Identifiability and Unmixing of Latent Parse Trees

Daniel Hsu
Microsoft Research

Sham M. Kakade
Microsoft Research

Percy Liang
Stanford University

Abstract

This paper explores unsupervised learning of parsing models along two directions. First, which models are identifiable from infinite data? We use a general technique for numerically checking identifiability based on the rank of a Jacobian matrix, and apply it to several standard constituency and dependency parsing models. Second, for identifiable models, how do we estimate the parameters efficiently? EM suffers from local optima, while recent work using spectral methods [1] cannot be directly applied since the topology of the parse tree varies across sentences. We develop a strategy, unmixing, which deals with this additional complexity for restricted classes of parsing models.

1 Introduction

Generative parsing models, which define joint distributions over sentences and their parse trees, are one of the core techniques in computational linguistics. We are interested in the unsupervised learning of these models [2–6], where the goal is to estimate the model parameters given only examples of sentences. Unsupervised learning can fail for a number of reasons [7]: model misspecification, non-identifiability, estimation error, and computation error. In this paper, we delve into two of these issues: identifiability and computation. In doing so, we confront a central challenge of parsing models—that the topology of the parse tree is unobserved and varies across sentences. This is in contrast to standard phylogenetic models [8] and other latent tree models for which there is a single fixed global tree across all examples [9].

A model is identifiable if there is enough information in the data to pinpoint the parameters (up to some trivial equivalence class); establishing the identifiability of a model is often a highly non-trivial task. A classic result of Kruskal [10] has been employed to prove the identifiability of a wide class of latent variable models, including hidden Markov models and certain restricted mixtures of latent tree models [11–13]. However, these techniques cannot be directly applied to parsing models since the tree topology varies over an exponential set of possible topologies. Instead, we turn to techniques from algebraic geometry [14–17]; we show that a simple numerical procedure can be used to check identifiability for a wide class of models in NLP. Using this tool, we discover that probabilistic context-free grammars (PCFGs) are non-identifiable, but that simpler PCFG variants and dependency models are identifiable.

The most common way to estimate unsupervised parsing models is by using local techniques such as EM [18] or MCMC sampling [19], but these methods can suffer from local optima and slow mixing. Meanwhile, recent work [1, 20–23] has shown that spectral methods can be used to estimate mixture models and HMMs with provable guarantees. These techniques express low-order moments of the observable distribution as a product of matrix parameters and use eigenvalue decomposition to recover these matrices. However, these methods are not directly applicable to parsing models because the tree topology again varies non-trivially. To address this, we propose a new technique, *unmixing*. The main idea is to express moments of the observable distribution as a mixture over the possible topologies. For restricted parsing models, the moments for a fixed tree structure can

E-mail: dahsu@microsoft.com, skakade@microsoft.com, pliang@cs.stanford.edu
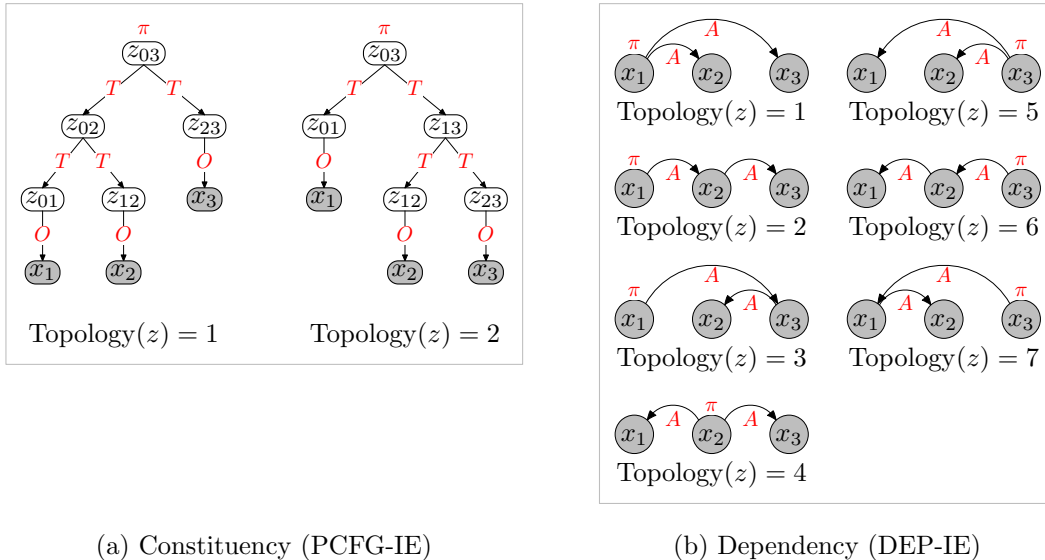
(a) Constituency (PCFG-IE) (b) Dependency (DEP-IE)

Figure 1: The two constituency trees and seven dependency trees over $L = 3$ words, x_1, x_2, x_3. (a) A constituency tree consists of a hierarchical grouping of the words with a latent state z_v for each node v. (b) A dependency tree consists of a collection of directed edges between the words. In both cases, we have labeled each edge from i to j with the parameters used to generate the state of node j given i.

be "unmixed", thereby reducing the problem to one with a fixed topology, which can be tackled using standard techniques [1]. Importantly, our unmixing technique does not require the training sentences be annotated with the tree topologies *a priori*, in contrast to recent extensions of [21] to learning PCFGs [24] and dependency trees [25, 26], which work on a fixed topology.

2 Notation

For a positive integer n, define $[n] \stackrel{\text{def}}{=} \{1, \ldots, n\}$ and $\langle n \rangle = \{e_1, \ldots, e_n\}$, where e_i is the vector which is 1 in component i and 0 elsewhere. For integers $a, b \in [n]$, let $a \otimes_n b = (a-1)n+b \in [n^2]$ be the integer encoding of the pair (a, b). For a pair of matrices, $A, B \in \mathbb{R}^{m \times n}$, define the columnwise tensor product $A \otimes^{\text{c}} B \in \mathbb{R}^{m^2 \times n}$ to be such that $(A \otimes^{\text{c}} B)_{(i_1 \otimes_m i_2)j} = A_{i_1 j} B_{i_2 j}$. For a matrix $A \in \mathbb{R}^{m \times n}$, let $A^\dagger$ denote the Moore-Penrose pseudoinverse.

3 Parsing models

A sentence is a sequence of L words, $\mathbf{x} = (x_1, \ldots, x_L)$, where each word $x_i \in \langle d \rangle$ is one of d possible word types. A (generative) *parsing model* defines a joint distribution $\mathbb{P}_\theta(\mathbf{x}, z)$ over a sentence $\mathbf{x}$ and its parse tree z (to be made precise later), where θ are the model parameters (a collection of multinomials). Each parse tree z has a *topology* $\text{Topology}(z) \in \text{Topologies}$, which is both unobserved and varying across sentences. The learning problem is to recover θ given only samples of $\mathbf{x}$.

Two important classes of models of natural language syntax are constituency models, which represent a hierarchical grouping and labeling of the phrases of a sentence (e.g., Figure 1(a)), and dependency models, which represent pairwise relationships between the words of a sentence (e.g., Figure 1(b)).

3.1 Constituency models

A *constituency tree* $z = (V, s)$ consists of a set of nodes V and a collection of hidden states $s = \{s_v\}_{v \in V}$. Each state $s_v \in \langle k \rangle$ represents one of k possible syntactic categories. Each node v has the form $[i : j]$ for $0 \leq i < j \leq L$ corresponding to the phrase between positions i and j of the sentence. These nodes form a binary tree as follows: the root node is $[0 : L] \in V$, and for each node $[i : j] \in V$ with $j - i > 1$, there exists a unique m with $i < m < j$ defining the two children nodes $[i : m] \in V$ and $[m : j] \in V$. Let $\text{Topology}(z)$ be an integer encoding of V.

PCFG. Perhaps the most well-known constituency parsing model is the probabilistic context-free grammar (PCFG). The parameters of a PCFG are $\theta = (\pi, B, O)$, where $\pi \in \mathbb{R}^k$ specifies the initial state distribution, $B \in \mathbb{R}^{k^2 \times k}$ specifies the binary production distributions, and $O \in \mathbb{R}^{d \times k}$ specifies the emission distributions.

A PCFG corresponds to the following generative process (see Figure 1(a) for an example): choose a topology $\text{Topology}(z)$ uniformly at random; generate the state of the root node using π; recursively generate pairs of children states given their parents using B; and finally generate words x_i given their parents using O. This generative process defines a joint probability over a sentence $\mathbf{x}$ and a parse tree z:

$$\mathbb{P}_\theta(\mathbf{x}, z) = |\text{Topologies}|^{-1} \pi^\top s_{[0:L]} \prod_{[i:m],[m:j] \in V} (s_{[i:m]} \otimes_k s_{[m:j]})^\top B s_{[i:j]} \prod_{i=1}^{L} x_i^\top O s_{[i-1:i]}, \quad (1)$$

We will also consider two variants of the PCFG with additional restrictions:

PCFG-I. The left and right children states are generated independently—that is, we have the following factorization: $B = T_1 \otimes^c T_2$ for some $T_1, T_2 \in \mathbb{R}^{k \times k}$.

PCFG-IE. The left and the right productions are independent and equal: $B = T \otimes^c T$.

3.2 Dependency tree models

In contrast to constituency trees, which posit internal nodes with latent states, dependency trees connect the words directly. A *dependency tree* z is a set of directed edges (i, j), where $i, j \in [L]$ are distinct positions in the sentence. Let $\text{Root}(z)$ denote the position of the root node of z. We consider only *projective* dependency trees [27]: z is projective if for every path from i to j to k in z, we have that j and k are on the same side of i (that is, $j - i$ and $k - i$ have the same sign). Let $\text{Topology}(z)$ be an integer encoding of z.

DEP-I. We consider the simple dependency model of [4]. The parameters of this model are $\theta = (\pi, A_\swarrow, A_\searrow)$, where $\pi \in \mathbb{R}^d$ is the initial word distribution and $A_\swarrow, A_\searrow \in \mathbb{R}^{d \times d}$ are the left and right argument distributions. The generative process is as follows: choose a topology $\text{Topology}(z)$ uniformly at random, generate the root word using π, and recursively generate argument words to the left to the right given the parent word using $A_\swarrow$ and $A_\searrow$, respectively. The corresponding joint probability distribution is as follows:

$$\mathbb{P}_\theta(\mathbf{x}, z) = |\text{Topologies}|^{-1} \pi^\top x_{\text{Root}(z)} \prod_{(i,j) \in z} x_j^\top A_{\text{dir}(i,j)} x_i, \quad (2)$$

where $\text{dir}(i, j) = \swarrow$ if $j < i$ and $\searrow$ if $j > i$.

We also consider the following two variants:

DEP-IE. The left and right argument distributions are equal: $A = A_\swarrow = A_\searrow$.

DEP-IES. $A = A_\swarrow = A_\searrow$ and π is the stationary distribution of A (that is, $\pi = A\pi$).

Usually a PCFG induces a topology via a state-dependent probability of choosing a binary production versus an emission. Our model is a restriction which corresponds to a state-independent probability.

4 Identifiability

Our goal is to estimate model parameters $\theta_0 \in \Theta$ given only access to sentences $\mathbf{x} \sim \mathbb{P}_{\theta_0}$. Specifically, suppose we have an *observation function* $\phi(\mathbf{x}) \in \mathbb{R}^m$, which is the only lens through which an algorithm can view the data. We ask a basic question: in the limit of infinite data, is it information-theoretically possible to identify θ_0 from the *observed moments* $\mu(\theta_0) \overset{\text{def}}{=} \mathbb{E}_{\theta_0}[\phi(\mathbf{x})]$?

To be more precise, define the *equivalence class* of θ_0 to be the set of parameters θ that yield the same observed moments:

$$\mathcal{S}_\Theta(\theta_0) = \{\theta \in \Theta : \mu(\theta) = \mu(\theta_0)\}. \tag{3}$$

It is impossible for an algorithm to distinguish among the elements of $\mathcal{S}_\Theta(\theta_0)$. Therefore, one might want to ensure that $|\mathcal{S}_\Theta(\theta_0)| = 1$ for all $\theta_0 \in \Theta$. However, this requirement is too strong for two reasons. First, models often have natural symmetries—e.g., the k states of any PCFG can be permuted without changing $\mu(\theta)$, so $|\mathcal{S}_\Theta(\theta_0)| \geq k!$. Second, $|\mathcal{S}_\Theta(\theta_0)| = \infty$ for some pathological θ_0's—e.g., PCFGs where all states have the same emission distribution O are indistinguishable regardless of the production distributions B. The following definition of identifiability accommodates these two exceptional cases:

Definition 1 (Identifiability). *A model family with parameter space Θ is* (globally) identifiable from ϕ if there exists a measure zero set $\mathcal{E}$ such that $|\mathcal{S}_\Theta(\theta_0)|$ is finite for every $\theta_0 \in \Theta\backslash\mathcal{E}$. It is locally identifiable from ϕ if there exists a measure zero set $\mathcal{E}$ such that, for every $\theta_0 \in \Theta\backslash\mathcal{E}$, there exists an open neighborhood $N(\theta_0)$ around θ_0 such that $\mathcal{S}_\Theta(\theta_0) \cap N(\theta_0) = \{\theta_0\}$.*

Example of non-identifiability. Consider the DEP-IE model with $L = 2$ with the full observation function $\phi(\mathbf{x}) = x_1 \otimes x_2$. The corresponding observed moments are $\mu(\theta) = 0.5A\operatorname{diag}(\pi) + 0.5\operatorname{diag}(\pi)A^\top$. Note that $A\operatorname{diag}(\pi)$ is an arbitrary $d \times d$ matrix whose entries sum to 1, which has $d^2 - 1$ degrees of freedom, whereas $\mu(\theta)$ is a symmetric matrix whose entries sum to 1, which has $\binom{d+1}{2} - 1$ degrees of freedom. Therefore, $\mathcal{S}_\Theta(\theta)$ has dimension $\binom{d}{2}$ and therefore the model is non-identifiable.

Parameter counting. It is important to compute the degrees of freedom correctly—simple parameter counting is insufficient. For example, consider the PCFG-IE model with $L = 2$. The observed moments with respect to $\phi(\mathbf{x}) = x_1 \otimes x_2$ is a $d \times d$ matrix, which places d^2 constraints on the $k^2 + (d-1)k$ parameters. When $d \geq 2k$, there are more constraints than parameters, but the PCFG-IE model with $L = 2$ is actually non-identifiable (as we will see later). The issue here is that the number of constraints does not reveal the fact that some of these constraints are redundant.

4.1 Observation functions

An observation function $\phi(\mathbf{x})$ and its associated observed moments $\mu(\theta_0) = \mathbb{E}_{\theta_0}[\phi(\mathbf{x})]$ reveals aspects of the distribution $\mathbb{P}_{\theta_0}(\mathbf{x})$. For example, $\phi(\mathbf{x}) = x_1$ would only reveal the marginal distribution of the first word, whereas $\phi(\mathbf{x}) = x_1 \otimes \cdots \otimes x_L$ reveals the entire distribution of $\mathbf{x}$. There is a tradeoff: Higher-order moments provide more information, but are harder to estimate reliably given finite data, and are also computationally more expensive. In this paper, we consider the following intermediate moments:

$$\phi_{12}(\mathbf{x}) \overset{\text{def}}{=} x_1 \otimes x_2 \qquad\qquad \phi_{**}(\mathbf{x}) \overset{\text{def}}{=} \big(x_i \otimes x_j : i, j \in [L]\big)$$

$$\phi_{123}(\mathbf{x}) \overset{\text{def}}{=} x_1 \otimes x_2 \otimes x_3 \qquad\qquad \phi_{***}(\mathbf{x}) \overset{\text{def}}{=} \big(x_i \otimes x_j \otimes x_k : i, j, k \in [L]\big)$$

$$\phi_{123\eta}(\mathbf{x}) \overset{\text{def}}{=} (x_1 \otimes x_2)(\eta^\top x_3) \qquad\qquad \phi_{***\eta}(\mathbf{x}) \overset{\text{def}}{=} \big((x_i \otimes x_j)(\eta^\top x_k) : i, j, k \in [L]\big)$$

$$\phi_{\text{all}}(\mathbf{x}) \overset{\text{def}}{=} x_1 \otimes \cdots \otimes x_L$$

Above, $\eta \in \mathbb{R}^d$ denotes a unit vector in $\mathbb{R}^d$ (e.g., e_1) which picks out a linear combination of matrix slices from a third-order $d \times d \times d$ tensor.

4.2 Automatically checking identifiability

One immediate goal is to determine which models in Section 3 are identifiable from which of the observed moments (Section 4.1). A powerful analytic tool that has been succesfully applied in

previous work is Kruskal's theorem [10, 11], but (i) it is does not immediately apply to models with random topologies, and (ii) only gives sufficient conditions for identifiability, and cannot be used to determine non-identifiability. Furthermore, since it is common practice to explore many different models for a given problem in rapid succession, we would like to check identifiability quickly and reliably. In this section, we develop an automatic procedure to do this.

To establish identifiability, let us examine the algebraic structure of $\mathcal{S}_\Theta(\theta_0)$ for $\theta_0 \in \Theta$, where we assume that the parameter space Θ is an open subset of $[0, 1]^n$. Recall that $\mathcal{S}_\Theta(\theta_0)$ is defined by the moment constraints $\mu(\theta) = \mu(\theta_0)$. We can write these constraints as $h_{\theta_0}(\theta) = 0$, where

$$h_{\theta_0}(\theta) \stackrel{\text{def}}{=} \mu(\theta) - \mu(\theta_0)$$

is a vector of m polynomials in θ.

Let us now compute the number of degrees of freedom of h_{θ_0} around θ_0. The key quantity is $J(\theta) \in \mathbb{R}^{m \times n}$, the Jacobian of h_{θ_0} at θ (note that the Jacobian of h_{θ_0} does not depend on θ_0; it is precisely the Jacobian of μ). This Jacobian criterion is well-established in algebraic geometry, and has been adopted in the statistical literature for testing model identifiability and other related properties [14–17].

Intuitively, each row of $J(\theta_0)$ corresponds to a direction of a constraint violation, and thus the row space of $J(\theta_0)$ corresponds to all directions that would take us outside the equivalence class $\mathcal{S}_\Theta(\theta_0)$. If $J(\theta_0)$ has less than rank n, then there is a direction orthogonal to all the rows along which we can move and still satisfy all the constraints—in other words, $|\mathcal{S}_\Theta(\theta_0)|$ is infinite, and therefore the model is non-identifiable. This intuition leads to the following algorithm:

CHECKIDENTIFIABILITY:
  1. Choose a point $\tilde{\theta} \in \Theta$ uniformly at random.
  2. Compute the Jacobian matrix $J(\tilde{\theta})$.
  3. Return "yes" if the rank of $J(\tilde{\theta}) = n$ and "no" otherwise.

The following theorem asserts the correctness of CHECKIDENTIFIABILITY. It is largely based on techniques in [16], although we have not seen it explicitly stated in this form.

Theorem 1 (Correctness of CHECKIDENTIFIABILITY). *Assume the parameter space Θ is a non-empty open connected subset of $[0, 1]^n$; and the observed moments $\mu \colon \mathbb{R}^n \to \mathbb{R}^m$, with respect to observation function ϕ, is a polynomial map. Then with probability 1, CHECKIDENTIFIABILITY returns "yes" iff the model family is locally identifiable from ϕ. Moreover, if it returns "yes", then there exists $\mathcal{E} \subset \Theta$ of measure zero such that the model family with parameter space $\Theta \setminus \mathcal{E}$ is identifiable from ϕ.*

The proof of Theorem 1 is given in Appendix A.

4.3 Implementation of CHECKIDENTIFIABILITY

Computing the Jacobian. The rows of J correspond to $\partial \mathbb{E}_\theta[\phi_j(\mathbf{x})]/\partial \theta$ and can be computed efficiently by adapting dynamic programs used in the E-step of an EM algorithm for parsing models. There are two main differences: (i) we must sum over possible values of $\mathbf{x}$ in addition to z, and (ii) we are not computing moments, but rather gradients thereof. Specifically, we adapt the CKY algorithm for constituency models and the algorithm of [27] for dependency models. See Appendix C.1 for more details.

Numerical issues. Because we implemented CHECKIDENTIFIABILITY on a finite precision machine, the results are subject to numerical precision errors. However, we verified that our numerical results are consistent with various analytically-derived identifiability results (e.g., from [11]).

While we initially defined θ to be a tuple of conditional probability matrices, we will now use its non-redundant vectorized form $\theta \in \mathbb{R}^n$.

Model \ Observation function	ϕ_{12}	ϕ_{**}	ϕ_{123e_1}	ϕ_{123}	ϕ_{***e_1}	ϕ_{***}
PCFG	No, even from ϕ_{all} for $L \in \{3,4,5\}$					
PCFG-I / PCFG-IE	No	Yes iff $L \geq 4$	Yes iff $L \geq 3$			
DEP-I	No	Yes iff $L \geq 3$				
DEP-IE / DEP-IES	Yes iff $L \geq 3$					

Figure 2: Local identifiability of parsing models. These findings are given by CHECKIDENTIFIABILITY have the semantics from Theorem 1. These were checked for $d \in \{2,3,\ldots,8\}$, $k \in \{2,\ldots,d\}$ (applies only for PCFG models), $L \in \{2,3,\ldots,9\}$.

4.4 Identifiability of constituency and dependency tree models

We checked the identifiability status of various constituency and dependency tree models using our implementation of CHECKIDENTIFIABILITY. We focus on the regime where $d \geq k$ for PCFGs; additional results for $d < k$ are given in Appendix B.

The results are reported in Figure 2. First, we found that the PCFG is not identifiable from ϕ_{all} (and therefore not identifiable from any ϕ) for $L \in \{3,4,5\}$; we believe that the same holds for all L. This negative result motivates exploring restricted subclasses of PCFGs, such as PCFG-I and PCFG-IE, which factorize the binary productions. For these classes, we found that the sentence length L and choice of observation function can influence identifiability: Both models are identifiable for large enough L (e.g., $L \geq 3$) and with a sufficiently rich observation function (e.g., $\phi_{123\eta}$).

The dependency models, DEP-I and DEP-IE, were all found to be identifiable for $L \geq 3$ from second-order moments ϕ_{**}. The conditions for identifiability are less stringent than their constituency counterparts (PCFG-I and PCFG-IE), which is natural since dependency models are simpler without the latent states. Note that in all identifiable models, second-order moments suffice to determine the distribution—this is good news because low-order moments are easier to estimate.

5 Unmixing algorithms

Having established which parsing models are identifiable, we now turn to parameter estimation for these models. We will consider algorithms based on moment matching—those that try to find a θ satisfying $\mu(\theta) = u$ for some u. Typically, u is an empirical estimate of $\mu(\theta_0) = \mathbb{E}_{\theta_0}[\phi(\mathbf{x})]$ based on samples $\mathbf{x} \sim \mathbb{P}_{\theta_0}$.

In general, solving $\mu(\theta) = u$ corresponds to finding solutions to systems of multivariate polynomials, which is NP-hard [28]. However, $\mu(\theta)$ often has additional structure which we can exploit. For instance, for an HMM, the sliced third-order moments $\mu_{123\eta}(\theta)$ can be written as a product of parameter matrices in θ, and each matrix can be recovered by decomposing the product [1].

For parsing models, the challenge is that the topology is random, so the moments is not a single product, but a mixture over products. To deal with this complication, we propose a new technique, which we call *unmixing*: We "unmix" the products from the mixtures, essentially reducing the problem to one with a fixed topology.

We will first present the general idea of unmixing (Section 5.1) and then apply it to the PCFG-IE model (Section 5.2) and the DEP-IES model (Section 5.3).

5.1 General case

We assume the observation function $\phi(\mathbf{x})$ consists of a collection of observation matrices $\{\phi_o(\mathbf{x})\}_{o \in \mathcal{O}}$ (e.g., for $o = (i,j)$, $\phi_o(\mathbf{x}) = x_i \otimes x_j$). Given an observation matrix $\phi_o(\mathbf{x})$ and a topology $t \in \text{Topologies}$, consider the mapping that computes the observed moment conditioned on

---

Note that these subclasses occupy measure zero subsets of the PCFG parameter space, which is expected given the non-identifiability of the general PCFG.

We will develop our algorithms assuming true moments ($u = \mu(\theta_0)$). The empirical moments converge to the true moments at $O_p(n^{-\frac{1}{2}})$, and matrix perturbation arguments (e.g., [1]) can be used derive sample complexity arguments for the parameter error.

that topology: $\Psi_{o,t}(\theta) = \mathbb{E}_\theta[\phi_o(\mathbf{x}) \mid \text{Topology} = t]$. As we range o over $\mathcal{O}$ and t over Topologies, we will enounter a finite number of such mappings. We call these mappings *compound parameters*, denoted $\{\Psi_p\}_{p \in \mathcal{P}}$.

Now write the observed moments as a weighted sum:

$$\mu_o(\theta) = \sum_{p \in \mathcal{P}} \underbrace{\mathbb{P}(\Psi_{o,\text{Topology}} = \Psi_p)}_{\overset{\text{def}}{=} M_{op}} \Psi_p \quad \text{for all } o \in \mathcal{O}, \tag{4}$$

where we have defined M_{op} to be the probability mass over tree topologies that yield compound parameter Ψ_p. We let $\{M_{op}\}_{o \in \mathcal{O}, p \in \mathcal{P}}$ be the *mixing matrix*. Note that (4) defines a system of equations $\mu = M\Psi$, where the variables are the compound parameters and the constraints are the observed moments. In a sense, we have replaced the original system of polynomial equations (in θ) with a system of linear equations (in Ψ).

The key to the utility of this technique is that the number of compound parameters can be polynomial in L even when the number of possible topologies is exponential in L. Previous analytic techniques [13] based on Kruskal's theorem [10] cannot be applied here because the possible topologies are too many and too varied.

Note that the mixing equation $\mu = M\Psi$ holds for each sentence length L, but many compound parameters p appear in the equations of multiple L. Therefore, we can combine the equations across all observed sentence lengths, yielding a more constrained system than if we considered the equations of each L separately.

The following proposition shows how we can recover θ by unmixing the observed moments μ:

Proposition 1 (Unmixing). *Suppose that there exists an efficient base algorithm to recover θ from some subset of compound parameters $\{\Psi_p(\theta) : p \in \mathcal{P}_0\}$, and that $e_p^\top$ is in the row space of M for each $p \in \mathcal{P}_0$. Then we can recover θ as follows:*

UNMIX(μ):
1. Compute the mixing matrix M (4).
2. Retrieve the compound parameters $\Psi_p(\theta) = (M^\dagger \mu)_p$ for each $p \in \mathcal{P}_0$.
3. Call the base algorithm on $\{\Psi_p(\theta) : p \in \mathcal{P}_0\}$ to obtain θ.

For all our parsing models, M can be computed efficiently using dynamic programming (Appendix C.2). Note that M is data-independent, so this computation can be done once in advance.

5.2   Application to the PCFG-IE model

As a concrete example, consider the PCFG-IE model over $L = 3$ words. Write $A = OT$. For any $\eta \in \mathbb{R}^d$, we can express the observed moments as a sum over the two possible topologies in Figure 1(a):

$$\mu_{123\eta} \overset{\text{def}}{=} \mathbb{E}[x_1 \otimes x_2(\eta^\top x_3)] = 0.5\Psi_{1;\eta} + 0.5\Psi_{2;\eta}, \qquad \Psi_{1;\eta} \overset{\text{def}}{=} A \operatorname{diag}(T \operatorname{diag}(\pi)A^\top \eta)A^\top,$$

$$\mu_{132\eta} \overset{\text{def}}{=} \mathbb{E}[x_1 \otimes x_3(\eta^\top x_2)] = 0.5\Psi_{3;\eta} + 0.5\Psi_{2;\eta}, \qquad \Psi_{2;\eta} \overset{\text{def}}{=} A \operatorname{diag}(\pi)T^\top \operatorname{diag}(A^\top \eta)A^\top,$$

$$\mu_{231\eta} \overset{\text{def}}{=} \mathbb{E}[x_2 \otimes x_3(\eta^\top x_1)] = 0.5\Psi_{3;\eta} + 0.5\Psi_{1;\eta}, \qquad \Psi_{3;\eta} \overset{\text{def}}{=} A \operatorname{diag}(A^\top \eta)T \operatorname{diag}(\pi)A^\top,$$

or compactly in matrix form:

$$\underbrace{\begin{pmatrix} \mu_{123\eta} \\ \mu_{132\eta} \\ \mu_{231\eta} \end{pmatrix}}_{\text{observed moments } \mu_\eta} = \underbrace{\begin{pmatrix} 0.5I & 0.5I & 0 \\ 0 & 0.5I & 0.5I \\ 0.5I & 0 & 0.5I \end{pmatrix}}_{\text{mixing matrix } M} \underbrace{\begin{pmatrix} \Psi_{1;\eta} \\ \Psi_{2;\eta} \\ \Psi_{3;\eta} \end{pmatrix}}_{\text{compound parameters } \Psi_\eta} .$$

Let us observe μ_η at two different values of η, say at $\eta = \mathbf{1}$ and $\eta = \tau$ for some random τ. Since the mixing matrix M is invertible, we can obtain the compound parameters $\Psi_{2;\mathbf{1}} = (M^{-1}\mu_1)_2$ and $\Psi_{2;\tau} = (M^{-1}\mu_\tau)_2$.

Now we will recover θ from $\Psi_{2;\mathbf{1}}$ and $\Psi_{2;\tau}$ by first extracting $A = OT$ via an eigenvalue decomposition, and then recovering π, T, and O in turn (all up to the same unknown permutation) via elementary matrix operations.

For the first step, we will use the following tool (adapted from Algorithm A of [1]), which allow us to decompose two related matrix products:

Lemma 1 (Spectral decomposition). *Let $M_1, M_2 \in \mathbb{R}^{d \times k}$ have full column rank and D be a diagonal matrix with distinct diagonal entries. Suppose we observe $X = M_1 M_2^\top$ and $Y = M_1 D M_2^\top$. Then* DECOMPOSE(X, Y) *recovers M_1 up to a permutation and scaling of the columns.*

DECOMPOSE(X, Y):
    1. Find $U_1, U_2 \in \mathbb{R}^{d \times k}$ such that $\mathrm{range}(U_1) = \mathrm{range}(X)$ and $\mathrm{range}(U_2) = \mathrm{range}(X^\top)$.
    2. Perform an eigenvalue decomposition of $(U_1^\top Y U_2)(U_1^\top X U_2)^{-1} = V S V^{-1}$.
    3. Return $(U_1^\top)^\dagger V$.

First, run DECOMPOSE$(X = \Psi_{2;\mathbf{1}}^\top, Y = \Psi_{2;\tau}^\top)$ (Lemma 1), which corresponds to $M_1 = A$ and $M_2 = A \operatorname{diag}(\pi) T^\top$. This produces $A\Pi S$ for some permutation matrix Π and diagonal scaling S. Since we know that the columns of A sum to one, we can identify $A\Pi$.

To recover the initial distribution π (up to permutation), take $\Psi_{2;\mathbf{1}}\mathbf{1} = A\pi$ and left-multiply by $(A\Pi)^\dagger$ to get $\Pi^{-1}\pi$. For T, put the entries of π in a diagonal matrix: $\Pi^{-1}\operatorname{diag}(\pi)\Pi$. Take $\Psi_{2;\mathbf{1}}^\top = AT \operatorname{diag}(\pi) A^\top$ and multiply by $(A\Pi)^\dagger$ on the left and $((A\Pi)^\top)^\dagger(\Pi^{-1}\operatorname{diag}(\pi)\Pi)^{-1}$ on the right, which yields $\Pi^{-1}T\Pi$. (Note that Π is orthogonal, so $\Pi^{-1} = \Pi^\top$.) Finally, multiply $A\Pi = OT\Pi$ and $(\Pi^{-1}T\Pi)^{-1}$, which yields $O\Pi$.

The above algorithm identifies the PCFG-IE from only length 3 sentences. To exploit sentences of different lengths, we can compute a mixing matrix M which includes constraints from sentences of length $1 \le L \le L_{\max}$ up to some upper bound $L_{\max}$. For example, $L_{\max} = 10$ results in a 990×2376 mixing matrix. We can retrieve the same compound parameters ($\Psi_{2;\mathbf{1}}$ and $\Psi_{2;\tau}$) from the pseudoinverse of M and as proceed as before.

5.3 Application to the DEP-IES model

We now turn to the DEP-IES model over $L = 3$ words. Our goal is to recover the parameters $\theta = (\pi, A)$. Let $D = \operatorname{diag}(\pi) = \operatorname{diag}(A\pi)$, where the second equality is due to stationarity of π.

$$\mu_1 \overset{\text{def}}{=} \mathbb{E}[x_1] = \pi,$$

$$\mu_{12} \overset{\text{def}}{=} \mathbb{E}[x_1 \otimes x_2] = 7^{-1}(DA^\top + DA^\top + DA^\top A^\top + AD + ADA^\top + AD + DA^\top),$$

$$\mu_{13} \overset{\text{def}}{=} \mathbb{E}[x_1 \otimes x_3] = 7^{-1}(DA^\top + DA^\top A^\top + DA^\top + ADA^\top + AD + AAD + AD),$$

$$\tilde{\mu}_{12} \overset{\text{def}}{=} \tilde{\mathbb{E}}[x_1 \otimes x_2] = 2^{-1}(DA^\top + AD),$$

where $\tilde{\mathbb{E}}[\cdot]$ is taken with respect to length 2 sentences. Having recovered π from μ_1, it remains to recover A. By selectively combining the moments above, we can compute $AA + A = [7(\mu_{13} - \mu_{12}) + 2\tilde{\mu}_{12}]\operatorname{diag}(\mu_1)^{-1}$. Assuming A is generic position, it is diagonalizable: $A = Q\Lambda Q^{-1}$ for some diagonal matrix $\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_d)$, possibly with complex entries. Therefore, we can recover $\Lambda^2 + \Lambda = Q^{-1}(AA + A)Q$. Since Λ is diagonal, we simply have d independent quadratic equations in λ_i, which can be solved in closed form. After obtaining Λ, we retrieve $A = Q\Lambda Q^{-1}$.

6 Discussion

In this work, we have shed some light on the identifiability of standard generative parsing models using our numerical identifiability checker. Given the ease with which this checker can be applied, we believe it should be a useful tool for analyzing more sophisticated models [6], as well as developing new ones which are expressive yet identifiable.

There is still a large gap between showing identifiability and developing explicit algorithms. We have made some progress on closing it with our unmixing technique, which can deal with models where the tree topology varies non-trivially.

References

[1] A. Anandkumar, D. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden Markov models. In *COLT*, 2012.

[2] F. Pereira and Y. Shabes. Inside-outside reestimation from partially bracketed corpora. In *ACL*, 1992.

[3] G. Carroll and E. Charniak. Two experiments on learning probabilistic dependency grammars from corpora. In *Workshop Notes for Statistically-Based NLP Techniques, AAAI*, pages 1–13, 1992.

[4] M. A. Paskin. Grammatical bigrams. In *NIPS*, 2002.

[5] D. Klein and C. D. Manning. Conditional structure versus conditional estimation in NLP models. In *EMNLP*, 2002.

[6] D. Klein and C. D. Manning. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *ACL*, 2004.

[7] P. Liang and D. Klein. Analyzing the errors of unsupervised learning. In *HLT/ACL*, 2008.

[8] J. T. Chang. Full reconstruction of Markov models on evolutionary trees: Identifiability and consistency. *Mathematical Biosciences*, 137:51–73, 1996.

[9] A. Anandkumar, K. Chaudhuri, D. Hsu, S. M. Kakade, L. Song, and T. Zhang. Spectral methods for learning multivariate latent tree structure. In *NIPS*, 2011.

[10] J. B. Kruskal. Three-way arrays: Rank and uniqueness of trilinear decompositions, with application to arithmetic complexity and statistics. *Linear Algebra and Applications*, 18:95–138, 1977.

[11] E. S. Allman, C. Matias, and J. A. Rhodes. Identifiability of parameters in latent structure models with many observed variables. *Annals of Statistics*, 37:3099–3132, 2009.

[12] E. S. Allman, S. Petrovi, J. A. Rhodes, and S. Sullivant. Identifiability of 2-tree mixtures for group-based models. *Transactions on Computational Biology and Bioinformatics*, 8:710–722, 2011.

[13] J. A. Rhodes and S. Sullivant. Identifiability of large phylogenetic mixture models. *Bulletin of Mathematical Biology*, 74(1):212–231, 2012.

[14] T. J. Rothenberg. Identification in parameteric models. *Econometrica*, 39:577–591, 1971.

[15] L. A. Goodman. Exploratory latent structure analysis using both identifiabile and unidentifiable models. *Biometrika*, 61(2):215–231, 1974.

[16] D. Bamber and J. P. H. van Santen. How many parameters can a model have and still be testable? *Journal of Mathematical Psychology*, 29:443–473, 1985.

[17] D. Geiger, D. Heckerman, H. King, and C. Meek. Stratified exponential families: graphical models and model selection. *Annals of Statistics*, 29:505–529, 2001.

[18] K. Lari and S. J. Young. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer Speech and Language*, 4:35–56, 1990.

[19] M. Johnson, T. Griffiths, and S. Goldwater. Bayesian inference for PCFGs via Markov chain Monte Carlo. In *HLT/NAACL*, 2007.

[20] E. Mossel and S. Roch. Learning nonsingular phylogenies and hidden Markov models. *Annals of Applied Probability*, 16(2):583–614, 2006.

[21] D. Hsu, S. M. Kakade, and T. Zhang. A spectral algorithm for learning hidden Markov models. In *COLT*, 2009.

[22] S. M. Siddiqi, B. Boots, and G. J. Gordon. Reduced-rank hidden Markov models. In *AISTATS*, 2010.

[23] A. Parikh, L. Song, and E. P. Xing. A spectral algorithm for latent tree graphical models. In *ICML*, 2011.

[24] S. B. Cohen, K. Stratos, M. Collins, D. P. Foster, and L. Ungar. Spectral learning of latent-variable PCFGs. In *ACL*, 2012.

[25] F. M. Luque, A. Quattoni, B. Balle, and X. Carreras. Spectral learning for non-deterministic dependency parsing. In *EACL*, 2012.

[26] P. Dhillon, J. Rodue, M. Collins, D. P. Foster, and L. Ungar. Spectral dependency parsing with latent variables. In *EMNLP-CoNLL*, 2012.

[27] J. Eisner. Three new probabilistic models for dependency parsing: An exploration. In *COLING*, 1996.

[28] S. Sahni. Computationally related problems. *SIAM Journal on Computing*, 3:262–279, 1974.

[29] J. Eisner. Bilexical grammars and their cubic-time parsing algorithms. In *Advances in Probabilistic and Other Parsing Technologies*, pages 29–62, 2000.

A   Proof of Theorem 1

Theorem 1 (restated). *Assume Θ is a non-empty open connected subset of $[0,1]^n$ and $\mu\colon \mathbb{R}^n \to \mathbb{R}^m$ is a polynomial map. With probability 1, the following holds.*

- CHECKIDENTIFIABILITY *returns "no" $\Rightarrow$ for almost all $\theta_0 \in \Theta$ and any open neighborhood $N(\theta_0)$ around θ_0, $|\mathcal{S}_\Theta(\theta_0) \cap N(\theta_0)|$ is infinite (not locally identifiable).*

- CHECKIDENTIFIABILITY *returns "yes" $\Rightarrow$ (i) for almost all $\theta_0 \in \Theta$, there exists an open neighborhood $N(\theta_0)$ around θ_0 such that $|\mathcal{S}_\Theta(\theta_0) \cap N(\theta_0)| = 1$ (locally identifiable); and (ii) there exists a set $\mathcal{E} \subset \Theta$ with measure zero such that $|\mathcal{S}_{\Theta \setminus \mathcal{E}}(\theta_0)|$ is finite for every $\theta_0 \in \Theta \setminus \mathcal{E}$ (identifiability of $\Theta \setminus \mathcal{E}$).*

The proof of Theorem 1 crucially relies on the following lemma from [16] which holds even in the case that μ is merely an analytic function (see Lemma 9 of [17] for a simpler proof in the case μ is a polynomial map); it states that the Jacobian achieves its maximal rank almost everywhere in Θ. To state this precisely, first define $r_{\max} \stackrel{\text{def}}{=} \max\{\text{rank}(J(\theta)) : \theta \in \Theta\}$ and $\Theta_{\max} \stackrel{\text{def}}{=} \{\theta \in \Theta : \text{rank}(J(\theta)) = r_{\max}\}$.

Lemma 2. *The set $\Theta \setminus \Theta_{\max}$ has Lebesgue measure zero. That is, $\Theta_{\max}$ is almost all of Θ.*

Proof of Theorem 1. By Lemma 2, CHECKIDENTIFIABILITY chooses a point $\tilde{\theta} \in \Theta_{\max}$ with probability 1. We henceforth condition on this event, so $\text{rank}(J(\tilde{\theta})) = r_{\max}$.

Case 1: $\text{rank}(J(\tilde{\theta})) < n$ (*i.e.*, "no" is returned). In this case, we have $r_{\max} < n$. We now employ an argument from the proof of Proposition 20 of [16]. Fix any $\theta_0 \in \Theta_{\max}$. Since Θ is open, Weyl's theorem implies that there is an open neighborhood U around θ_0 in Θ on which $\text{rank}(J(\theta)) = r_{\max}$ for all $\theta \in U$ (*i.e.*, $\text{rank}(J(\cdot))$ is constant on U). Therefore, by the constant rank theorem, there is an open neighborhood $N(\theta_0)$ around θ_0 in Θ such that $\mu^{-1}(\mu(\theta_0)) \cap N(\theta_0)$ is homeomorphic with an open set in $\mathbb{R}^{n-r_{\max}}$. Therefore $\mathcal{S}_\Theta(\theta_0) \cap N(\theta_0)$ is uncountably infinite.

Case 2: $\text{rank}(J(\tilde{\theta})) = n$ (*i.e.*, "yes" is returned). In this case, we have $r_{\max} = n$. Therefore for every $\theta_0 \in \Theta_{\max}$, the Jacobian $J(\theta_0)$ has full column rank, and thus by the inverse function theorem, μ is injective on a neighborhood of θ_0. This in turn implies that for all $\theta_0 \in \Theta_{\max}$, there exists an open neighborhood $N(\theta_0)$ around θ_0 such that $\mathcal{S}_\Theta(\theta_0) \cap N(\theta_0) = \{\theta_0\}$. This proves (i).

To show (ii), define $\mathcal{E} \stackrel{\text{def}}{=} \Theta \setminus \Theta_{\max}$, and now claim that for every $\theta_0 \in \Theta_{\max}$, the equivalence class $\mathcal{S}_{\Theta_{\max}}(\theta_0)$ is finite. Observe that by (i), the set $\mathcal{S}_{\Theta_{\max}}(\theta_0)$ contains only geometrically isolated solutions to the system of polynomial equations given by $\mu(\theta) = \mu(\theta_0)$. Therefore the claim follows immediately from Bézout's Theorem, which implies that the number of geometrically isolated solutions is finite. $\square$

Remark.   All the models considered in this paper have moments μ which correspond to a polynomial map. However, for some models (e.g., exponential families), μ will not be a polynomial map, but rather, a general analytic function. In this case, Theorem 1 holds with one modification to (ii). If CHECKIDENTIFIABILITY returns "yes", then we have the following weaker guarantee in place of (ii): $\mathcal{S}_{\Theta_{\max}}(\theta_0)$ is *countable* (but not necessarily finite) for all $\theta_0 \in \Theta_{\max}$. The above proof does not require the fact that μ is a polynomial map except in the invocation of Bézout's Theorem. In place of Bézout's Theorem, we use the following argument. If $\mathcal{S}_{\Theta_{\max}}(\theta_0)$ is uncountable, then it contains a limit point $\theta^* \in \mathcal{S}_{\Theta_{\max}}(\theta_0)$; thus for any small enough neighborhood $N(\theta^*)$ of θ^*, there is some $\theta \in \mathcal{S}_{\Theta_{\max}}(\theta_0) \cap N(\theta^*)$. This contradicts (i) as applied to θ^*, and thus we conclude that $\mathcal{S}_{\Theta_{\max}}(\theta_0)$ is countable.

B   Additional results from the identifiability checker

PCFG models with $d < k$.   The PCFG models that we've considered so far assume that the number of words d is at least the number of hidden states k, which is a realistic assumption for natural language. However, there are applications, e.g., computational biology, where the vocabulary size d is relatively small. In this regime, identifiability becomes trickier because the data doesn't

reveal as much about the hidden states, and brings us closer to the boundary between identifiability and non-identifiability. In this section, we consider the $d < k$ regime.

The following table gives additional identifiability results from CHECKIDENTIFIABILITY for values of d, k, and L where $d < k$ (recall that the results reported in Section 4.4 only considered values where $d \geq k$). In each cell, we show the (k, d, L) values for which CHECKIDENTIFIABILITY returned "yes"; the values checked were $k \in \{3, 4, \ldots, 8\}$, $d \in \{2, \ldots, k-1\}$, $L \in \{3, 4, \ldots, 9\}$.

	ϕ_{12}	ϕ_{**}	ϕ_{123e_1}	ϕ_{123}	ϕ_{***e_1}	ϕ_{***}
PCFG				None		
PCFG-I	None	$(3,2,\geq 6)$ $(4,2,\geq 8)$ $(4,3,\geq 5)$ $(5,3,\geq 6)$ $(5,4,\geq 4)$ $(6,3,\geq 7)$ $(6,4,\geq 5)$ $(6,5,\geq 4)$ $(7,3,\geq 8)$ $(7,4,\geq 6)$ $(7,5,\geq 5)$ $(7,6,\geq 4)$	None	$(5,4,\geq 4)$ $(6,5,\geq 4)$ $(7,5,\geq 4)$ $(7,6,\geq 4)$	$(3,2,\geq 5)$ $(4,2,\geq 6)$ $(4,3,\geq 4)$ $(5,2,\geq 7)$ $(5,\geq 3,\geq 4)$ $(6,2,\geq 8)$ $(6,3,\geq 5)$ $(6,\geq 4,\geq 4)$ $(7,2,\geq 9)$ $(7,3,\geq 5)$ $(7,\geq 4,\geq 4)$	
PCFG-IE	None	$(3,2,\geq 6)$ $(4,2,\geq 8)$ $(4,3,\geq 5)$ $(5,3,\geq 6)$ $(5,4,\geq 5)$ $(6,3,\geq 7)$ $(6,4,\geq 5)$ $(6,5,\geq 4)$ $(7,3,\geq 8)$ $(7,4,\geq 6)$ $(7,5,\geq 5)$ $(7,6,\geq 4)$	$(5,4,\geq 4)$ $(6,5,\geq 4)$ $(7,5,\geq 5)$ $(7,6,\geq 4)$	$(4,3,\geq 4)$ $(5,4,\geq 4)$ $(6,\geq 4,\geq 4)$ $(7,\geq 5,\geq 4)$	$(3,2,\geq 5)$ $(4,2,\geq 6)$ $(4,3,\geq 4)$ $(5,2,\geq 7)$ $(5,3,\geq 5)$ $(5,4,\geq 4)$ $(6,2,\geq 8)$ $(6,3,\geq 5)$ $(6,\geq 4,\geq 4)$ $(7,2,\geq 9)$ $(7,3,\geq 5)$ $(7,\geq 4,\geq 4)$	$(3,2,\geq 5)$ $(4,2,\geq 6)$ $(4,3,\geq 4)$ $(5,2,\geq 7)$ $(5,\geq 3,\geq 4)$ $(6,2,\geq 8)$ $(6,3,\geq 5)$ $(6,\geq 4,\geq 4)$ $(7,2,\geq 9)$ $(7,3,\geq 5)$ $(7,\geq 4,\geq 4)$

Fixed topology models. We now present some results for latent class models (LCMs) and hidden Markov models (HMMs). While identifiability for these models are more developed than for parsing models, we show that the identifiability checker can refine the results even for the classic models.

The parameters of an HMM are $\theta = (\pi, T, O)$, where $\pi \in \mathbb{R}^k$ specifies the initial state distribution, $T \in \mathbb{R}^{k \times k}$ specifies the state transition probabilities, and $O \in \mathbb{R}^{d \times k}$ specifies the emission distributions. The probability over a sentence $\mathbf{x}$ is:

$$\mathbb{P}_\theta(\mathbf{x}) = \mathbf{1}^\top T \operatorname{diag}(O^\top x_L) \cdots T \operatorname{diag}(O^\top x_2) T \operatorname{diag}(O^\top x_1) \pi. \tag{5}$$

The parameters of an LCM are $\theta = (\pi, O)$—the same as that of an HMM except with $T \equiv I$. The probability over a sentence $\mathbf{x}$ is also given by (5) (with $T = I$).

The following table summarizes some identifiability results obtained by CHECKIDENTIFIABILITY (for $d \geq k$); these results have all been proven analytically in previous work (e.g., [8, 10, 11, 20, 21]) except for the identifiability of HMMs from ϕ_{**}.

	ϕ_{12}	ϕ_{**}	ϕ_{123e_1}	ϕ_{123}	ϕ_{***e_1}	ϕ_{***}
LCM		No			Yes iff $L \geq 3$	
HMM	No			Yes iff $L \geq 3$		

It is known that LCMs are not identifiable from ϕ_{**} for any value of L [8]. However, LCMs constitute a subfamily of HMMs arising from a measure zero subset of the HMM parameter space. Therefore the identifiability of HMMs from ϕ_{**} (for $L \geq 3$) does not contradict this result. The result does not appear to be covered by application of Kruskal's theorem in previous work [11], so we prove the result rigorously below.

11

It can be checked using (5) that

$$\mathbb{E}_\theta[\phi_{12}(\mathbf{x})] = O \operatorname{diag}(\pi) T^\top O^\top$$
$$\mathbb{E}_\theta[\phi_{34}(\mathbf{x})] = O \operatorname{diag}(T\pi) T^\top O^\top.$$

Let $M_1 \stackrel{\text{def}}{=} O$, $M_2 \stackrel{\text{def}}{=} OT \operatorname{diag}(\pi)$, and $D \stackrel{\text{def}}{=} \operatorname{diag}(T\pi) \operatorname{diag}(\pi)^{-1}$. Provided that

1. $\pi > 0$,
2. O has full column rank,
3. T is invertible,
4. the ratios of probabilities $(T\pi)_i / \pi_i$, ranging over $i \in [k]$, are distinct

(all of which are true for all but a measure zero set of parameters in Θ), the matrices M_1 and M_2 have full column rank and the diagonal matrix D has distinct diagonal entries. Therefore Lemma 1 can be applied with $X = \mathbb{E}_\theta[\phi_{12}(\mathbf{x})] = M_1 M_2^\top$ and $Y = \mathbb{E}_\theta[\phi_{34}(\mathbf{x})] = M_1 D M_2^\top$ to recover $M_1 = O$. It is easy to see that π and T can also easily be recovered.

Note that the fourth condition above, that $T\pi$ be entry-wise distinct from π, is violated when a LCM distribution is cast as an HMM distribution (by setting $T = I$ so $T\pi = \pi$). However, the set of HMM parameters satisfying this equation is a measure zero set.

Discussion. CHECKIDENTIFIABILITY tests for local identifiability. If it finds that a model family is not locally identifiable, then it is not globally identifiable. However the inverse claim is not necessarily true: if it finds that a model family is locally identifiable, it is not necessarily globally identifiable. Theorem 1 provides the somewhat weaker guarantee that a restricted model family is globally identifiable, where the equivalence classes $\mathcal{S}_{\Theta \setminus \mathcal{E}}(\theta_0)$ are only taken with respect to a subset $\Theta \setminus \mathcal{E} \subseteq \Theta$ of the parameter space. However, there is a gap between this property (which is with respect to $\Theta \setminus \mathcal{E}$) and true global identifiability (which is with respect to Θ).

On the other hand, having explicit estimators guarantees us proper global identifiability with respect to the original model family Θ. In fact, the exceptional set $\mathcal{E}$ can typically be characterized explicitly. For instance, in the case of PCFG-IE, the set $\Theta \setminus \mathcal{E}$ contains those $\theta = (\pi, T, O)$ that satisfy full rank conditions:

$$\Theta \setminus \mathcal{E} = \{(\pi, T, O) : \pi \succ 0, T \text{ is invertible}, O \text{ has full column rank}\}. \tag{6}$$

Additionally, the explicit estimators also provides an explicit characterization of the elements in the equivalence class $\mathcal{S}_\Theta(\theta_0)$ for each $\theta_0 \in \Theta \setminus \mathcal{E}$: the set $\mathcal{S}_\Theta(\theta_0)$ contains exactly $k!$ elements corresponding to permutation of the hidden states. Specifically,

$$\mathcal{S}_\Theta((\pi, T, O)) = \{(\Pi^{-1}\pi, \Pi^{-1}T\Pi, O\Pi) : \Pi \text{ is a permutation matrix}. \tag{7}$$

Note that this is shaper than Theorem 1, which only says that the equivalence classes have to be finite.

C Dynamic programs

For a sentence of length L, the number of parse trees is exponential in L. Therefore, dynamic programming is often employed to efficiently compute expectations over the parse trees, the core computation in the E-step of the EM algorithm. In the case of PCFG, this dynamic program is referred to as the CKY algorithm, which runs in $O(L^3 k^3)$ time, where k is the number of hidden states. For simple dependency models, a $O(L^3)$ dynamic program was developed by [29]. At a high-level, the states of the dynamic program in both cases are the spans $[i : j]$ of the sentence (and for the PCFG, the these states include the hidden states $z_{[i:j]}$ of the nodes).

In this paper, we need to compute (i) the Jacobian matrix for checking identifiability (Section 4.2) and (ii) the mixing matrix for recovering compound parameters (Section 5.1). Both computations can be performed efficiently with a modified version of the classic dynamic programs, which we will describe in this section.

C.1 Computing the Jacobian matrix

Recall that the j-th row of the Jacobian matrix J is (the transpose of) the gradient of $h_j(\theta) = \mu_j(\theta) - \mu_j(\theta_0)$. Specifically, entry J_{ji} is the derivative of the j-th moment with respect to the i-th parameter:

$$J_{ji} = \frac{\partial h_j(\theta)}{\partial \theta_i} \tag{8}$$

$$= \frac{\partial \mathbb{E}_\theta[\phi_j(\mathbf{x})]}{\partial \theta_i} \tag{9}$$

$$= \sum_{\mathbf{x},z} \frac{\partial p_\theta(\mathbf{x}, z)}{\partial \theta_i} \phi_j(\mathbf{x}). \tag{10}$$

We can encode the sum over the exponential set of possible sentences $\mathbf{x}$ and parse trees z using a directed acyclic hypergraph so that each hyperpath through the hypergraph corresponds to a $(\mathbf{x}, z)$ pair. Specifically, a *hypergraph* consists of the following:

- a set of nodes $\mathcal{V}$ with a designated start node $\text{START} \in \mathcal{V}$ and an end node $\text{END} \in \mathcal{V}$, and
- a set of hyperedges $\mathcal{E}$ where each hyperedge $e \in \mathcal{E}$ has a source node $e.a \in \mathcal{V}$ and a pair of target nodes $(e.b, e.c) \in \mathcal{V} \times \mathcal{V}$ (we say that e connects $e.a$ to $e.b$ and $e.c$) and an index $e.i \in [n]$ corresponding to a component of the parameter vector $\theta \in \mathbb{R}^n$.

Define a *hyperpath P* to be a subset of the edges $\mathcal{E}$ such that:

- $(\text{START}, a, b) \in P$ for some $a, b \in \mathcal{V}$;
- if $(a, b, c) \in P$ and $b \neq \text{END}$, then $(b, d, e) \in P$ for some $d, e \in \mathcal{V}$; and
- if $(a, b, c) \in P$ and $c \neq \text{END}$, then $(c, d, e) \in P$ for some $d, e \in \mathcal{V}$.

Each hyperpath P, encoding $(\mathbf{x}, z)$, is associated with a probability equal to the product of all of the parameters on that hyperpath:

$$p_\theta(\mathbf{x}, z) = p_\theta(P) = \prod_{e \in P} \theta_{e.i}. \tag{11}$$

In this way, the hypergraph compactly defines a distribution over exponentially many hyperpaths.

Now, we assume that each moment $\phi_j(\mathbf{x})$ corresponds to a function $f_j : \mathcal{E} \mapsto \mathbb{R}$ mapping each hyperedge e to a real number so that the moment is equal to the product over function values:

$$\phi_j(\mathbf{x}) = \prod_{e \in P} f_j(e), \tag{12}$$

where P is any hyperpath that encodes the sentence $\mathbf{x}$ and some parse tree z (we assume that the product is the same no matter what z is).

Now, let us write out the Jacobian matrix entries in terms of hyperpaths:

$$J_{ji} = \sum_P \sum_{e_0 \in P} \frac{\partial \theta_{e_0.i}}{\partial \theta_i} \prod_{e \in P, e \neq e_0} \theta_{e.i} f_j(e). \tag{13}$$

The sum over hyperpaths P can be computed efficiently as follows. For each hypergraph node a, we compute an inside score $\alpha(a)$, which sums over all possible partial hyperpaths terminating at the target node, and an outside score $\beta(a)$, which sums over all possible partial hyperpaths from the source node:

$$\alpha(a) \stackrel{\text{def}}{=} \sum_{e \in \mathcal{E}: e.a = a} \theta_{e.i} \alpha(e.b) \alpha(e.c), \tag{14}$$

$$\beta(a) \stackrel{\text{def}}{=} \sum_{e \in \mathcal{E}: e.b = a} \theta_{e.i} \alpha(e.c) \beta(e.a) \sum_{e \in \mathcal{E}: e.c = a} \theta_{e.i} \alpha(e.b) \beta(e.a). \tag{15}$$

The Jacobian entry J_{ji} can be computed as follows:

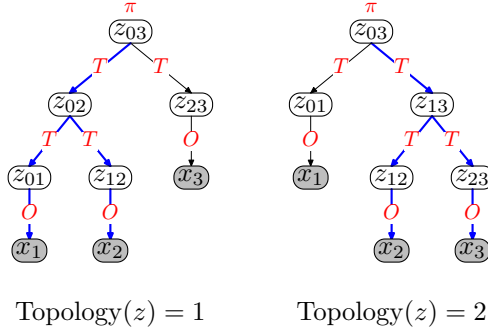$$J_{ji} = \sum_{e \in \mathcal{E}} \beta(e.a) \alpha(e.b) \alpha(e.c) \mathbb{I}[i = e.i]. \tag{16}$$

Topology$(z) = 1$ Topology$(z) = 2$

Figure 3: An example of a backbone structure in blue corresponding to the compound parameter $OT \operatorname{diag}(T\pi)T^\top O^\top$, which appears in two different topologies, for two observation matrices, ϕ_{12} and ϕ_{23}, respectively.

Example: PCFG. For a PCFG, nodes $\mathcal{V}$ have the form $(i, j, s) \in [L] \times [L] \times [k]$, corresponding to a hidden state s over span $[i : j]$. For each hidden state s, we have a hyperedge e connecting $e.a =$ START to $e.b = (s, 0, L)$ and $e.c =$ END; this hyperedge has parameter index $e.i$ corresponding to π_s. For each span $[i : j]$ with $j - i > 1$, split point $i < m < j$, and hidden states $s_1, s_2, s_3 \in [k]$, $\mathcal{E}$ contains a hyperedge e connecting $e.a = (i, j, s_1)$ to $e.b = (i, m, s_2)$ and $e.c = (m, j, s_3)$; the parameter index $e.i$ corresponds to the binary production $B_{(s_2 \otimes_k s_3)s_1}$. For each span $[i - 1 : i]$, hidden state $s \in [k]$ and word $x \in [d]$, we have a hyperedge e connecting $e.a = (i - 1, i, s)$ to $e.b =$ END and $e.c =$ END with parameter index $e.i$ corresponding to the emission O_{xs}.

The moments can be encoded as follows: For example, if $\phi_j(\mathbf{x}) = \mathbb{I}[x_i = t]$, then we define $f_j(e)$ to be 0 if the source node corresponds to position i ($e.a = (i - 1, i, s)$) and the parameter index $e.i$ does not correspond to O_{ts} for some $s \in [k]$, and 1 otherwise. In this way, $\prod_{e \in P} f_j(e)$ is zero if P encodes a sentence with $x_i \neq t$.

Higher-order moments simply correspond to hyperedge-wise multiplication of these first-order moments. For example, if $\phi_{j_1}(\mathbf{x}) = \mathbb{I}[x_{i_1} = t_1]$ and $\phi_{j_2}(\mathbf{x}) = \mathbb{I}[x_{i_2} = t_2]$, then the second-order moment $\phi_j(\mathbf{x}) = \mathbb{I}[x_{i_1} = t_1, x_{i_2} = t_2]$ corresponds to $f_j(e) = f_{j_1}(e)f_{j_2}(e)$.

C.2 Computing the mixing matrix

Recall that the mixing matrix M includes a row for each observation matrix $o \in \mathcal{O}$ and a column for each compound parameter $p \in \mathcal{P}$. Assuming a uniform distribution over topologies, computing each entry of M reduces to counting the number of topologies t consistent with a particular compound parameter Ψ_p:

$$M_{op} = \mathbf{P}(\Psi_{o,\text{Topology}} = \Psi_p) \tag{17}$$

$$= |\text{Topologies}|^{-1} \sum_t \mathbb{I}[\Psi_{o,t} = \Psi_p]. \tag{18}$$

First, we will characterize the set of compound parameters graphically in terms of backbone structures. As an example, consider the PCFG-IE model and the observation matrix ϕ_{12} ($o = 12$) corresponding to the marginal distribution over the first two words of the sentence. Given a topology t, consider starting at the root, descending to the lowest common ancestor of x_1 and x_2, and then following both paths down to x_1 and x_2, respectively. We refer to this traversal as the *backbone structure* with respect to topology t and observation matrix ϕ_{12}. See Figure 3 for an example of the backbone structure, outlined in blue.

Note that the compound parameter $\Psi_{12,t}(\theta) = \mathbb{E}_\theta[\phi_{12}(\mathbf{x}) \mid \text{Topology} = t]$ can be written as a product over the parameter matrices, one for each edge of the backbone structure. For Figure 3, this would yield

$$\Psi_{12,1}(\theta) = OT \operatorname{diag}(T\pi)T^\top O^\top. \tag{19}$$

For general trees, we would have

$$\Psi_{12,t}(\theta) = OT^{n_1} \operatorname{diag}(T^{n_3}\pi)(T^\top)^{n_2} O^\top. \tag{20}$$

for some positive integers n_1, n_2, n_3 corresponding to the number of edges (in t) from the common node to the preterminal node z_{01}, the preterminal node z_{12}, and the root z_{0L}, respectively.

Note that the compound parameter does not depend on the structure of t outside the backbone—that part of the topology is effectively marginalized out—so the compound parameter $\Psi_{12,t}(\theta)$ will be identical for all topologies sharing that same backbone structure. Therefore, there are only a polynomial number of compound parameters despite an exponential number of topologies t.

We define a dynamic program that recursively computes M_{op} for the PCFG-IE model under a fixed second-order observation matrix $\phi_{i_0 j_0}$. Specifically, for each span $[i:j]$ define $H(i,j)$ to be the set of pairs $\langle t,n \rangle$ where t is a partial backbone structure t and n is the number of partial topologies over span $[i:j]$ which are consistent with t.

In the base case $H(i-1,i)$, if i is either of the designated leaf positions defined by the observation matrix (i_0 or j_0), then we return the single-node backbone structure $\bullet$; otherwise, we return the null backbone structure $\emptyset$:

$$H(i-1,i) = \begin{cases} \{\langle \emptyset, 1 \rangle\} & \text{if } i = i_0 \text{ or } i = j_0 \\ \{\langle \bullet, 1 \rangle\} & \text{otherwise.} \end{cases} \tag{21}$$

In the recursive case $H(i,j)$, we consider all split points m, partial backbones t_1 and t_2 from $H(i,m)$ and $H(m,j)$, respectively, and create a new tree with t_1 and/or t_2 as the subtrees if they are not null:

$$H(i,j) = \bigcup_{i<m<j}^{+} \bigcup_{\langle t_1, n_1 \rangle \in H(i,m)}^{+} \bigcup_{\langle t_2, n_2 \rangle \in H(m,j)}^{+} \langle \textsc{Combine}(t_1, t_2), n_1 n_2 \rangle, \tag{22}$$

$$\textsc{Combine}(t_1, t_2) = \begin{cases} (T:t_1, T:t_2) & \text{if } t_1 \neq \emptyset \text{ and } t_2 \neq \emptyset, \\ T:t_1 & \text{if } t_1 \neq \emptyset, \\ T:t_2 & \text{if } t_2 \neq \emptyset, \\ \emptyset & \text{otherwise.} \end{cases} \tag{23}$$

Here, we use the notation $\bigcup^{+}$ to denote a multi-set union: $\{\langle t, n_1 \rangle\} \cup^{+} \{\langle t, n_2 \rangle\} = \{\langle t, n_1 + n_2 \rangle\}$. In this notation, the backbone structure in Figure 3 would be represented as $T:(T:O:\bullet, T:O:\bullet)$, which can be easily converted to the compound parameter $OT \operatorname{diag}(T\pi)T^\top O^\top$.

For third-order observation matrices (e.g., $\phi_{i_0 j_0 k_0 \eta}$), we add an additional case to $H(i-1,i)$ to return $\langle \circ, 1 \rangle$ if $i = k_0$; note that k_0 is represented by a special node $\circ$ because that observation is projected using η. The first case of $\textsc{Combine}(t_1, t_2)$ undergoes one change: if t_2 is a chain ending in $\circ$, then we return $(T:t_2, T:t_1)$. The reason for this is best demonstrated by an example: consider topology 1 in Figure 3, and the two observation matrices $\phi_{132\eta}$ and $\phi_{231\eta}$. Without the reordering, we would have the backbone structure: $(T:(T:\bullet, T:\circ), T:\bullet)$ and $(T:(T:\circ, T:\bullet), T:\bullet)$. However, they have the same compound parameter $OT \operatorname{diag}(T^\top O^\top \eta)T^\top \operatorname{diag}(\pi)TO$. This is because the contribution of a subtree ending in $\circ$ is simply a diagonal matrix ($\operatorname{diag}(T^\top O^\top \eta)$ in this case) which is applied on the hidden state regardless of whether it came from the left or right side.

One might also see why the unmixing technique does not directly apply to the PCFG-I model, where T is replaced with T_1 for left edges and T_2 for right edges. In that case, there are many backbone structures (and thus more compound parameters) due to the different interleavings of left and right edges.