

Efficient Sampling of Dependency Structures

Ran Zmigrod  Tim Vieira  Ryan Cotterell  
 University of Cambridge  Johns Hopkins University  ETH Zürich
rz279@cam.ac.uk tim.f.vieira@gmail.com
ryan.cotterell@inf.ethz.ch

Abstract

Probabilistic distributions over spanning trees in directed graphs are a fundamental model of dependency structure in natural language processing, syntactic dependency trees. In NLP, dependency trees often have an additional root constraint: only one edge may emanate from the root. However, no sampling algorithm has been presented in the literature to account for this additional constraint. In this paper, we adapt two spanning tree sampling algorithms to faithfully sample dependency trees from a graph subject to the root constraint. Wilson (1996)’s sampling algorithm has a running time of $\mathcal{O}(H)$ where H is the mean hitting time of the graph. Colbourn et al. (1996)’s sampling algorithm has a running time of $\mathcal{O}(N^3)$, which is often greater than the mean hitting time of a directed graph. Additionally, we build upon Colbourn’s algorithm and present a novel extension that can sample K trees without replacement in $\mathcal{O}(KN^3 + K^2N)$ time. To the best of our knowledge, no algorithm has been given for sampling spanning trees without replacement from a directed graph.¹

1 Introduction

Spanning trees in directed graphs² are fundamental combinatorial structures in natural language processing where they are used to represent dependency structures—especially syntactic dependency structure (Kübler et al., 2009). Additionally, probabilistic models over spanning trees are common in the NLP literature with applications primarily in non-projective dependency parsing (Pei et al., 2015; Wang and Chang, 2016; Dozat and Manning, 2017; Ma and Hovy, 2017), but also in recovering

¹Our implementation of these algorithms is publicly available at <https://github.com/rycolab/treesample>.

²Directed spanning trees are known as *arborescences* in the graph theory literature (Williamson, 1985). However, we will simply refer to them as spanning trees.

phylogenetic structures (Andrews et al., 2012), and event extraction (McClosky et al., 2011).

Given the prevalence of such probabilistic models, efficient dependency tree sampling algorithms deserve study. Indeed, some work has been done in transition-based dependency parsing (Keith et al., 2018) as well as graph-based dependency parsing (Nakagawa, 2007; Mareček and Žabokrtský, 2011). Sampling has also been utilized in an abundance of NLP tasks, such as text generation (Clark et al., 2018; Fedus et al., 2018), co-reference resolution (Singh et al., 2012), and language modeling (Mnih and Hinton, 2007; Logan IV et al., 2020).

The theoretical computer science literature has several efficient algorithms for sampling directed spanning trees. These algorithms come in two flavors. First, random walks through Markov chains have been used to sample spanning trees from both undirected (Broder, 1989; Aldous, 1990) and directed graphs (Wilson, 1996). The algorithm of Wilson (1996) is linear in the mean hitting time of the graph and is currently the fastest sampling algorithm for directed spanning trees. It has been used in dependency parsing inference by Zhang et al. (2014a,b). Second, several algorithms have leveraged the matrix–tree theorem (MTT; Kirchhoff, 1847; Tutte, 1984). The MTT has been frequently used to perform inference on non-projective graph-based dependency parsers (Koo et al., 2007; McDonald and Satta, 2007; Smith and Smith, 2007; Zmigrod et al., 2021). This theorem was first used for sampling by Guénoche (1983) who gave an $\mathcal{O}(N^5)$ algorithm which was then improved by Kulkarni (1990) and Colbourn et al. (1996). Colbourn et al. (1996) give an $\mathcal{O}(N^3)$ algorithm to sample spanning trees from an *unweighted* directed graph. We generalize their algorithm to the weighted case.

While directed spanning tree sampling algorithms exist, an important constraint of many dependency tree schemes, such as the Universal De-

pendency (UD) scheme (Nivre et al., 2018), is that a dependency tree may only have one edge emanating from the designated root symbol. Algorithms exist for enforcing this constraint in decoding (Gabow and Tarjan, 1984; Zmigrod et al., 2020; Stanojević and Cohen, 2021) and inference (Koo et al., 2007; Zmigrod et al., 2021). However, to the best of our knowledge, no sampling algorithm exists which enforces the root constraint.

In this paper, we adapt the algorithms of Wilson (1996) and Colbourn et al. (1996)³ to efficiently sample directed spanning trees subject to a root constraint while maintaining the runtime of the original algorithms. Additionally, we provide a further extension to Colbourn et al. (1996)’s algorithm that allows us to sample trees without replacement. Sampling without replacement (SWOR) algorithms are useful when distributions are skewed, which is often the case in a trained system. To the best of our knowledge, no SWOR algorithm has been presented in the literature for directed spanning trees, though Shi et al. (2020) provides a general framework that enables a SWOR algorithm to be adapted for particular kinds of sampling algorithms.

2 Distributions over Trees

We consider distributions over spanning trees in **rooted directed weighted graphs** (graphs for short). A graph is denoted by $G = (\rho, \mathcal{N}, \mathcal{E})$ where \mathcal{N} is a set of $N + 1$ nodes including a designated root node ρ and \mathcal{E} is a set of ordered pairs between two nodes $(i \rightarrow j)$. We note that the non-root nodes $\mathcal{N} \setminus \{\rho\}$ can be enumerated as $[1, \dots, N]$. For each edge $(i \rightarrow j) \in \mathcal{E}$, we associate a non-negative weight $w_{i \rightarrow j} \in \mathbb{R}_{\geq 0}$. Note that $w_{i \rightarrow j} = 0$ when $(i \rightarrow j) \notin \mathcal{E}$.

A **directed spanning tree**, denoted by t , is a collection of edges in a graph G such that each node $j \in \mathcal{N} \setminus \{\rho\}$ has exactly one incoming edge and t contains no cycles. Moreover, we specifically examine distributions over **dependency trees**, which are spanning trees with an additional constraint that exactly one edge must emanate from the root ρ . This additional constraint is common amongst most dependency tree annotation schemes, e.g.,

³Note that Colbourn et al. (1996) presents two algorithms for sampling directed spanning trees. In this work, we focus on their first algorithm, which runs in $\mathcal{O}(N^3)$. While the second algorithm is based on a reduction to fast matrix multiplication, which is typically impractical. We did not extend this algorithm because it is not amendable to sampling without replacement and it is generally slower than the algorithm of Wilson (1996).

Universal Dependencies.⁴ When the type of tree (spanning or dependency) is clear from context, we will simply use **trees**. The set of all spanning trees of a graph G is given by $\mathcal{T}(G)$ and the set of all dependency trees of a graph G is given by $\mathcal{D}(G)$. Note that $\mathcal{D}(G) \subseteq \mathcal{T}(G)$. When the graph G is clear from context, we will refer to these sets \mathcal{T} and \mathcal{D} respectively.

The **weight** of a tree t is the product of its edge weights:

$$w(t) \stackrel{\text{def}}{=} \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \quad (1)$$

The **probability** of a spanning tree is then given by

$$p(t) \stackrel{\text{def}}{=} \frac{w(t)}{Z} \quad \text{where} \quad Z \stackrel{\text{def}}{=} \sum_{t \in \mathcal{T}} w(t) \quad (2)$$

The normalizing constant in the case of dependency trees sums over \mathcal{D} instead of \mathcal{T} .

3 Random Walk Sampling

In this section, we present the spanning tree sampling algorithm of Wilson (1996) and adapt it to sample dependency trees. The algorithm is based on a random walk through the nodes of the graph until a tree is formed. In order to do this, we require a graph G to be **stochastic**. A stochastic graph is such that the weights of all incoming edges to a node sum to one. That is, for all non-root nodes $j \in \mathcal{N} \setminus \{\rho\}$ we have that

$$\sum_{i \in \mathcal{N}} w_{i \rightarrow j} = 1 \quad (3)$$

Wilson (1996) shows that any graph may be converted into a stochastic graph, by adjusting the edge weights to be locally normalized along all incoming edges. Therefore, the edge $(i \rightarrow j)$ has a weight $w'_{i \rightarrow j}$ defined by

$$w'_{i \rightarrow j} \stackrel{\text{def}}{=} \frac{w_{i \rightarrow j}}{\sum_{i' \in \mathcal{N}} w_{i' \rightarrow j}} \quad (4)$$

We can also define this new weight over a tree $w'(t) \stackrel{\text{def}}{=} \prod_{(i \rightarrow j) \in t} w'_{i \rightarrow j}$.

Lemma 1. For any tree $t \in \mathcal{T}$,

$$w'(t) \propto w(t) \quad (5)$$

⁴However, we note that there are exceptions that do not require the root constraint, such as the Prague Treebank (Bejček et al., 2013).

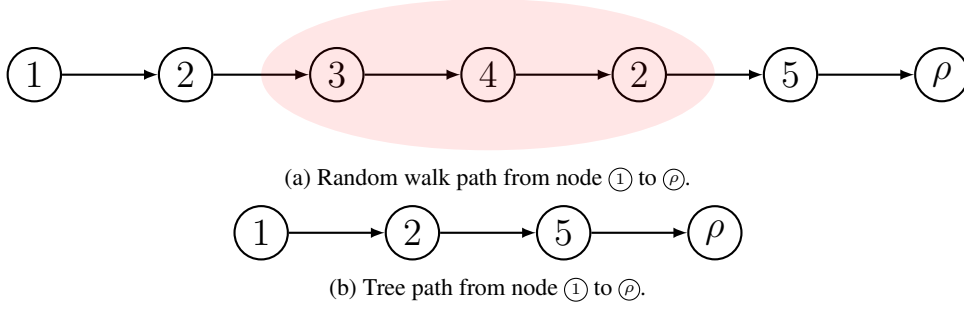


Figure 1: Cycle erasure in a random walk of a graph. The associated graph of the above walk has six nodes (including the root node ρ) and we start the walk at $\textcircled{1}$ with the tree only containing the root node. The random walk includes a cycle with nodes $\textcircled{2}$, $\textcircled{3}$, and $\textcircled{4}$, this cycle is erased when we create the path from $\textcircled{1}$ back up to ρ . Note that the arrows here mark the path upwards rather than the edges in the tree, the tree edges are reversed (e.g. $\rho \rightarrow \textcircled{5}$).

Proof.

$$w'(t) = \prod_{(i \rightarrow j) \in t} w'_{i \rightarrow j} \quad (6a)$$

$$= \prod_{(i \rightarrow j) \in t} \frac{w_{i \rightarrow j}}{\sum_{i' \in \mathcal{N}} w_{i' \rightarrow j}} \quad (6b)$$

$$= \left(\prod_{(i \rightarrow j) \in t} \frac{1}{\sum_{i' \in \mathcal{N}} w_{i' \rightarrow j}} \right) \left(\prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \right)$$

$$= \underbrace{\left(\prod_{(i \rightarrow j) \in t} \frac{1}{\sum_{i' \in \mathcal{N}} w_{i' \rightarrow j}} \right)}_{\text{constant}} w(t) \quad (6c)$$

$$\propto w(t) \quad (6d)$$

Note that the left-hand term in (6c) is a constant since every non-root node has in incoming edge in the tree, and so the constant is equal to $\prod_{i \in \mathcal{N} \setminus \{\rho\}} \left(\sum_{i' \in \mathcal{N}} w_{i' \rightarrow j} \right)^{-1}$. ■

A stochastic graph then defines a **Markov chain** that we can perform a random walk on. We begin with a tree that is just composed of the root ρ . Then while there exists a node connect to the tree, we start a random walk where we add each edge used to the tree, until we encounter a node that is connected to the tree. When we have connected to the tree, we can proceed to start a new random walk from a node not in the tree.

Of course, during a random walk, we may go through a cycle. Whenever we walk through a cycle, we simply forget the cycle as part of the walk. That is, if we find ourselves visiting node i in our walk for a second time, we erase the cyclic path formed at i and continue our walk. This can be seen visually in Fig. 1. This type of walk is known as a loop-erased random walk (Lawler, 1979). Given

1. **def wilson(G):**

2. \triangleright Sample a spanning tree from a graph G ; requires $\mathcal{O}(H)$ time, $\mathcal{O}(N^2)$ space.

3. $t \leftarrow \mathbf{0}$

4. $\text{visited} \leftarrow \{\rho\}$

5. \triangleright The following for and while loops take $\mathcal{O}(H)$ to execute. More specifically, line 9 is called $\mathcal{O}(H)$ times.

6. **for** $i \in \mathcal{N} \setminus \{\rho\}$:

7. $u \leftarrow i$

8. **while** $u \notin \text{visited}$:

9. Sample $v \in \mathcal{N}$ with weight $w'_{v \rightarrow u}$

10. $t_u \leftarrow v$

11. $u \leftarrow v$

12. $u \leftarrow i$

13. **while** $u \notin \text{visited}$:

14. $\text{visited.add}(u)$

15. $u \leftarrow v$ such that $t_u = v$

16. **return** t

Figure 2: Wilson (1996)'s algorithm to sample spanning trees from a graph.

we sample a sequence of edges \mathcal{S} in a random walk, we can split the We can split the edges into those corresponding to the loop-erased random walk, $\widehat{\mathcal{S}}$, and those that do not, $\overline{\mathcal{S}}$.

Lemma 2. For any sequence of edges \mathcal{S} sampled from a random walk, we have

$$p(\mathcal{S}) = p(\widehat{\mathcal{S}}) p(\overline{\mathcal{S}}) \quad (7)$$

Proof.

$$p(\mathcal{S}) = \prod_{(i \rightarrow j) \in \mathcal{S}} w'_{i \rightarrow j} \quad (8a)$$

$$= \left(\prod_{(i \rightarrow j) \in \widehat{\mathcal{S}}} w'_{i \rightarrow j} \right) \left(\prod_{(i \rightarrow j) \in \overline{\mathcal{S}}} w'_{i \rightarrow j} \right)$$

$$= p(\widehat{\mathcal{S}}) p(\overline{\mathcal{S}}) \quad (8b)$$

■

Pseudocode for [Wilson \(1996\)](#)'s algorithm is given in Fig. 2. Examining the pseudocode, one can see it is possible to infinitely encounter cycles. Fortunately, [Wilson \(1996\)](#) proves that this is not the case and that the algorithm has a probabilistic bound of $\mathcal{O}(H)$ where H is the mean hitting time of the graph. The **mean hitting time** of a Markov chain is defined as

$$H \stackrel{\text{def}}{=} \sum_{i,j} \pi_i \pi_j h(i,j) \quad (9)$$

where π is the stationary distribution of the Markov chain, and $h(i,j)$ is the expected number of steps to reach node j starting at node i . [Wilson \(1996\)](#) and [Broder \(1989\)](#) demonstrate that the mean hitting time for directed graphs is usually small (sometimes as low as linear in N). We will compare the empirical runtime of the algorithm against an $\mathcal{O}(N^3)$ algorithm in §4.2.

Theorem 1. *For any graph G , $\text{wilson}(G)$ samples a directed spanning tree $t \in \mathcal{T}$ with probability*

$$p(t) \propto \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \quad (10)$$

Furthermore, wilson runs in $\mathcal{O}(H)$ time.

Proof. It is clear that if $\text{wilson}(G)$ terminates, t will contain a directed spanning tree. Let \mathcal{S} be the set of edges sampled in line 9 of wilson . As we sample t using several independent random walks, by Lemma 2, we sample the edges in t independently of the edges in $\mathcal{S} \setminus t$. In particular, we sample t with probability

$$p(t) = \prod_{(i \rightarrow j) \in t} w'_{i \rightarrow j} \quad (11)$$

By Lemma 1, t is a sample from our desired distribution over trees.

[Wilson \(1996, Theorem 3\)](#) proves that line 9 of wilson executes $\mathcal{O}(H)$ times before the program terminates with tree t . ■

3.1 Root Constraint Modification

[Wilson \(1996\)](#)'s algorithm does not ensure that only one edge emanates from the root. However, a simple modification to the algorithm allows us to do this. The original algorithm samples a spanning

```

1. def wilson_rc( $G$ ):
2.   ▷ Sample a dependency tree from a graph  $G$ ;
   requires  $\mathcal{O}(H)$  time,  $\mathcal{O}(N^2)$  space.
3.   Sample  $j$  with weight  $w_{\rho \rightarrow j}$ 
4.    $t \leftarrow \mathbf{0}$ 
5.    $t_j \leftarrow \rho$ 
6.   visited  $\leftarrow \{j\}$ 
7.   ▷ The following for and while loops take  $\mathcal{O}(H)$ 
   to execute. More specifically, line 11 is called
    $\mathcal{O}(H)$  times.
8.   for  $i \in \mathcal{N} \setminus \{\rho\}$  :
9.      $u \leftarrow i$ 
10.    while  $u \notin$  visited :
11.      Sample  $v \in \mathcal{N} \setminus \{\rho\}$  with weight  $w'_{v \rightarrow u}$ 
12.       $t_u \leftarrow v$ 
13.       $u \leftarrow v$ 
14.     $u \leftarrow i$ 
15.    while  $u \notin$  visited :
16.      visited.add( $u$ )
17.       $u \leftarrow v$  such that  $t_u = v$ 
18.  return  $t$ 

```

Figure 3: Modification of [Wilson \(1996\)](#)'s algorithm to sample dependency trees from a graph. The lines that differ to wilson are highlighted.

tree rooted at ρ for a graph G . Suppose we know we want the edge $(\rho \rightarrow j)$ to be the single edge emanating from the root. Then we can run Wilson 's algorithm on the graph G' , which is G with node ρ removed and node j defined as the new root. By adding $(\rho \rightarrow j)$ to the newly sampled tree, we clearly have a dependency tree in G . We can sample the root edge $(\rho \rightarrow j)$ from the root weights ρ to sample an unbiased dependency tree from the distribution. The pseudocode for this algorithm is given as wilson_rc in Fig. 3

Theorem 2. *For any graph G , $\text{wilson_rc}(G)$ samples a dependency tree $t \in \mathcal{D}$ with probability*

$$p(t) \propto \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \quad (12)$$

Furthermore, wilson_rc runs in $\mathcal{O}(H)$ in.

Proof. We note that executing lines line 6 to line 18 is equivalent to running wilson on the graph G' that is rooted at j (and does not have ρ). By Theorem 1, this results in a tree $t' \in \mathcal{T}(G')$ and was sampled with probability proportional to $\prod_{(i \rightarrow j) \in t'} w_{i \rightarrow j}$. We sample the edge emanating from the root $(\rho \rightarrow j)$ with probability proportional to $w_{\rho \rightarrow j}$ and so $t = t' \cup \{\rho \rightarrow j\}$ is a dependency

tree sampled with probability

$$p(t) \propto \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j}$$

Furthermore, as line 3 of `wilson_rc` takes $\mathcal{O}(N)$ which is less than $\mathcal{O}(H)$, by Theorem 1, `wilson_rc` has a runtime of $\mathcal{O}(H)$. ■

4 Ancestral Sampling

In this section, we present an extension to the ancestral sampling algorithm of Colbourn et al. (1996) to the weighted graph case. This algorithm relies on the efficient computation of Z using the MTT (Kirchhoff, 1847; Tutte, 1984), allows us to compute Z in $\mathcal{O}(N^3)$ by taking the determinant of the Laplacian matrix, $\mathbf{L} \in \mathbb{R}^{N \times N}$. We use Koo et al. (2007)'s adaptation of the MTT to dependency trees.⁵

Theorem 3 (Proposition 1, Koo et al. (2007)). *For any graph G , the normalization constant Z over the distribution of dependency trees \mathcal{D} is given by $Z = |\mathbf{L}|$ where*

$$L_{ij} = \begin{cases} w_{\rho \rightarrow j} & \text{if } i = 1 \\ \sum_{i' \in \mathcal{N} \setminus \{\rho, i\}} w_{i' \rightarrow j} & \text{if } i = j \\ -w_{i \rightarrow j} & \text{otherwise} \end{cases} \quad (13)$$

We present the algorithm using the above Laplacian matrix to sample dependency trees rather than spanning trees. However, one can easily modify this algorithm to sample spanning trees.⁶ The premise of the algorithm is that we iteratively sample an incoming edge to a non-root node of the graph until we have a tree. Without loss of generality, we can enumerate the edges of any sampled tree t as e_1 to e_N . Therefore, at time step n of our sampling algorithm, we will have a subset of our tree

$$t_{<n} = [e_1, \dots, e_{n-1}] \quad (14)$$

⁵Koo et al. (2007)'s adaptation constructs the Laplacian matrix without considering edges emanating from the root. They then arbitrarily replace the first row of the Laplacian matrix with the root edge weights. One can see that the first row is chosen for convenience by examining the proof of Proposition 1 in Koo et al. (2007). Indeed, the desired Laplacian can be obtained by replacing any row by the root edge weights.

⁶The Laplacian matrix for spanning trees is given by

$$L_{ij} = \begin{cases} \sum_{i' \in \mathcal{N} \setminus \{i\}} w_{i' \rightarrow j} & \text{if } i = j \\ -w_{i \rightarrow j} & \text{otherwise} \end{cases}$$

Note that $t_{<1} = \emptyset$ and $t_{<N+1} = t$. We can then express the probability of a tree as

$$p(t) = \prod_{n=1}^N p(e_n \mid t_{<n}) \quad (15)$$

We sample the first edge e_1 with probability $p(e_1)$. We can find its marginal probability by taking the derivative of the log partition function $\log Z$ in Theorem 3 which Koo et al. (2007) show to be⁷

$$p(i \rightarrow j) = \begin{cases} w_{\rho \rightarrow j} \mathbf{B}_{1j} & \text{if } i = \rho \\ w_{i \rightarrow j} (\delta_{j \neq 1} \mathbf{B}_{jj} - \delta_{i \neq 1} \mathbf{B}_{ij}) & \text{otherwise} \end{cases} \quad (16)$$

where $\mathbf{B} = \mathbf{L}^{-\top}$ (the transpose of the inverse of \mathbf{L}) and $\delta_x = 1 \iff x$ is true, otherwise, $\delta_x = 0$. Therefore, after computing \mathbf{B} once, we can compute each $p(i \rightarrow j)$ in $\mathcal{O}(1)$ time. Finding \mathbf{B} requires us to take a matrix inverse, and so runs in $\mathcal{O}(N^3)$ time.⁸

Each subsequent edge that we sample must be conditioned by all previously sampled edges ($t_{<n}$). At the n^{th} step, we have sampled $t_{<n}$, and so our final sampled tree t will be such that $t_{<n} \subseteq t$. Therefore, sampling e_n from G is equivalent to sampling e_n from the subgraph $G \succ t_{<n}$, which is defined as the largest graph such that $t \in \mathcal{T}(G \succ t_{<n}) \implies t_{<n} \subseteq t$. Consequently, if $(i \rightarrow j) \in t_{<n}$, then $G \succ t_{<n}$ does not contain any other incoming edges to node j other than $(i \rightarrow j)$.⁹

A correct ancestral sampling algorithm will sample an edge e for each non-root node from the graph using (16), and then update graph to be $G \succ e$ and repeat. This algorithm will have to recompute \mathbf{B} $\mathcal{O}(N)$ times and so will have a runtime of $\mathcal{O}(N^4)$. We show a graphical example of the algorithm in Fig. 4. We give pseudocode for this as `colbourn` in Fig. 5. The function `sample_edge` samples from the distribution defined in (16) and the function condition updates \mathbf{B} to contain the transpose of the Laplacian inverse of the conditioned graph. We

⁷For spanning trees, the marginal probability can be similarly derived as $p(i \rightarrow j) = w_{i \rightarrow j} (\mathbf{B}_{jj} - \delta_{i \neq \rho} \mathbf{B}_{ij})$ where \mathbf{B} is the transpose of the inverse of the Laplacian matrix in footnote 6.

⁸This runtime is also true by automatic differentiation (Griewank and Walther, 2008) as finding Z takes $\mathcal{O}(N^3)$ time.

⁹When sampling dependency trees, one would think that if $(\rho \rightarrow j) \in t_{<n}$, we would need to remove all outgoing edges from the root. However, by the construction of our Laplacian, Z only accounts for dependency trees and so the marginals already enforce this restriction. Therefore, we only need to remove all other incoming edges to j .




Figure 4: Consider sampling a tree from the fully connected graph G given in (a). We do this by sampling an incoming edge to each non-root node. We first sample an incoming edge to $\textcircled{1}$, the possible edges are dashed in (a). Suppose we sample e_1 with probability $p(e_1)$, then we have $t_{<2} = \{e_1\}$. If we include e_1 in our graph as in (b), and repeat the process, we will sample edge e_2 with probability $p(e_2 | t_{<2})$. We now have $t_{<3} = \{e_1, e_2\}$, and we can sample an incoming edge e_3 to $\textcircled{3}$ with probability $p(e_3 | t_{<3})$ as in (c). We can similarly find $t_{<4}$ in d. Finally, in (e), we have $t_{<5} = t = \{e_1, e_2, e_3, e_4\}$, which is a tree in $\mathcal{T}(G)$. Note $\mathcal{T}(G \succ t) = \{t\}$.

describe an efficient procedure for this conditioning step in the following section.

4.1 Efficiently Computing Marginals

Colbourn et al. (1996) show that we can update the marginals in $\mathcal{O}(N^2)$ rather than $\mathcal{O}(N^3)$ by using rank-one updates on \mathbf{L} . Namely, Colbourn et al. (1996)'s adds an outer-product $\mathbf{u}\mathbf{v}^\top$ to \mathbf{L} for some column vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$ for each conditioning operation. We extend this to the weighted Laplacian for dependency trees.

Lemma 3. For any graph G with Laplacian \mathbf{L} and any edge $e = (i \rightarrow j) \in \mathcal{E}$, the j^{th} column of the Laplacian \mathbf{L}^e of $G \succ e$ is given by

$$\mathbf{L}^e_{[:,j]} = \begin{cases} w_{\rho \rightarrow j} \vec{\mathbf{1}}_1 & \text{if } i = \rho \\ w_{i \rightarrow j} (\delta_{j \neq 1} \vec{\mathbf{1}}_j - \delta_{i \neq 1} \vec{\mathbf{1}}_i) & \text{otherwise} \end{cases} \quad (17)$$

where $\vec{\mathbf{1}}_j$ is the one-hot vector such that the j^{th} element is 1. Furthermore, the k^{th} column \mathbf{L}^e , where $k \neq j$, is equivalent to the k^{th} column of \mathbf{L} .

Proof. Consider the column $\mathbf{L}^e_{[:,k]}$.

Case $k = j$: Then the only incoming edge to j in $G \succ e$ is $(i \rightarrow j)$.

Case $i = \rho$: Then element $\mathbf{L}^e_{1j} = w_{\rho \rightarrow j}$ by Koo et al. (2007)'s construction. As there are no other incoming edges to kj , the remainder of the column is filled with zeros and so $\mathbf{L}^e_{[:,j]} = w_{\rho \rightarrow j} \vec{\mathbf{1}}_1$.

Case $i \neq \rho$: Then $\mathbf{L}^e_{ij} = -w_{i \rightarrow j}$ as long as $i \neq 0$. Since $i \neq \rho$, we also have that $\mathbf{L}^e_{jj} = w_{i \rightarrow j}$ as long as $j \neq 0$. Therefore, we can represent the j^{th} column by $w_{i \rightarrow j} (\delta_{j \neq 1} \vec{\mathbf{1}}_j - \delta_{i \neq 1} \vec{\mathbf{1}}_i)$.

Case $k \neq j$: Then all incoming edges to node k are still in $G \succ e$ and so $\mathbf{L}^e_{[:,k]} = \mathbf{L}_{[:,k]}$. ■

Lemma 3 shows that conditioning by an edge $(i \rightarrow j)$ is equivalent to a column replacement for \mathbf{L} . A column replacement is equivalent to a rank-one update where $\mathbf{v} = \vec{\mathbf{1}}_j$ and $\mathbf{u} = w_{i \rightarrow j} (\vec{\mathbf{1}}_j - \vec{\mathbf{1}}_i) - \mathbf{L}_{[:,j]}$ and $\mathbf{L}_{[:,j]}$ is the j^{th} column of \mathbf{L} . Performing such a rank-one update speeds up the conditioning of \mathbf{L} from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$. More importantly, it lets us update \mathbf{B} in $\mathcal{O}(N^2)$ using the Sherman–Morrison formula (Sherman and Morrison, 1950), which states that for any matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and column vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^\top)^{-1} = \mathbf{A}^{-1} - \mathbf{T} \quad (18)$$

where

$$\mathbf{T} = \frac{\mathbf{A}^{-1} \mathbf{u} \mathbf{v}^\top \mathbf{A}^{-1}}{1 + \mathbf{v}^\top \mathbf{A}^{-1} \mathbf{u}} \quad (19)$$

Recalling that \mathbf{B} requires the inverse transpose, and our choice of $\mathbf{v} = \vec{\mathbf{1}}_j$, we can simplify the expression for \mathbf{T}^\top to be

$$\mathbf{T}^\top = \left(\frac{\mathbf{L}^{-1} \mathbf{u} \vec{\mathbf{1}}_j^\top \mathbf{L}^{-1}}{1 + \vec{\mathbf{1}}_j^\top \mathbf{L}^{-1} \mathbf{u}} \right)^\top \quad (20a)$$

$$= \left(\frac{\mathbf{L}^{-1} \mathbf{u} \mathbf{L}^{-1}_{[j,:]}^\top}{1 + \mathbf{L}^{-1}_{[j,:]} \mathbf{u}} \right)^\top \quad (20b)$$

$$= \frac{\mathbf{L}^{-\top}_{[j,:]} \mathbf{u}^\top \mathbf{L}^{-\top}}{1 + \mathbf{u}^\top \mathbf{L}^{-\top}_{[j,:]}} \quad (20c)$$

$$= \frac{\mathbf{B}_{[j,:]} \mathbf{u}^\top \mathbf{B}}{1 + \mathbf{u}^\top \mathbf{B}_{[j,:]}} \quad (20d)$$

Therefore, we can update \mathbf{B} in $\mathcal{O}(N^2)$. We give pseudocode for this efficient update as condition in Fig. 5.

¹⁰The Sherman–Morrison formula can be computed in $\mathcal{O}(N^2)$ due to the associativity of matrix multiplication.

Global Variables: \mathbf{L} , \mathbf{B}

```

1. def colbourn( $G$ ):
2.    $\triangleright$  Sample a dependency tree from a graph  $G$ ;
      requires  $\mathcal{O}(N^3)$  time,  $\mathcal{O}(N^2)$  space.
3.    $\mathbf{L} \leftarrow \text{Laplacian}(G)$ 
4.    $\mathbf{B} \leftarrow \mathbf{L}^{-\top}$   $\triangleright \mathcal{O}(N^3)$ 
5.    $t \leftarrow []$ 
6.   for  $j \in \mathcal{N} \setminus \{\rho\}$ :
7.      $e \leftarrow \text{sample\_edge}(j)$ 
8.      $t.append(e)$ 
9.     condition( $e$ )
10.  return  $t$ 
11. def Laplacian( $G$ ):
12.   $\triangleright$  Construct the Laplacian of Koo et al. (2007)
      for dependency trees as in (13); requires
       $\mathcal{O}(N^2)$  time,  $\mathcal{O}(N^2)$  space.
13.   $\mathbf{L} \leftarrow \mathbf{0}$ 
14.  for  $j \in \mathcal{N} \setminus \{\rho\}$ :
15.    for  $i \in \mathcal{N} \setminus \{\rho, j\}$ :
16.       $L_{ij} \leftarrow -w_{i \rightarrow j}$ 
17.       $L_{jj} += w_{i \rightarrow j}$ 
18.     $L_{1j} \leftarrow w_{\rho \rightarrow j}$ 
19.     $\triangleright$  For spanning trees, we can construct the Laplacian
      in footnote 6 by replacing line 18 with
       $L_{jj} += w_{\rho \rightarrow j}$ .
20.  return  $\mathbf{L}$ 
21. def sample_edge( $j$ ):
22.   $\triangleright$  Sample an incoming edge to  $j$  using global
      variable  $\mathbf{B}$  as in (16); requires  $\mathcal{O}(N)$  time,
       $\mathcal{O}(N)$  space.
23.   $\mathbf{m} \leftarrow \mathbf{0}$ 
24.   $\mathbf{m}_\rho += w_{\rho \rightarrow j} \mathbf{B}_{1j}$ 
25.  for  $i \in \mathcal{N} \setminus \{\rho, j\}$ :
26.     $\mathbf{m}_i += w_{i \rightarrow j} (\delta_{j \neq 1} \mathbf{B}_{jj} - \delta_{i \neq 1} \mathbf{B}_{ij})$ 
27.   $\triangleright$  For spanning trees, we can construct the
      marginals in footnote 7 by replacing line 24
      with  $\mathbf{m}_\rho += w_{\rho \rightarrow j} \mathbf{B}_{jj}$  and line 26 with
       $\mathbf{m}_i += w_{i \rightarrow j} (\mathbf{B}_{jj} - \mathbf{B}_{ij})$ .
28.  return sample from  $\mathbf{m}$ 
29. def condition( $e$ ):
30.   $\triangleright$  Condition the Laplacian and the transpose of
      its inverse to always include  $e$  in any tree; re-
      quires  $\mathcal{O}(N^2)$  time,  $\mathcal{O}(N^2)$  space.
31.  Let  $e = (i \rightarrow j)$ 
32.  if  $i = \rho$ :
33.     $\mathbf{u} \leftarrow w_{\rho \rightarrow j} \vec{\mathbf{1}}_1 - \mathbf{L}_{[:,j]}$ 
34.  else:
35.     $\mathbf{u} \leftarrow w_{i \rightarrow j} (\delta_{j \neq 1} \vec{\mathbf{1}}_j - \delta_{i \neq 1} \vec{\mathbf{1}}_i) - \mathbf{L}_{[:,j]}$ 
36.   $\mathbf{L}_{[:,j]} += \mathbf{u}$ 
37.   $\mathbf{B} -= (\mathbf{B}_{[j,:]} \mathbf{u}^\top \mathbf{B}) / (1 + \mathbf{u}^\top \mathbf{B}_{[j,:]})$ 

```

Figure 5: Algorithm for sampling dependency trees using the method of Colbourn et al. (1996). We describe the changes required to sample spanning trees in the comments.

Theorem 4. For any graph G , $\text{colbourn}(G)$ samples a dependency tree with probability

$$p(t) = \frac{1}{Z} \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \quad (21)$$

Furthermore, colbourn runs in $\mathcal{O}(N^3)$ time.

Proof. The probability of a tree $p(t)$ can equivalently be written as the product of the conditional edge marginals as in (15). To prove correctness, we prove by induction that for all $n \leq N$, at the n^{th} call to `sample_edge`, \mathbf{B} contains the transpose of the Laplacian inverse of $G \succ t_{<n}$ and an edge e with probability $p(e \mid t_{<n})$.

Base case: Then $n = 1$ and $t_{<1} = \emptyset$. \mathbf{B} contains the transpose of the Laplacian inverse of G as expected and so `sample_edge` will sample an edge e with probability $p(e)$ as expected.

Inductive step: Then $t_{<n} = [e_1, \dots, e_{n-1}]$. At the $(n - 1)^{\text{th}}$ call to `sample_edge`, \mathbf{B} contains the transpose of the Laplacian inverse of $G \succ t_{<n-1}$. We then call `condition`, which by Lemma 3 and the Sherman–Morrison formula updates \mathbf{B} to be the transpose of the Laplacian inverse of $G \succ t_{<n}$. Therefore, at the n^{th} call to `sample_edge`, \mathbf{B} contains the correct values and so `sample_edge` will sample an edge e with probability $p(e \mid t_{<n})$.

Therefore, colbourn samples a dependency tree t with the correct probability. We have N iterations of the main loop, each call to `sample_edge` takes $\mathcal{O}(N)$ times and each call to `condition` takes $\mathcal{O}(N^2)$ time. These runtimes are easily observed from the pseudocode. Therefore, colbourn has a runtime of $\mathcal{O}(N^3)$. ■

4.2 Runtime Experiment

We conduct a brief runtime experiment for `colbourn` and `wilson_rc` (see §3) whose runtimes are $\mathcal{O}(N^3)$ and $\mathcal{O}(H)$ respectively. We artificially generate random complete graphs of increasing size and measure the average sample time of each algorithm.¹¹ The results of the experiment are shown in Fig. 6. We note that despite `colbourn` being slower, the best-fit line for `colbourn`’s runtime has a slope of 1.42, suggesting it is much faster in practice than its complexity bound $\mathcal{O}(N^3)$.¹²

¹¹The experiment was conducted using an Intel(R) Core(TM) i7-7500U processor with 16GB RAM.

¹²We would expect the slope to be ≈ 3 to match the complexity bound.

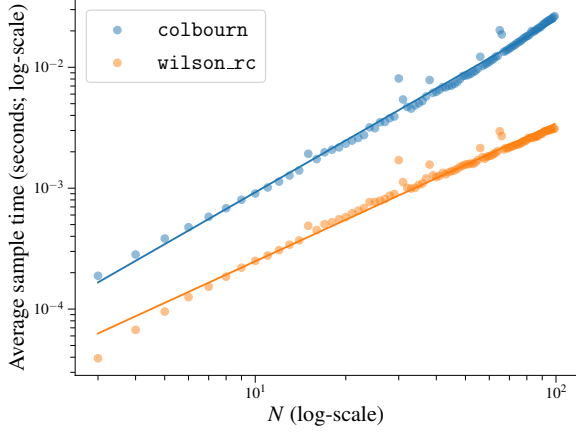


Figure 6: Runtime experiment for sampling using wilson_rc and colbourn. For each graph size, we randomly generated 100 graphs and took 20 samples from each graph. The best fit lines for colbourn and wilson_rc have slopes of 1.42 and 1.14 respectively.

5 Sampling Without Replacement

In this section, we present a novel extension to colbourn that can sample dependency trees without replacement. SWOR algorithms are useful when we must sample multiple trees from the same graph. Specifically, when the distribution of trees over the graph is skewed so that a normal sampling algorithm frequently samples the same trees. This is often the case when the edge weights have been learned using a neural model (Dozat and Manning, 2017; Ma and Hovy, 2017). The SWOR algorithm we present follows the scheme of Shi et al. (2020).

In order to use colbourn to sample without replacement, we need an expression of the edge marginals conditioned on the set of previously sampled trees. If D is the set of previously sampled trees, then we need to compute the following marginal probability efficiently

$$p(i \rightarrow j | D) = \frac{1}{Z_D} \sum_{t \in \mathcal{D}_{ij} \setminus D} w(t) \quad (22)$$

where \mathcal{D}_{ij} is the set of all dependency trees that contain edge $(i \rightarrow j)$ and

$$Z_D \stackrel{\text{def}}{=} \sum_{t \in \mathcal{D} \setminus D} w(t) = Z - \sum_{t \in D} w(t) \quad (23)$$

Lemma 4. For any graph G , set of trees D , and edge $(i \rightarrow j) \in \mathcal{E}$,

$$p(i \rightarrow j | D) = \frac{1}{Z_D} \left[Z \cdot p(i \rightarrow j) - \sum_{t \in D_{ij}} w(t) \right] \quad (24)$$

where $D_{ij} \subseteq D$ is the set of trees in D that contain the edge $(i \rightarrow j)$.

Proof.

$$p(i \rightarrow j | D) \quad (25a)$$

$$= \frac{1}{Z_D} \sum_{t \in \mathcal{D}_{ij} \setminus D} w(t) \quad (25b)$$

$$= \frac{1}{Z_D} \left[\sum_{t \in \mathcal{D}_{ij}} w(t) - \sum_{t \in D_{ij}} w(t) \right] \quad (25c)$$

$$= \frac{1}{Z_D} \left[Z p(i \rightarrow j) - \sum_{t \in D_{ij}} w(t) \right] \quad (25d)$$

■

Lemma 4 then gives a new formula to use for sampling edges by re-weighting the probability of an edge.¹³ We can then compute the marginal distribution for an incoming edges to a node in $\mathcal{O}(N + K)$ time. Note that (24) makes explicit use of Z which is not needed for the original marginals in (16). Consequently, as we sample an edge from the tree, we must condition Z as well as \mathbf{B} . Fortunately, this can be done in $\mathcal{O}(N)$ using the matrix determinant Lemma, which states that for any matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ and column vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^N$

$$|\mathbf{A} + \mathbf{u}\mathbf{v}^\top| = |\mathbf{A}|(1 + \mathbf{v}^\top \mathbf{A}^{-1} \mathbf{u}) \quad (26)$$

Furthermore, at each conditioning step, we must also update D (and Z_D), to only include the sampled trees containing the new sampled edge. These can both be achieved in $\mathcal{O}(K)$ time where K is the number of trees that we sample. The pseudocode for the sampling and conditioning steps are given as `sample_edge'` and `condition'` in Fig. 7. The sampling algorithm itself, `swor` is similar to colbourn in Fig. 5. However, it samples K dependency trees rather than a single dependency tree and stores additional variables in order to cache frequently used values such as the original Z , \mathbf{L} , and \mathbf{B} values.

Theorem 5. For any graph G and $K > 0$, `swor`(G, K) samples K dependency trees without replacement, where t is sampled with probability

$$p(t | D) = \frac{\delta_{t \notin D}}{Z_D} \prod_{(i \rightarrow j) \in t} w_{i \rightarrow j} \quad (27)$$

¹³Re-weighting has been recently used by Stanojević and Cohen (2021) to speed-up algorithms for single-root tree decoding algorithms.

Global Variables: \mathbf{L} , \mathbf{B} , D , Z , and Z_D

```

1. def swor( $G, K$ ):
2.    $\triangleright$  Sample  $K$  dependency trees without replacement from a graph  $G$ ; requires  $\mathcal{O}(KN^3 + K^2N)$  time,  $\mathcal{O}(N^2 + KN)$  space.
3.    $\mathbf{L}' \leftarrow \text{Laplacian}(G)$ 
4.    $\mathbf{Z}' \leftarrow |\mathbf{L}'|$ ;  $\mathbf{B}' \leftarrow \mathbf{L}'^{-\top}$   $\triangleright \mathcal{O}(N^3)$ 
5.    $D' \leftarrow []$ ;  $Z'_D \leftarrow Z'$ 
6.   for  $k \in \{1, \dots, K\}$  :
7.      $D \leftarrow D'$ ;  $Z \leftarrow Z'$ 
8.      $Z_D \leftarrow Z'_D$ 
9.      $\mathbf{L} \leftarrow \mathbf{L}'$ ;  $\mathbf{B} \leftarrow \mathbf{B}'$ ;  $t \leftarrow []$ 
10.    for  $j \in \mathcal{N} \setminus \{\rho\}$  :
11.       $e \leftarrow \text{sample\_edge}'(j)$ 
12.       $t.append(e)$ 
13.       $\text{condition}'(e)$ 
14.       $D'.append(t)$ ;  $Z'_D -= w(t)$ 
15.    return  $D'$ 
16. def sample\_edge'( $j$ ):
17.    $\triangleright$  Sample an incoming edge to  $j$  using global variables  $\mathbf{B}$ ,  $Z$ , and  $Z_D$  as in (24); requires  $\mathcal{O}(N)$  time,  $\mathcal{O}(N)$  space.
18.    $\mathbf{m} \leftarrow \mathbf{0}$ 
19.    $\mathbf{m}_\rho += Z w_{\rho \rightarrow j} \mathbf{B}_{1j} - \sum_{t \in T_{\rho j}} w(t)$ 
20.   for  $i \in \mathcal{N} \setminus \{\rho, j\}$  :  $\triangleright \mathcal{O}(N)$ 
21.      $\mathbf{m}_i += Z w_{i \rightarrow j} (\delta_{j \neq 1} \mathbf{B}_{jj} - \delta_{i \neq 1} \mathbf{B}_{ij})$ 
22.      $\mathbf{m}_i -= \sum_{t \in D_{ij}} w(t)$ 
23.   return sample from  $\frac{1}{Z_D} \mathbf{m}$ 
24. def condition'( $e$ ):
25.    $\triangleright$  Condition the Laplacian, the transpose of the Laplacian inverse, the partition partition function, and the set of previously sampled trees to always include  $e$  in any tree; requires  $\mathcal{O}(N^2 + K)$  time,  $\mathcal{O}(N^2)$  space.
26.   Let  $e = (i \rightarrow j)$ 
27.   if  $i = \rho$  :
28.      $\mathbf{u} \leftarrow w_{\rho \rightarrow j} \vec{\mathbf{1}}_1 - \mathbf{L}_{[:,j]}$ 
29.   else:
30.      $\mathbf{u} \leftarrow w_{i \rightarrow j} (\delta_{j \neq 1} \vec{\mathbf{1}}_j - \delta_{i \neq 1} \vec{\mathbf{1}}_i) - \mathbf{L}_{[:,j]}$ 
31.    $\mathbf{L}_{[:,j]} += \mathbf{u}$ 
32.    $\mathbf{Z} *= 1 + \mathbf{u}^\top \mathbf{B}_{[j,:]}$ 
33.    $\mathbf{B} -= (\mathbf{B}_{[j,:]} \mathbf{u}^\top \mathbf{B}) / (1 + \mathbf{u}^\top \mathbf{B}_{[j,:]})$ 
34.    $D \leftarrow D_{ij}$ 
35.    $Z_D \leftarrow Z - \sum_{t \in D} w(t)$ 

```

Figure 7: Algorithm for sampling dependency trees without replacement.

where D is the set of trees sampled prior to t . Furthermore, swor runs in $\mathcal{O}(KN^3 + K^2N)$ time.

Proof. To prove correctness, we prove by induction that the K^{th} sampled tree is sampled with the

probability in (27).

Base case: Then $K = 1$ and so $D' = \emptyset$ and $Z'_D = Z'$. Therefore, swor executes colbourn and samples tree t with probability $p(t)$ as expected.

Inductive step: Assume that D' contains $K - 1$ trees, which were each sampled with the correct probability. Then $Z'_D = Z' - \sum_{t \in D'} w(t)$ and so by Lemma 4, sample_edge' will sample the first edge of the new tree with the correct probability. We can then prove that all edges of the new tree are sampled with the correct probability using an inductive proof analogous to Theorem 4.

We require $\mathcal{O}(N^3)$ to find Z' and \mathbf{B}' . Then for each of the K sampled trees, we have N iterations of the main loop. Each call to sample_edge' takes $\mathcal{O}(N + K)$ times and each call to condition' takes $\mathcal{O}(N^2 + K)$ time. These runtimes are easily observed from the pseudocode. Therefore, swor has a runtime of $\mathcal{O}(KN^3 + K^2N)$. ■

6 Conclusion

We presented two efficient approaches to sample spanning trees subject to a root constraint, which were based on prior algorithms by Wilson (1996) and Colbourn et al. (1996). While Wilson (1996)'s $\mathcal{O}(H)$ algorithm was more rapid, Colbourn et al. (1996)'s $\mathcal{O}(N^3)$ algorithm is extendable to a novel sampling without replacement algorithm. To the best of our knowledge, not much work has been done in graph-based dependency parsing to sample dependency trees, and none has used sampling without replacement. We hope that this paper serves as a tutorial for how this can be done and encourages the use of sampling in future work.

Acknowledgments

We would like to thank all the reviewers for their invaluable feedback and time spent engaging with our work. The first author is supported by the University of Cambridge School of Technology Vice-Chancellor's Scholarship as well as by the University of Cambridge Department of Computer Science and Technology's EPSRC.

References

David J. Aldous. 1990. The random walk construction of uniform spanning trees and uniform labelled trees. *Society for Industrial and Applied Mathematics Journal of Discrete Mathematics*, 3(4):450–465.

- Nicholas Andrews, Jason Eisner, and Mark Dredze. 2012. [Name phylogeny: A generative model of string variation](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 344–355, Jeju Island, Korea. Association for Computational Linguistics.
- Eduard Bejček, Eva Hajičová, Jan Hajič, Pavlína Jínová, Václava Kettnerová, Veronika Kolářová, Marie Mikulová, Jiří Mírovský, Anna Nedoluzhko, Jarmila Panevová, Lucie Poláková, Magda Ševčíková, Jan Štěpánek, and Šárka Zikánová. 2013. [Prague dependency treebank 3.0](#).
- Andrei Z. Broder. 1989. [Generating random spanning trees](#). In *30th Annual Symposium on Foundations of Computer Science*, pages 442–447.
- Elizabeth Clark, Yangfeng Ji, and Noah A. Smith. 2018. [Neural text generation in stories using entity representations as context](#). 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 2250–2260, New Orleans, Louisiana. Association for Computational Linguistics.
- Charles J. Colbourn, Wendy J. Myrvold, and Eugene Neufeld. 1996. [Two algorithms for unranking arborescences](#). *Journal of Algorithms*, 20(2):268–281.
- Timothy Dozat and Christopher D. Manning. 2017. [Deep biaffine attention for neural dependency parsing](#). In *Proceedings of the International Conference on Learning Representations*.
- William Fedus, Ian J. Goodfellow, and Andrew M. Dai. 2018. [MaskGAN: Better text generation via filling in the _____](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Harold N. Gabow and Robert Endre Tarjan. 1984. [Efficient algorithms for a family of matroid intersection problems](#). *Journal of Algorithms*, 5(1).
- Andreas Griewank and Andrea Walther. 2008. *Evaluating Derivatives—Principles and Techniques of Algorithmic Differentiation*, 2 edition. Society for Industrial and Applied Mathematics.
- A. Guénoche. 1983. [Random spanning tree](#). *Journal of Algorithms*, 4(3):214–220.
- Katherine Keith, Su Lin Blodgett, and Brendan O’Connor. 2018. [Monte Carlo syntax marginals for exploring and using dependency parses](#). 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 917–928, New Orleans, Louisiana. Association for Computational Linguistics.
- Gustav Kirchhoff. 1847. [Über die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird](#). *Annalen der Physik*, 148(12).
- Terry Koo, Amir Globerson, Xavier Carreras, and Michael Collins. 2007. [Structured prediction models via the matrix-tree theorem](#). In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*.
- Sandra Kübler, Ryan T. McDonald, and Joakim Nivre. 2009. *Dependency Parsing*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- Vidyadhar G. Kulkarni. 1990. [Generating random combinatorial objects](#). *Journal of Algorithms*, 11(2):185–207.
- Gregory F. Lawler. 1979. *A Self-Avoiding Random Walk*. Ph.D. thesis, Princeton University.
- Robert L. Logan IV, Matt Gardner, and Sameer Singh. 2020. [On importance sampling-based evaluation of latent language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2171–2176, Online. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 2017. [Neural probabilistic model for non-projective MST parsing](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- David Mareček and Zdeněk Žabokrtský. 2011. [Gibbs sampling with treeness constraint in unsupervised dependency parsing](#). In *Proceedings of Workshop on Robust Unsupervised and Semisupervised Methods in Natural Language Processing*, pages 1–8, Hissar, Bulgaria. Association for Computational Linguistics.
- David McClosky, Mihai Surdeanu, and Christopher Manning. 2011. [Event extraction as dependency parsing for BioNLP 2011](#). In *Proceedings of BioNLP Shared Task 2011 Workshop*, pages 41–45, Portland, Oregon, USA. Association for Computational Linguistics.
- Ryan McDonald and Giorgio Satta. 2007. [On the complexity of non-projective data-driven dependency parsing](#). In *Proceedings of the Tenth International Conference on Parsing Technologies*, Prague, Czech Republic. Association for Computational Linguistics.
- Andriy Mnih and Geoffrey E. Hinton. 2007. [Three new graphical models for statistical language modelling](#). In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007*, pages 641–648.

Tetsuji Nakagawa. 2007. [Multilingual dependency parsing using global features](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 952–956, Prague, Czech Republic. Association for Computational Linguistics.

Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Katya Aplonova, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, Victoria Basmov, John Bauer, Sandra Bellato, Kepa Bengoetxea, Yevgeni Berzak, Irshad Ahmad Bhat, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Rogier Blokland, Victoria Bobicev, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Flavio Massimiliano Cecchini, Giuseppe G. A. Celano, Slavomír Čéplö, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilaraza, Carly Dickerson, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Tomáš Erjavec, Aline Etienne, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Sebastian Garza, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta Gonzáles Saavedra, Matias Gioni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena Hwang, Radu Ion, Elena Irimia, Olajídé Ishola, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Boris Katz, Tolga Kayadelen, Jessica Kenney, Václava Kettnerová, Jesse Kirchner, Kamil Kopacewicz, Natalia Kotsyba, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Lucia Lam, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phuong Lê Hồng, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Nikola Ljubešić, Olga Logina, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Gustavo Mendonça, Niko Miekka, Margarita Misirpashayeva, Anna Missilä, Cătălin Mititelu, Yusuke

Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Keiko Sophie Mori, Shinsuke Mori, Bjartur Mortensen, Bohdan Moskalevskiy, Kadri Muischnek, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horňiáček, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Luong Nguyễn Thị, Huyền Nguyễn Thị Minh, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Adédayo Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Guilherme Paulino-Passos, Siyao Peng, Cene Augusto Perez, Guy Perrier, Slav Petrov, Jussi Piitulainen, Emily Pitler, Barbara Plank, Thierry Poibeau, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tina Puolakainen, Sampo Pyysalo, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Michael Rießler, Larissa Rinaldi, Laura Rituma, Luísa Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roşca, Olga Rudina, Jack Rueter, Shoval Sadde, Benoît Sagot, Shadi Saleh, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamel Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Muh Shohibussirri, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Carolyn Spadine, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Zsolt Szántó, Dima Taji, Yuta Takahashi, Takaaki Tanaka, Isabelle Tellier, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uribe, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Eric Villemonte de la Clergerie, Veronika Vincze, Lars Wallin, Jing Xian Wang, Jonathan North Washington, Seyi Williams, Mats Wirén, Tsegay Wolde-mariam, Tak-sum Wong, Chunxiao Yan, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, Manying Zhang, and Hanzhi Zhu. 2018. [Universal dependencies 2.3](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.

Wenzhe Pei, Tao Ge, and Baobao Chang. 2015. [An effective neural network model for graph-based dependency parsing](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 313–322, Beijing, China. Association for Computational Linguistics.

Jack Sherman and Winifred J. Morrison. 1950. [Adjustment of an inverse matrix corresponding to a change in one element of a given matrix](#). *The Annals of Mathematical Statistics*, 21(1):124–127.

- Kensen Shi, David Bieber, and Charles Sutton. 2020. [Incremental sampling without replacement for sequence models](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, pages 8785–8795.
- Sameer Singh, Michael Wick, and Andrew McCallum. 2012. [Monte Carlo MCMC: Efficient inference by approximate sampling](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1104–1113, Jeju Island, Korea. Association for Computational Linguistics.
- David A. Smith and Noah A. Smith. 2007. [Probabilistic models of nonprojective dependency trees](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, Prague, Czech Republic. Association for Computational Linguistics.
- Miloš Stanojević and Shay B. Cohen. 2021. [A root of a problem: Optimizing single-root dependency parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online. Association for Computational Linguistics.
- W. T. Tutte. 1984. *Graph Theory*. Addison-Wesley Publishing Company.
- Wenhui Wang and Baobao Chang. 2016. [Graph-based dependency parsing with bidirectional LSTM](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2306–2315, Berlin, Germany. Association for Computational Linguistics.
- Stanley G. Williamson. 1985. *Combinatorics for Computer Science*. Computer Science Press.
- David Bruce Wilson. 1996. [Generating random spanning trees more quickly than the cover time](#). In *Proceedings of the Twenty-Eighth Annual Association for Computing Machinery Symposium on the Theory of Computing*, pages 296–303.
- Yuan Zhang, Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2014a. [Greed is good if randomized: New inference for dependency parsing](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1013–1024, Doha, Qatar. Association for Computational Linguistics.
- Yuan Zhang, Tao Lei, Regina Barzilay, Tommi Jaakkola, and Amir Globerson. 2014b. [Steps to excellence: Simple inference with refined scoring of dependency trees](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 197–207, Baltimore, Maryland. Association for Computational Linguistics.
- Ran Zmigrod, Tim Vieira, and Ryan Cotterell. 2020. [Please mind the root: Decoding arborescences for dependency parsing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4809–4819, Online. Association for Computational Linguistics.
- Ran Zmigrod, Tim Vieira, and Ryan Cotterell. 2021. [Efficient computation of expectations under spanning tree distributions](#). *Transactions of the Association for Computational Linguistics*, 9:675–690.