

Tractable Bayesian Learning of Tree Belief Networks

Marina Meilä
Carnegie Mellon University
mmp@cs.cmu.edu

Tommi Jaakkola
Massachusetts Institute of Technology
tommi@ai.mit.edu

May, 2000
CMU-RI-TR-00-15

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This work was funded by ONR contract number N00014-98-C-0326 and by NSF KDI award number DMS-9873442.

Keywords: graphical models, spanning tree, Bayesian learning

Abstract

In this paper we present *decomposable priors*, a family of priors over structure and parameters of tree belief nets for which Bayesian learning with complete observations is tractable, in the sense that the posterior is also decomposable and can be completely determined analytically in polynomial time. This follows from two main results: First, we show that factored distributions over spanning trees in a graph can be integrated in closed form. Second, we examine priors over tree parameters and show that a set of assumptions similar to (Heckerman and al., 1995) constrain the tree parameter priors to be a compactly parametrized product of Dirichlet distributions. Besides allowing for exact Bayesian learning, these results permit us to formulate a new class of tractable latent variable models in which the likelihood of a data point is computed through an ensemble average over tree structures.

1 Introduction

In the framework of graphical models, tree distributions stand out by their special computational advantages. Inference and sampling from a tree are linear in the number of variables n . While it is known that for many classes of graphical models, as for example junction trees with cliquewidth > 2 , the problem of learning the optimal structure is NP-hard, for trees this problem is solvable in only quadratic time. The latter result is due to [1] who present an algorithm for finding the structure and parameters of the tree that best fits a given distribution in the Maximum Likelihood (ML) framework. This algorithm was generalized to Maximum A-Posteriori (MAP) learning [7, 5].

In this paper we present another remarkable property of tree graphical models: the fact that Bayesian learning for a certain class of priors, called *decomposable*¹ priors, is also tractable. Essentially, decomposable priors are priors that can be represented as a product of factors corresponding to the edges of the tree. We show that if the prior is decomposable and we have a data set consisting of N complete i.i.d. observations, then the posterior distribution over all tree structures and parameters is also decomposable, is expressible with a quadratic number of parameters that can be computed exactly from data in $\mathcal{O}(n^3 + n^2N)$ operations. Evaluating the posterior for a given tree takes then $\mathcal{O}(n)$ time. The first two results come from the fact that, with the standard assumptions of likelihood equivalence, parameter independence and parameter modularity, the prior for tree parameters is constrained to be a product of Dirichlet distributions whose parameters satisfy a set of consistency relations. The last result, i.e. the possibility of computing the posterior exactly, is consequence of the fact that a factored distribution over tree structures can be integrated exactly, using a theorem from combinatorics called *the Matrix tree theorem*.

The paper starts by defining tree distributions and the problem of Bayesian learning in section 2; it presents decomposable priors over tree structures and parameters in sections 3 and 4; the pieces of the puzzle are put together in section 5 where Bayesian learning is described; the previous results are generalized to the case of several sets of independent variables in section 6; the next section, 7 exploits a different set of possibilities opened by our tractability results: it defines a new model, *ensembles of trees*, and shows that it can be learned by gradient ascent in the ML framework; section 8

¹The term decomposable prior will refer here to a prior over a family of graphical models. It should not be confused with a *decomposable model* which is a distribution over V .

contains the final remarks.

2 Tree distributions and the Bayesian learning problem

In this section we introduce the tree model and the notation that will be used throughout the paper. Let $V = \{1, \dots, n\}$ denote the set of variables of interest. Let r_v be the number of values of variable $v \in V$, $r_{MAX} = \max r_v$, x_v a particular value of v , and x an assignment to all the variables in V .

According to the graphical model paradigm, each variable is viewed as a vertex of a graph. We shall call a graph that has no cycles a *tree*² and shall denote by E its edge set. If the tree is connected, e.g. it *spans* all the nodes in V , it is called a *spanning tree*.

Now we define a probability distribution T that is *factored according to a tree*. Let us denote by T_{uv} and T_v the marginals of T :

$$\begin{aligned} T_{uv}(x_u, x_v) &= \sum_{x: u=x_u, v=x_v} T(x) \\ T_v(x_v) &= \sum_{x: v=x_v} T(x). \end{aligned}$$

Let $\deg v$ be the *degree* of vertex v , e.g. the number of edges incident to $v \in V$. Then, the distribution T is factored according to the tree (V, E) if it can be represented as:

$$T(x) = \frac{\prod_{uv \in E} T_{uv}(x_u, x_v)}{\prod_{v \in V} T_v(x_v)^{\deg v - 1}} \quad (1)$$

The distribution itself will be called a tree when no confusion is possible. An equivalent representation for T in terms of conditional probabilities is

$$T(x) = \prod_{v \in V} T_{v|\text{pa}(v)}(x_v | x_{\text{pa}(v)}) \quad (2)$$

where $\text{pa}(v)$ represents the parent of v in the thus directed tree or the empty set if v is the root of a connected component. The form (2) can be obtained from (1) by choosing an arbitrary root in each connected component and recursively substituting $\frac{T_{v,\text{pa}(v)}}{T_{\text{pa}(v)}}$ by $T_{v|\text{pa}(v)}$ starting from the root. We denote

²In the graph theory literature, our definition corresponds to a *forest*. The connected components of a forest are called trees.

such a *directed* tree structure by \overline{E} . The directed tree representation has the advantage of having independent parameters. The total number of free parameters in either representation for a tree with p connected components is

$$\sum_{uv \in E} r_u r_v - \sum_{v \in V} (\deg v - 1) r_v - p$$

In the forthcoming we shall use both representations. Which representation we consider will be clear from the context in all cases of relevance.

We now turn to the problem of learning trees in the Bayesian framework. In this framework, one assumes a prior $P_0(T)$ over the set \mathcal{T}_V of all tree distributions defined on the domain V . Learning from a dataset of complete and independently generated observations $\mathcal{D} = \{x^1, x^2, \dots, x^N\}$ means finding the *posterior* distribution $P(T|\mathcal{D})$ over the set of models \mathcal{T}_V . The solution to this problem is given by the well known Bayes' formula

$$P(T|\mathcal{D}) \propto P_0(T) \prod_{t=1}^N T(x^t) \quad (3)$$

Practically however, Bayesian learning poses a number of significant challenges. First, one needs to define a distribution over the space of all models to play the role of the prior. Such a distribution is composed of a discrete distribution over the set of tree structures $P_0(E)$ and a probability density over the continuous set of tree parameters $P_0(\theta|E)$. Here θ consists of all the parameters of T in some representation.

$$P_0(T) = P_0(E)P_0(\theta|E) \quad (4)$$

The discrete space of all spanning tree structures over V has a super-exponential number of trees (n^{n-2}) which makes defining a distribution over such a space a non-trivial task. Moreover, the second factor in the above formula requires us to define a prior distribution for the tree parameters for each possible structure E . Thus, the first practical requirement is to have a tractable representation for the prior. Even with a tractable representation, the explicit computation of the posterior $P(T|\mathcal{D})$ is usually intractable due to the difficulty of computing the normalization constant in (3). Therefore common practices in Bayesian learning are Maximum A-Posteriori (MAP) estimation and approximations of the posterior around its peaks. An exception from this situation are the so-called *conjugate priors*. If a given (graphical) model has a family of conjugate priors \mathcal{P} then for $P_0 \in \mathcal{P}$ the

posterior is also in \mathcal{P} . The property of having conjugate priors is characteristic of the exponential family of distributions [2]. In this paper we set out to find the conjugate prior for the family of spanning tree models \mathcal{T}_V .

According to (4), to define a prior over \mathcal{T}_V one needs to define a prior over spanning tree structures and a prior for tree parameters, given the structure. While it is not hard to see that for a fixed structure E a tree distribution over discrete variables is an exponential model and thus has conjugate priors, realizing the same fact when E also varies is by far less obvious and constitutes the main contribution of this paper. In the next section we establish the core theorem that allows us to do so.

3 Decomposable priors over tree structures

A *decomposable* distribution P over spanning tree structures E depends on a set of parameters $\beta_{uv} = \beta_{vu} \geq 0$; $\beta_{vv} = 0$, $u, v \in V$ by

$$P(E) = \frac{1}{Z} \prod_{uv \in E} \beta_{uv}. \quad (5)$$

In the above, Z is the normalization constant

$$Z = \sum_E \prod_{uv \in E} \beta_{uv}. \quad (6)$$

Note that in the distribution (6), each parameter β_{uv} can be interpreted as the weight of edge uv , and the probability of a structure E is the product of the weights of all edges in E . Although this distribution is expressible in a product form, it does not imply that the edges' occurrences in E are independent, since the set E as a whole is constrained to be a tree structure.

This prior is simple and compactly parametrized, but to be completely defined one needs to evaluate the normalization constant Z . Using formula (6) is intractable, but the following theorems develop a practical and exact method for doing so.

Theorem 1 (Matrix Tree Theorem) [9] *Let $G = (V, E)$ be a multigraph and denote by $a_{uv} = a_{vu}$ the number of undirected edges between vertices u and v . Then the number of all spanning trees of G is given by $|A_{uv}|(-1)^{u+v}$ the value of the determinant obtained from the following matrix by removing*

row u and column v ³.

$$A = \begin{bmatrix} \deg v_1 & -a_{12} & -a_{13} & \dots & -a_{1,n} \\ -a_{21} & \deg v_2 & -a_{23} & \dots & -a_{2,n} \\ \dots & \dots & \dots & \dots & \dots \\ -a_{n,1} & -a_{n,2} & -a_{n,3} & \dots & \deg v_n \end{bmatrix} \quad (7)$$

In the following, we shall use the simplifying notation below to refer to a set of real values each corresponding to a pair of variables in V

$$a = \{a_{uv}, u, v \in V, u \neq v\} \quad (8)$$

In addition, $a \geq 0$ will mean that $a_{uv} \geq 0$, $a_{uv} \in a$ and ab will denote $\{a_{uv}b_{uv}, u, v \in V, u \neq v\}$ for a, b defined as above. By extending the Matrix Tree theorem to continuous valued A and letting the weights β play the role of a in (7), one can prove

Theorem 2 [6, 4] *Let $P(E)$ be a distribution over spanning tree structures defined by (5,6). Then the normalization constant Z is equal to $|Q(\beta)|$ with $Q(\beta)$ being the first $(n-1)$ lines and columns of the matrix $\overline{Q}(\beta)$ given by:*

$$\overline{Q}_{uv}(\beta) = \overline{Q}_{vu}(\beta) = \begin{cases} -\beta_{uv} & 1 \leq u < v \leq n \\ \sum_{v'=1}^n \beta_{v'v} & 1 \leq u = v \leq n \end{cases} \quad (9)$$

This shows that summing over the distribution of all spanning trees, when this distribution factors according to the trees' edges, can be done in closed form by computing the value of a order $n-1$ determinant, operation that involves $\mathcal{O}(n^3)$ operations. The proof of theorem 2 as well as the other proofs appear in the appendices.

In the following it will be useful to think of $\overline{Q}(\beta)$ and $Q(\beta)$ as functions mapping a set of parameters β each corresponding to a pair of variables in V to a matrix in the ways described by theorem 2.

The support graph. The factored form of the decomposable distribution makes it easy to test whether a given structure has non-zero probability. If all the β parameters are strictly positive, then every tree structure is possible. Otherwise, the structures that will never appear are the structures containing one or more zero-weight edges. We denote by E^{sup} the set of

³Note that A as a whole is a singular matrix

edges uv for which $\beta_{uv} > 0$. The graph $G^{sup} = (V, E^{sup})$ is called *the support graph* of $P(E)$. If enough edges have zero weights, then G^{sup} may be disconnected. In the following we shall assume that the support graph is connected, leaving the discussion of the general case for section 6.

In the remainder of this section we develop a number of consequences of theorem 2.

Computing averages under a decomposable distribution A decomposable distribution is a (curved) [8] exponential model and $\ln Z$ represents its *cumulant generating function* or *partition function*; many quantities of interest, like averages under $P(E)$ can be expressed as derivatives of the partition function. The next series of results exemplifies these possibilities. We assume that G^{sup} is connected.

Lemma 3 [6] *Let Z be given by equation (6) with $\beta \geq 0$, $Q(\beta)$ be given by theorem 2, Q^{-1} be the inverse of Q and $M(\beta)$ be a symmetric matrix with 0 diagonal defined by*

$$\begin{aligned} M_{uv} &= (Q^{-1})_{uu} + (Q^{-1})_{vv} - 2(Q^{-1})_{uv}, \quad u, v < n \\ M_{nv} &= M_{vn} = (Q^{-1})_{vv}, \quad v < n \\ M_{vv} &= 0 \end{aligned} \tag{10}$$

Then the partial derivative of Z with respect to β_{uv} is

$$\frac{\partial Z}{\partial \beta_{uv}} = M_{uv}(\beta) |Q(\beta)|. \tag{11}$$

We shall denote by $\langle f \rangle_P$ the average of a function f under distribution P . The following lemma states a useful fact about averages of additive functions. An *additive* function $f(E)$ satisfies

$$f(E) = \sum_{uv \in E} f_{uv} \tag{12}$$

Lemma 4 [6] *Let $P(E)$, Q and M be given by (5), theorem 2 and (10) respectively and f be an additive function of the structure E . Then the average of f under P is*

$$\langle f(E) \rangle_P = \sum_E f(E) P(E) \tag{13}$$

$$= \sum_{u < v} f_{uv} \beta_{uv} M_{uv}(\beta) \tag{14}$$

$$= \text{trace } Q(f\beta) Q^{-1}(\beta) \tag{15}$$

In (15), f is an overloaded notation representing the set $\{f_{uv}, u, v \in V\}$ in the sense of (8). A similar but more obvious result holds for functions $g(E)$ that are *multiplicative*, i.e. $g(E) = \prod_{uv \in E} g_{uv}$. For such functions we obtain

$$\langle g(E) \rangle_P = \frac{|Q(\beta g)|}{|Q(\beta)|} \quad (16)$$

4 Decomposable priors over tree parameters

Now we examine priors over tree parameters, with the goal of finding conditions under which the priors can be tractably represented. The assumptions we make are similar to those of Heckerman, Geiger and Chickering in [5] (called HGC in the forthcoming) and so will be some of the results. In addition, we will show that in the case of trees these assumptions are also sufficient for tractable representation and learning.

In the following, without loss of generality, we will consider that both the directed and the undirected tree representation are in the *probability table* parametrizations, and we denote respectively

$$\theta_v(j) = T_v(j) \quad (17)$$

$$\theta_{uv}(ij) = T_{uv}(ij) \quad (18)$$

$$\theta_E = \{\theta_{uv}(ij), uv \in E, i = 1, \dots, r_u, j = 1, \dots, r_v\} \quad (19)$$

$$\theta_{\overline{E}} = \{\theta_{u|v}(ij), \overline{vu} \in \overline{E}, i = 1, \dots, r_u, j = 1, \dots, r_v\}$$

We will assume by convention that if $v \in V$ has no parent then $\theta_{v|\text{pa}(v)} = \theta_v$ and that $\text{pa}(v)$ takes one value only.

First let us keep the distribution T fixed. As shown in section 2 this distribution can be represented either by (1) or by (2), the latter representation having a distinct form for each possible choice of the root(s). These representations however will assign exactly the same probability $T(x)$ to an observation x , so there is no way to distinguish between them from the point of view of the data. Thus we shall require that the corresponding parameter sets are also the same from the point of view of the prior. This leads to the assumption of *Likelihood equivalence*:

Assumption 1 (Likelihood equivalence) *Let T be a tree distribution having structure E , \overline{E} a directed tree structure obtained from E and $\theta_E, \theta_{\overline{E}}$ the*

respective parameters of T . Denote by $|\frac{\partial \theta_{\overline{E}}}{\partial \theta_E}|$ the Jacobian of the transformation $\theta_E \rightarrow \theta_{\overline{E}}$. Then $P_0(\theta_{\overline{E}}(\theta_E)|\overline{E})|\frac{\partial \theta_{\overline{E}}}{\partial \theta_E}| = P_0(\theta_E|E)$.

This assumption states that in all possible parametrizations consistent with a given structure E the prior will assign the same probability mass to any given (measurable) subset in parameter space. Thus, the prior treats likelihood equivalent parametrizations as indistinguishable.

Likelihood equivalence has somewhat compressed the space that we have to define P_0 on, but it still leaves us with the task of assigning a separate prior for each (undirected) tree structure. We now transform this problem into one of assigning a prior for each of the possible tree edges by making the following additional assumptions:

Assumption 2 (Parameter independence) *For any structure \overline{E} and any $\overline{vu} \in \overline{E}, j, j' = 1, \dots, r_v, j' \neq j$ the parameter vectors $\theta_{u|v}(\cdot|j)$ and $\theta_{u|v}(\cdot|j')$ are independent under P_0 . The parameters $\theta_{u|v}(\cdot|j)$ are also independent under (P_0) of the parameter sets $\theta_{u'|v'}(\cdot|j')$ corresponding to any other edge in \overline{E} .*

Assumption 3 (Parameter modularity) *The prior $P_0(\theta_{u|v}|\overline{E})$ is the same for all structures \overline{E} that contain the edge \overline{vu} .*

In other words, parameter independence states that the prior over parameters factors into a product over the edges; by stating in addition that the prior for an edge is the same for all tree structures that contain that edge, we have effectively removed the dependence on E from the parameters prior. From now on, we will write $P_0(\theta_E), P_0(\theta_{\overline{E}})$ instead of $P_0(\theta_E|E)$ and $P_0(\theta_{\overline{E}}|\overline{E})$ respectively. We shall call a prior P_0 satisfying assumptions 1, 2 and 3 a *decomposable prior* for tree parameters. If both $P_0(E)$ and $P_0(\theta_E)$ are decomposable, the resulting prior over tree distribution is also called *decomposable*. A decomposable prior is defined by a number of order n^2 parameters only. Moreover, as we shall see next, the same assumptions also constrain the functional form the prior can have. Again we assume that G^{sup} is connected.

Assumption 4 (Connectivity) *The support graph of $P_0(E)$ is connected.*

Theorem 5 *Let $P(T) = P(E)P(\theta_E)$ be a decomposable distribution over tree parameters, for which the support graph of $P(E)$ is connected and*

$P(\theta_E) > 0$ for $\theta_E > 0$. Then for any tree T in any directed representation $\overline{E}, \theta_{\overline{E}}$:

$$P(\theta_{\overline{E}}|\overline{E}) = \prod_{v \in V} P(\theta_{v|\text{pa}(v)}) \quad (20)$$

$$P(\theta_{v|u}) = \prod_{i=1}^{r_u} D(\theta_{v|u}(\cdot|i); N'_{vu}(\cdot i)) \quad (21)$$

where D is the Dirichlet distribution and $N'_{vu}(ij) > 0$ are its parameters. The numbers $N'_{uv}(ij) = N'_{vu}(ji)$ are defined for all edges uv with $\beta_{uv} > 0$ and satisfy

$$\sum_{i=1}^{r_u} N'_{uv}(ij) = N'_v(j) \quad (22)$$

$$\sum_{j=1}^{r_u} N'_v(j) = N' \quad (23)$$

The Dirichlet prior is defined over the parameter space $\theta_1, \dots, \theta_r$, ($\sum_j \theta_j = 1$, $\theta_j > 0, j = 1, \dots, r$) of a distribution over a discrete set by

$$D(\theta_1, \dots, \theta_r; N'_1, \dots, N'_r) = \frac{1}{Z_D} \prod_{j=1}^r \theta_j^{N'_j-1} \quad (24)$$

The numbers $N'_1, \dots, N'_r > 0$ are the hyper-parameters of the Dirichlet prior, their sum is denoted by N' . The normalization constant Z_D has the form

$$Z_D = \frac{\prod_{j=1}^r \Gamma(N'_j)}{\Gamma(N')} \quad (25)$$

with Γ denoting the Euler function, $\Gamma(p) = \int_0^\infty x^{p-1} e^{-x} dx$.

The above line of reasoning parallels the one in HGC. The assumptions 1 – 3 are the specialization for tree structures of their homonyms in HGC. But unlike the case of general Bayes nets, where the prior is in general specified by an exponential number of parameters, in the case of tree graphical models the prior can be specified by a set of only $\mathcal{O}(n^2 r_{MAX}^2)$ “pairwise marginal counts” $N'_{uv}(ij)$. This is possible because in the space of tree structures the likelihood equivalence classes can be explicitly represented and the number of possible parents for a variable is no larger than one. Therefore, not only

the tree belief net itself, but also any decomposable distribution over trees can be completely defined in terms of pairwise interactions⁴.

The same properties allow us to replace a fourth assumption made by HGC, namely structure possibility, with the weaker assumption 4. Note that if all $\beta_{uv} > 0$, then all tree structures are possible, G^{sup} is connected and our theorem 5 is an exact rewrite of the similar result in HGC. This last assumption is not essential for our results. In section 6 we shall give a general formulation of the above theorem that dispenses with the connectivity assumption.

To summarize, starting with the assumptions 1–3 and aiming mainly at obtaining a tractable and consistent prior representation, we have arrived at the conclusion that the prior has to be a product of Dirichlet distributions. This demonstrates that our initial requirement is essentially a drastic one; the restrictions on the prior should be understood as restrictions on the type of prior information about the model we are allowed to have. A Dirichlet distribution means essentially that we only have knowledge about the values of the parameters' means. This issue is further developed in HGC to which we refer the reader. On the computational side however, the advantage is enormous, since with the Dirichlet distribution Bayesian learning is possible in closed form. The next section will exploit exactly this property to find the posterior over tree distributions.

5 Bayesian learning with decomposable priors

From equations (1,2) we know that the likelihood can be written as a product over tree edges. Theorem 5 proves the same thing about the decomposable prior. It follows then that the posterior $P(T|\mathcal{D})$ in equation (3) can also be factored over the edges of T . We shall see that in addition $P(T|\mathcal{D})$ is decomposable and the normalization constant $P(\mathcal{D}) = Z_{\mathcal{D}}$ can be computed tractably.

We shall use the following important property of a Dirichlet distribution: Assume a discrete variable z that takes values $1 \dots r$ with probabilities $\theta = (\theta_1, \dots, \theta_r)$, a prior for θ that is Dirichlet with parameters $N'(1), \dots, N'(r)$ and a set \mathcal{D}_z of N independent observations for z , such that the value j appears $N(j)$ times in \mathcal{D}_z . Then, the posterior of the parameters θ is (see

⁴It is important to note that the parameters $N'_{uv}(ij)$ cannot be set arbitrarily. They have to be proportional to the marginals of some distribution over V .

e.g. [2]) is also a Dirichlet distribution with parameters $N'(j) + N(j)$.

$$P(\theta|\mathcal{D}_z) = D(\theta(\cdot); N'(\cdot) + N(\cdot)) \quad (26)$$

This result applies immediately to the posterior of a tree. Let us denote by $N_{uv}(ij)$ and $N_v(j)$ the *sufficient statistics* of the sample \mathcal{D} , i.e. the number of times $u = i, v = j$ and respectively $v = j$ in \mathcal{D} . Then, from (26) and theorem 5 we obtain

$$\begin{aligned} P(T|\mathcal{D}) &= \quad (27) \\ &= \frac{P(T)P(\mathcal{D}|T)}{Z_{\mathcal{D}}} \\ &= \frac{1}{Z_{\mathcal{D}}} \frac{1}{Z_{\beta}} \prod_{uv \in E} \beta_{uv} \prod_{v \in V} \prod_{i=1}^{r_{\text{pa}(v)}} \left[\frac{(N'_{\text{pa}(v)}(i)) \prod_{j=1}^{r_v}, (N_{\text{vpa}(v)}(ji) + N'_{\text{vpa}(v)}(ji))}{(N_{\text{pa}(v)}(i) + N'_{\text{pa}(v)}(i)) \prod_{j=1}^{r_v}, (N'_{\text{vpa}(v)}(ji))} \right] \\ &\quad \cdot D(\theta_{v|\text{pa}(v)}(\cdot|i); N_{\text{vpa}(v)}(i.) + N'_{\text{vpa}(v)}(i.)) \end{aligned}$$

Hence, $P(T|\mathcal{D})$ is also a decomposable distribution over tree structures, and its parameters are available directly from the parameters of the prior and the sufficient statistics of the sample. It remains to show how to compute the normalization constant $Z_{\mathcal{D}}$. For this, we will first keep the structure E fixed and integrate over the parameters $\theta_{\bar{E}}$ in some directed structure \bar{E} obtained from E . Since the parameters' posterior in (27) is already in the form of a normalized distribution, which consequently integrates to 1, we obtain after a simple calculation:

$$\begin{aligned} \int P(T|\mathcal{D}) d\theta_{\bar{E}} &= \quad (28) \\ &= \frac{1}{Z_{\mathcal{D}}} \frac{1}{Z_{\beta}} \frac{(N')}{(N' + N)} \prod_{v \in V} \prod_{j=1}^{r_v} \frac{(N_v(j) + N'_v(j))}{(N'_v(j))} \\ &\quad \cdot \prod_{uv \in E} \left[\beta_{uv} \prod_{i=1}^{r_u} \frac{(N'_u(i))}{(N_u(i) + N'_u(i))} \prod_{j=1}^{r_v} \frac{(N'_v(j))}{(N_v(j) + N'_v(j))} \prod_{i=1}^{r_u} \prod_{j=1}^{r_v} \frac{(N_{uv}(ij) + N'_{uv}(ij))}{(N'_{uv}(ij))} \right] \end{aligned}$$

This quantity represents the marginal posterior $P(E|\mathcal{D})$; as required by likelihood equivalence, it is the same no matter how \bar{E} is obtained from E . Note also that $P(E|\mathcal{D})$ decomposes into a product over the edges in E preceded by a factor independent of E . We define the weights W_{uv} as

$$W_{uv} = \prod_{i=1}^{r_u} \frac{(N'_u(i))}{(N_u(i) + N'_u(i))} \prod_{j=1}^{r_v} \frac{(N'_v(j))}{(N_v(j) + N'_v(j))} \prod_{i=1}^{r_u} \prod_{j=1}^{r_v} \frac{(N_{uv}(ij) + N'_{uv}(ij))}{(N'_{uv}(ij))} \quad (29)$$

Now we can apply theorem 2 again, to the weights βW obtaining

$$1 = \sum_E P(E|\mathcal{D}) = \frac{1}{Z_{\mathcal{D}}} \frac{1}{Z_{\beta}}, \frac{1}{(N'+N)}, \binom{N'}{(N'+N)} \prod_{v \in V} \prod_{j=1}^{r_v} \frac{\binom{N_v(j) + N'_v(j)}{(N'_v(j))}}{|Q(\beta W)|} \quad (30)$$

From this we obtain that

$$Z_{\mathcal{D}} = \frac{|Q(\beta W)|}{|Q(\beta)|}, \frac{1}{(N'+N)} \prod_{v \in V} \prod_{j=1}^{r_v} \frac{\binom{N_v(j) + N'_v(j)}{(N'_v(j))}}{\quad} \quad (31)$$

This completely defines the posterior distribution $P(T|\mathcal{D})$. The posterior probability of any tree distribution T can be now computed analytically based on equations (27) and (31) while (28) and (31) give the posterior of any tree structure E . Note that the weights W_{uv} are never 0, so that the support graph of the posterior distribution coincides with the G^{sup} of the prior.

To compute the posterior representation from the data set we need $\mathcal{O}(n^2 r_{MAX}^2 N)$ operations to obtain the sufficient statistics and to evaluate the edge weights W_{uv} and an additional $\mathcal{O}(n^3)$ to evaluate the normalization constant $Z_{\mathcal{D}}$ for a total of $\mathcal{O}(n^2 r_{MAX}^2 N + n^3)$ operations. Computing the posterior of a tree (or tree structure) is now $\mathcal{O}(nr_{MAX})$.

Furthermore, to perform Bayesian averaging in computing the probability of a new data point x one has to evaluate

$$P(x|\mathcal{D}) = \int T(x) P(T|\mathcal{D}) dT \quad (32)$$

Just as before, we can first integrate the above expression over the parameters for a fixed structure E and then perform a summation over structures. The former step yields

$$\begin{aligned} \int \prod_{v \in V} T_{v|\text{pa}(v)}(x_v | x_{\text{pa}(v)}) \prod_{v \in V} \prod_{i=1}^{r_{\text{pa}(v)}} D(\theta_{v|\text{pa}(v)}(\cdot|i); N'_{v\text{pa}(v)}(\cdot,i) + N_{v\text{pa}(v)}(\cdot,i)) &= \\ &= \prod_{v \in V} \frac{N'_{v,\text{pa}(v)}(x_v x_{\text{pa}(v)}) + N_{v,\text{pa}(v)}(x_v x_{\text{pa}(v)})}{N'_{\text{pa}(v)}(x_{\text{pa}(v)}) + N_{\text{pa}(v)}(x_{\text{pa}(v)})} \end{aligned} \quad (33)$$

$$= \frac{1}{N' + N} \underbrace{\prod_{v \in V} [N'_{\text{pa}(v)}(x_{\text{pa}(v)}) + N_{\text{pa}(v)}(x_{\text{pa}(v)})]}_{w_0(x)} \quad (34)$$

$$\prod_{uv \in E} \frac{N'_{uv}(x_v x_u) + N_{uv}(x_v x_u)}{\underbrace{[N'_u(x_u) + N_u(x_u)][N'_v(x_v) + N_v(x_v)]}_{w_{uv}(x)}}$$

Again, we note that the result includes a structure independent factor $w_0(x)$ and a product of factors corresponding to the tree edges $w_{uv}(x)$. Also, we note that the final result is invariant to the particular orientation \overline{E} of E . Summing now over tree structures is a mere exercise now; we have

$$P(x|\mathcal{D}) = \sum_E \frac{w_0(x)}{Z_{\mathcal{D}}} \prod_{uv \in E} \beta_{uv} w_{uv}(x) \quad (35)$$

$$= \frac{w_0(x) |Q(\beta w(x))|}{Z_{\mathcal{D}}} \quad (36)$$

The averaging involves computing the edge weights $w(x)$ and evaluating a determinant, so that the total computation is $\mathcal{O}(n^3)$, a relatively large value compared to the $\mathcal{O}(n)$ demands of the ML and MAP tree likelihood.

The result generalizes readily to the Bayesian averaging of the probability of a set of more than one independent observations.

6 Disconnected support graph

Here we generalize the previous results to the case when the support graph is disconnected. In other words, we discuss the case when the prior $P(E)$ enforces probabilistic independence between the edges and parameters corresponding to edges in different components.

The intuition behind the following results stems from the fact that for a disconnected support graph the domain V is effectively partitioned into K subsets of variables V^k , $k = 1, \dots, K$ each corresponding to a connected component of G^{sup} . We denote these components by $G^k = (V^k, E^k)$, $k = 1, \dots, K$. We also introduce a notation similar to (8), to denote a set of values corresponding to pairs of variables in a subset U of V .

$$a_U = \{a_{uv}, u \neq v, u, v \in U \subseteq V\} \quad (37)$$

As we shall see, the subsets V^k behave as separate and independent domains from all points of view.

Because there can be no edges between the subsets, it is easy to see that

$$\beta = \bigcup_{k=1}^K \beta_{V^k} \quad (38)$$

This holds in general for a function f defined over E^{sup} .

A disconnected graph will have no spanning tree (hence the value returned by the Matrix Tree theorem will be 0) but it will have *maximal trees*, i.e. trees having a maximal number of edges⁵. A maximal tree is composed of spanning trees corresponding to each G^k . Therefore the number of distinct maximal trees in G^{sup} is the product of the numbers of spanning trees in its connected components:

$$\# \text{ maximal trees}(G^{sup}) = \prod_{k=1}^K \# \text{ spanning trees}(G^k) \quad (39)$$

The above remarks allow us to prove a generalization of theorem 2.

Theorem 6 *Let $P(E)$ be a distribution over maximal tree structures defined by (5,6) with $\beta \geq 0$. Then the normalization constant Z is equal to*

$$Z = \prod_{k=1}^K |Q(\beta_{V^k})| \quad (40)$$

The proof of this theorem is an easy consequence of theorem 2 and of the previous remarks and therefore it is omitted.

Intuitively, one can imagine redefining $Q(\beta)$ as a block diagonal matrix of dimension $n - K$ consisting of blocks $Q(\beta_{V^k})$, $k = 1, \dots, K$. Then, one can rewrite (40) so as to obtain a form identical with the one in theorem 2

$$Z = |Q(\beta)| \quad (41)$$

If one defines $M(\beta_{V^k})$ to correspond to each $Q(\beta_{V^k})$ in a way similar to equation (10), then one has

$$\frac{\partial}{\partial \beta_{uv}} |Q(\beta_{V^k})| = M_{uv}(\beta_{V^k}) |Q(\beta_{V^k})| \quad \text{for } u, v \in V^k \quad (42)$$

Now assembling $M(\beta_{V^k})$ for $k = 1, \dots, K$ in a block diagonal matrix $M(\beta)$ with

$$M_{uv}(\beta) = \begin{cases} M_{uv}(\beta_{V^k}), & \text{for } u, v \in V^k \\ 0, & \text{otherwise} \end{cases} \quad (43)$$

one can formally recover equation (11):

$$\frac{\partial Z}{\partial \beta_{uv}} = M_{uv}(\beta) \prod_{k=1}^K |Q(\beta_{V^k})| = M_{uv}(\beta) |Q(\beta)|. \quad (44)$$

⁵This number is $n - K$.

It is also worth making the following point: $M(\beta_{V^k})$ (or $M(\beta)$ in the case of a connected G^{sup}) are defined only when $|Q(\beta_{V^k})|$ is non-zero. However, the derivative of $|Q(\beta_{V^k})|$ exists in all cases and is defined in terms of the minor determinants of the elements of $Q(\beta_{V^k})$. The minor A_{uv}^* of a square matrix A is obtained by deleting row u and column v from A and computing the determinant of the remaining matrix. Assume that to obtain $Q(\beta_{V^k})$ from $\bar{Q}(\beta_{V^k})$ we delete row and column n_k . Then, for every $u, v \in V^k$ we have

$$\frac{\partial}{\partial \beta_{uv}} |Q(\beta_{V^k})| = \begin{cases} Q_{uu}^*(\beta_{V^k}) + Q_{vv}^*(\beta_{V^k}) - 2Q_{uv}^*(\beta_{V^k}), & \text{if } u, v \neq n_k \\ Q_{vv}^*(\beta_{V^k}), & \text{if } v \neq n_k, u = n_k \\ Q_{uu}^*(\beta_{V^k}), & \text{if } u \neq n_k, v = n_k \end{cases} \quad (45)$$

Similarly to (13), for the average of an additive function f we obtain

$$\langle f(E) \rangle_P = \sum_{k=1}^K \sum_{u, v \in V^k, u < v} f_{uv} \beta_{uv} M_{uv}(\beta_{V^k}) = \sum_{u < v} f_{uv} \beta_{uv} M_{uv}(\beta) \quad (46)$$

Finally, if g is multiplicative, equation (16) becomes

$$\langle g(E) \rangle_P = \frac{\prod_{k=1}^K |Q(\beta_{V^k} g_{V^k})|}{\prod_{k=1}^K |Q(\beta_{V^k})|} = \frac{|Q(\beta g)|}{|Q(\beta)|} \quad (47)$$

7 Ensembles of trees

In this section we consider a new probability model, called *ensembles of trees* that naturally extends the tree graphical model. To best describe this model, imagine that a tree distribution is defined in two steps: first a set of parameters θ and second the structure E . Because E is not known at the time when we choose θ , we need to specify a parameter set that is sufficiently large, so that for any E we can afterwards extract from θ the actual set of parameters θ_E . This can be done easily via the same idea that allowed us to define a decomposable prior in section 4; we choose

$$\begin{aligned} \theta &= \{ \theta_{uv}(ij), u, v \in V, i = 1, \dots, r_u, j = 1, \dots, r_v \} \\ &\cup \{ \theta_v(j), v \in V, j = 1, \dots, r_v \} \end{aligned} \quad (48)$$

such that

$$\sum_{i=1}^{r_u} \theta_{uv}(ij) = \theta_v(j) \quad \forall u \in V \quad (49)$$

$$\sum_{j=1}^{r_v} \theta_v(j) = 1 \quad \forall v \in V$$

Now, changing the notation of equation (1) to emphasize the dependence on θ and E , and rearranging the factors, we write the tree distribution as

$$T(x|\theta, E) = \prod_{uv \in E} \frac{\theta_{uv}(x_u, x_v)}{\theta_u(x_u)\theta_v(x_v)} \prod_{v \in V} \theta_v(x_v) \quad (50)$$

The ensemble of trees $R(x)$ is a weighted average of all the possible tree distributions sharing the same parameters θ . To ensure tractability, the weights will represent a decomposable distribution over spanning tree structures as in (5).

$$R(x) = \sum_E P(E) T(x|\theta, E) \quad (51)$$

If we use the notations

$$w_{uv}(x) = \frac{\theta_{uv}(x_u, x_v)}{\theta_u(x_u)\theta_v(x_v)} \quad (52)$$

$$w_0(x) = \prod_{v \in V} \theta_v(x_v) \quad (53)$$

for the edge dependent and respectively edge independent factors in (50) then, by theorem 2, $R(x)$ has an alternative, tractable form

$$R(x) = w_0(x) \frac{|Q(w\beta)|}{|Q(\beta)|} \quad (54)$$

The ensemble of trees can be seen as a mixture model whose components are the trees over V parametrized by θ . The weighted averaging corresponds then to the presence of a hidden variable z taking as many values as there are structures, each with probability $Pr[z = E] = P(E)$. Therefore, the (generalized) EM algorithm [3] can be considered as a possibility for learning the parameters. We shall not pursue this issue in detail, but we will mention the following: the E step of the algorithm is tractable and straightforward given equation (51); the M-step however cannot be performed exactly and it is not known if the expression to be maximized has a unique local maximum.

But if we assume a set of complete observations \mathcal{D} as before, the likelihood of this data set, denoted by $R(\mathcal{D})$, can be optimized w.r.t. the parameters θ and β by gradient ascent. We shall denote by $M_{uv}(\beta)$ and

$M_{uv}(\beta w(x^t))$ respectively the values in equation (10) that correspond to $Q(\beta)$ and $Q(\beta w(x^t))$. Using lemma 3 we obtain

$$\frac{\partial \log R(\mathcal{D})}{\partial \beta_{uv}} = \sum_{t=1}^N \frac{w_{uv}(x^t) M_{uv}(x^t)}{|Q(\beta w(x^t))|} - N M_{uv}(\beta) \quad (55)$$

$$\frac{\partial \log R(\mathcal{D})}{\partial \theta_{uv}(ij)} = \frac{\beta_{uv}}{\theta_u(i)\theta_v(j)} \sum_{t: x_u^t=i, x_v^t=j} M_{uv}(x^t) \quad (56)$$

$$\frac{\partial \log R(\mathcal{D})}{\partial \theta_v(j)} = \frac{1}{\theta_v(j)} \sum_{t: x_v^t=j} [1 - \sum_{v' \in V} w_{vv'}(x) M_{vv'}(x^t)] \quad (57)$$

Note that the parameters θ need to satisfy (49) and therefore we will need to perform a constrained maximization of $R(\mathcal{D})$ using e.g. Lagrange multipliers and that this method will converge to only a local optimum of the log-likelihood.

8 Discussion

This paper has presented decomposable priors, a class of priors over tree structures and parameters that makes exact Bayesian learning tractable. A decomposable prior is expressed as a product of factors, each corresponding to an edge of the tree. The same edge contributes the same amount in every tree structure that includes it. This property allows (1) representing the prior by order n^2 parameters and (2) using the Matrix tree theorem to integrate the prior in closed form.

It is remarkable that for trees, the standard assumptions of HGC are sufficient to ensure tractability. In fact, these assumptions are no stronger than the assumptions of *functional independence* implicit in the original Chow and Liu algorithm [1, 7].

But is worth highlighting again that these assumptions are restrictive, in the sense of drastically limiting the type of prior knowledge that can be used efficiently in the Bayesian learning of trees. Prior knowledge that violates assumptions 1–3 is e.g. knowledge that two edges are more likely to appear simultaneously than separately in a tree structure, or knowledge that two edges have the same parameters. This problem is not specific to trees, but to Bayes nets in general. Therefore, a worthwhile area of future research is discovering tractable methods to deal with such type of knowledge in the case of tree structures or in the case of general Bayes nets.

We have also introduced ensembles of trees as a tractable extension to the tree model. Ensembles of trees can be learned in the ML framework. Exploring the properties of the new model and of the learning algorithm itself are areas of continuing research.

A Proof of Theorem 2

Theorem 2 *Let $P(E)$ be a distribution over spanning tree structures defined by (5,6). Then the normalization constant Z is equal to $|Q(\beta)|$ with $Q(\beta)$ being the first $(n - 1)$ lines and columns of the matrix $\overline{Q}(\beta)$ given by:*

$$\overline{Q}_{uv}(\beta) = \overline{Q}_{vu}(\beta) = \begin{cases} -\beta_{uv} & 1 \leq u < v \leq n \\ \sum_{v'=1}^n \beta_{v'v} & 1 \leq u = v \leq n \end{cases} \quad (58)$$

Proof We shall prove the result first for positive integer values of β , then for positive rational values. Then, because the determinant is a continuous function it will follow that the theorem is true for any real, positive values β .

Assume β are integers. Then β_{uv} can represent the number of edges between u and v of a multigraph over V . The degree of node v equals the number of edges incident to v ; this number is $\sum_{u \neq v} \beta_{uv}$. Thus, by theorem 1, the total number of spanning trees in this graph is $Z_1 = |Q(\beta)|$. We now show that $Z_1 = Z$ (where Z is defined by (6)).

$$\begin{aligned} Z_1 &= \sum_E \# \text{distinct trees having structure } E & (59) \\ &= \sum_E \prod_{uv \in E} \beta_{uv} \\ &= Z \end{aligned}$$

Assume now that β are all rational. Let m be the common denominator of β , such that $\beta_{uv} = a_{uv}/m$ with a_{uv} integer. Then $Q(a) = m\overline{Q}(\beta)$ is a matrix of integers and by virtue of (59) we have that

$$|Q(a)| = \sum_E \prod_{uv \in E} a_{uv} \quad (60)$$

But $|mQ(\beta)| = m^{n-1}|Q(\beta)|$ and therefore

$$|Q(\beta)| = \frac{1}{m^{n-1}} \sum_E \prod_{uv \in E} a_{uv}$$

$$\begin{aligned}
&= \sum_E \prod_{uv \in E} \frac{a_{uv}}{m} \\
&= \sum_E \prod_{uv \in E} \beta_{uv} \\
&= Z
\end{aligned}$$

■

B Proof of lemma 3

Lemma 3 *Let Z be given by equation (6) with $\beta \geq 0$, $Q(\beta)$ be given by theorem 2, Q^{-1} be the inverse of Q and $M(\beta)$ be a symmetric matrix with 0 diagonal defined by*

$$\begin{aligned}
M_{uv} &= (Q^{-1})_{uu} + (Q^{-1})_{vv} - 2(Q^{-1})_{uv}, \quad u, v < n \\
M_{nv} &= M_{vn} = (Q^{-1})_{vv}, \quad v < n \\
M_{vv} &= 0
\end{aligned} \tag{61}$$

Then the partial derivative of Z with respect to β_{uv} is

$$\frac{\partial Z}{\partial \beta_{uv}} = M_{uv}(\beta) |Q(\beta)|. \tag{62}$$

Proof The proof uses the fact that, for any nonsingular matrix with elements A_{ij}

$$\frac{\partial |A|}{\partial A_{ij}} = |A| (A^{-1})_{ij} \tag{63}$$

Then, for $u, v < n$, taking into account that the only elements of $Q(\beta)$ that depend on β_{uv} are Q_{uu} , Q_{vv} , Q_{uv} and Q_{vu} we have successively:

$$\frac{\partial Z}{\partial \beta_{uv}} = \frac{\partial |Q(\beta)|}{\partial \beta_{uv}} = \sum_{i,j \in V} \frac{\partial |Q(\beta)|}{\partial Q_{ij}} \frac{\partial Q_{ij}}{\partial \beta_{uv}} = |Q(\beta)| [(Q^{-1})_{uu} + (Q^{-1})_{vv} - 2(Q^{-1})_{uv}] \tag{64}$$

Similarly, if $u = n$, $v < n$, then β_{vn} appears only in Q_{vv} . Hence

$$\frac{\partial Z}{\partial \beta_{uv}} = |Q(\beta)| (Q^{-1})_{vv} \tag{65}$$

■

C Proof of lemma 4

Lemma 4 *Let $P(E)$, Q and M be given by (5), theorem 2 and (10) respectively and f be an additive function of the structure E . Then the average of f under P is*

$$\langle f(E) \rangle_P = \sum_E f(E)P(E) \quad (66)$$

$$= \sum_{u<v} f_{uv}\beta_{uv}M_{uv}(\beta) \quad (67)$$

$$= \text{trace } Q(f\beta)Q^{-1}(\beta) \quad (68)$$

Proof We first introduce the following lemma:

Lemma 7 [6] *If $P(E)$ is given by equation (5) and f is an additive function of E then*

$$\langle f(E) \rangle_P = \frac{1}{Z} \left. \frac{\partial \log Q(\beta e^{\alpha f})}{\partial \alpha} \right|_{\alpha=0} \quad (69)$$

This lemma can be easily proved by equating $Q(\beta e^{\alpha f})$ with its definition (9) and then taking derivatives of both sides.

Then, to obtain the result of lemma 4, equation (67) we use (63) again, conveniently grouping the terms involving f_{uv} afterwards. To derive to compact form (68) we notice that (67) represents the sum of element-wise products of $Q(f\beta)$ and $Q^{-1}(\beta)$ and use the well-known matrix identity:

$$\sum_{ij} A_{ij}B_{ij} = \text{trace } AB^T \quad (70)$$

■

D Proof of theorem 5

Theorem 5 *Let $P(T) = P(E)P(\theta)$ be a decomposable distribution over tree parameters, for which the support graph of $P(E)$ is connected and $P(\theta) > 0$ for $\theta > 0$. Then for any tree T in any directed representation $\overline{E}, \theta_{\overline{E}}$:*

$$P(\theta_{\overline{E}}|\overline{E}) = \prod_{v \in V} P(\theta_{v|\text{pa}(v)}) \quad (71)$$

$$P(\theta_{v|u}) = \prod_{i=1}^{r_u} D(\theta_{v|u}(\cdot|i); N'_{vu}(\cdot i)) \quad (72)$$

where D is the Dirichlet distribution and $N'_{vu}(ij) > 0$ are its parameters. The numbers $N'_{uv}(ij) = N'_{vu}(ji)$ are defined for all edges uv with $\beta_{uv} > 0$ and satisfy

$$\sum_{i=1}^{r_u} N'_{uv}(ij) = N'_v(j) \quad (73)$$

$$\sum_{j=1}^{r_v} N'_v(j) = N' \quad (74)$$

Proof Begin by fixing two variables $u, v \in V$ such that $\beta_{uv} > 0$ and a structure E containing edge uv . Denote by \overline{E}^1 and \overline{E}^2 the orientations of E where u and respectively v are root. Thus \overline{E}^1 and \overline{E}^2 differ only in the orientation of edge uv . Let θ^1 and θ^2 be two parametrizations corresponding to \overline{E}^1 and \overline{E}^2 respectively, such that θ^1 and θ^2 produce the same distribution T . For every θ^1 there will be a unique θ^2 (obtained by applying (1) and (2)) satisfying this condition. Moreover,

$$\theta^1_{w|\text{pa}(w)} = \theta^2_{w|\text{pa}(w)} \quad \text{for all } w \neq u, v. \quad (75)$$

The prior distribution of θ^1 is

$$P_0^1(\theta^1) = P_0^1(\theta_u^1) P_0^1(\theta_v^1|u) \prod_{w \neq u, v} P_0^1(\theta_{w|\text{pa}(w)}^1) \quad (76)$$

Similarly, for θ^2 we have

$$P_0^2(\theta^2) = P_0^2(\theta_v^2) P_0^2(\theta_u^2|v) \prod_{w \neq u, v} P_0^2(\theta_{w|\text{pa}(w)}^2) \quad (77)$$

The likelihood equivalence assumption, together with the change of variable formula, imply

$$P_0^1(\theta^1) = P_0^2(\theta^2(\theta^1)) \left| \frac{\partial \theta^2}{\partial \theta^1} \right| \quad (78)$$

In the above, $\left| \frac{\partial \theta^2}{\partial \theta^1} \right|$ denotes the Jacobian of the transformation $\theta^1 \rightarrow \theta^2$. Due to the equality (75) this Jacobian has the following structure:

$$\left| \frac{\partial \theta^2}{\partial \theta^1} \right| = \begin{vmatrix} \frac{\partial \theta_v^2}{\partial \theta_u^1} & \frac{\partial \theta_{u|v}^2}{\partial \theta_u^1} & 0 & 0 & \dots & 0 \\ \frac{\partial \theta_v^2}{\partial \theta_v^1|u} & \frac{\partial \theta_{u|v}^2}{\partial \theta_v^1|u} & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & 0 & \dots & 1 \end{vmatrix} = \begin{vmatrix} \frac{\partial \theta_v^2}{\partial \theta_u^1} & \frac{\partial \theta_{u|v}^2}{\partial \theta_u^1} \\ \frac{\partial \theta_v^2}{\partial \theta_v^1|u} & \frac{\partial \theta_{u|v}^2}{\partial \theta_v^1|u} \end{vmatrix}$$

Let us denote the resulting determinant by $\left| \frac{\partial(\theta_v^2, \theta_{u|v}^2)}{\partial(\theta_u^1, \theta_{v|u}^1)} \right|$. The above equality together with (76–78) implies

$$P_0^1(\theta_u^1, \theta_{v|u}^1) = P_0^2(\theta_v^2(\theta_u^1, \theta_{v|u}^1), \theta_{u|v}^2(\theta_u^1, \theta_{v|u}^1)) \left| \frac{\partial(\theta_v^2, \theta_{u|v}^2)}{\partial(\theta_u^1, \theta_{v|u}^1)} \right| \quad (79)$$

Now we have reached our first partial goal, because by theorem 7 in HGC (79) implies that $P_0(\theta_{uv})$ as well as $P_0^1(\theta_u^1, \theta_{v|u}^1)$ and $P_0^2(\theta_v^2, \theta_{u|v}^2)$ are Dirichlet. Let us denote by $N'_{uv}(ij)$, $i = 1, \dots, r_u$, $j = 1, \dots, r_v$ the parameters of $P(\theta_{uv})$ and by $N'_u(i)$ and $N'_v(j)$ their sums over j and i respectively. It is easy to show then that

$$P_0^1(\theta_u^1) = D(\theta_u^1(\cdot); N'_u(\cdot)) \quad (80)$$

$$P_0^1(\theta_{v|u}^1) = \prod_{i=1}^{r_u} D(\theta_{v|u}^1(\cdot|i); N'_{uv}(i \cdot)) \quad (81)$$

$$P_0^2(\theta_v^2) = D(\theta_v^2(\cdot); N'_v(\cdot)) \quad (82)$$

$$P_0^2(\theta_{u|v}^2) = \prod_{j=1}^{r_v} D(\theta_{u|v}^2(\cdot|j); N'_{uv}(\cdot j)) \quad (83)$$

By parameter modularity, these identities are true for any tree structure containing an edge between u and v . Therefore, all that remains to be shown is that $N' = \sum_{j \in \Omega_v} N'_v(j)$ has the same value for all $v \in V$.

Let $u \neq v$ be another variable in V . We prove that

$$\sum_{i \in \Omega_u} N'_u(i) = \sum_{j \in \Omega_v} N'_v(j) \quad (84)$$

We distinguish two cases: $\beta_{uv} > 0$ and $\beta_{uv} = 0$. The former is already solved, by virtue of the first part of the proof. It remains to show that (84) holds when $\beta_{uv} = 0$ and no tree structure will contain edge uv .

Indeed, because G^{sup} is connected, there will exist at least a path between u and v in G^{sup} . Now we can easily show (84) by induction over the length of the path. ■

This theorem and its proof are easily extended to multiply connected G^{sup} . In that case, each component of the graph will have its own equivalent sample size N'^k .

References

- [1] C. K. Chow and C. N. Liu. Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, IT-14(3):462–467, May 1968.
- [2] Morris H. DeGroot. *Probability and Statistics*. Addison–Wesley Pub. Co., 1975.
- [3] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, B*, 39:1–38, 1977.
- [4] Frank Harary. Graphs and matrices. *SIAM Review*, 9(1):83–90, 1967.
- [5] David Heckerman, Dan Geiger, and David M. Chickering. Learning Bayesian networks: the combination of knowledge and statistical data. *Machine Learning*, 20(3):197–243, 1995.
- [6] Tommi Jaakkola, Marina Meilă, and Tony Jebara. Maximum entropy discrimination. In Sara A. Solla, Todd K. Leen, and Klaus-R. Müller, editors, *Neural Information Processing Systems*, volume 12, pages 470–476. MIT Press, 2000.
- [7] Marina Meilă-Predoviciu. *Learning with Mixtures of Trees*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [8] Michael K. Murray and John W. Rice. *Differential geometry and statistics*. Chapman & Hall, 1993.
- [9] Douglas B. West. *Introduction to Graph Theory*. Prentice Hall, Upper Saddle River, NJ, 1996.