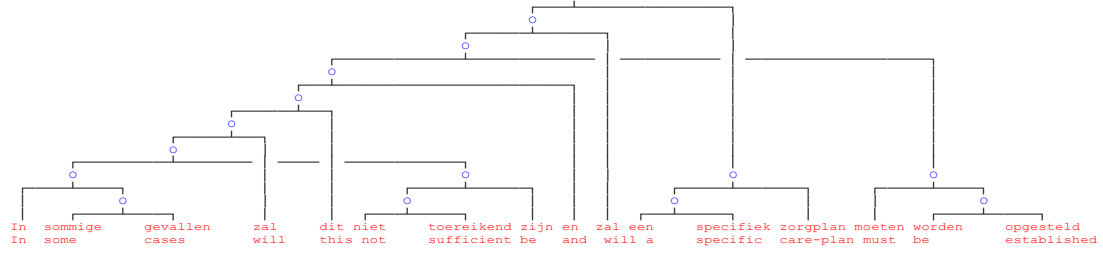
Individual 2: $F_1^{\text{overall}} = 54.6\%$, $F_1^{\text{cont}} = 63.2\%$, $F_1^{\text{disco}} = 00.0\%$



Individual 3: $F_1^{\text{overall}} = 27.3\%$, $F_1^{\text{cont}} = 31.6\%$, $F_1^{\text{disco}} = 00.0\%$



Individual 4: $F_1^{\text{overall}} = 63.6\%$, $F_1^{\text{cont}} = 70.6\%$, $F_1^{\text{disco}} = 40.0\%$



Individual 5: $F_1^{\text{overall}} = 36.4\%$, $F_1^{\text{cont}} = 40.0\%$, $F_1^{\text{disco}} = 00.0\%$




Figure 9: Case studies with an example in Dutch from the LASSY dataset.
15156