

The curved exponential family of a staged tree*

Christiane Görgen

Mathematisches Institut, Universität Leipzig, Germany
e-mail: goergen@math.uni-leipzig.de

Manuele Leonelli

School of Human Sciences and Technology, IE University, Madrid, Spain
e-mail: manuele.leonelli@ie.edu

Orlando Marigliano

KTH Royal Institute of Technology, Stockholm, Sweden
e-mail: orlandom@kth.se

Abstract: Staged tree models are a discrete generalization of Bayesian networks. We show that these form curved exponential families and derive their natural parameters, sufficient statistic, and cumulant-generating function as functions of their graphical representation. We give necessary and sufficient graphical criteria for classifying regular subfamilies and discuss implications for model selection.

MSC2020 subject classifications: Primary 62H99; secondary 68T99.

Keywords and phrases: Bayesian information criterion, chain event graph, curved exponential family, graphical model, staged tree.

Received December 2020.

1. Introduction

Staged trees define statistical models which can account for a variety of partial and asymmetric conditional independence statements between discrete events (Collazo, Görgen and Smith, 2018). The use of these graphical models in applications is constantly increasing (Barclay, Hutton and Smith, 2013; Collazo and Smith, 2015; Keeble et al., 2017) and free software for practitioners is newly available (Carli et al., 2020). However, their formal properties have only recently been studied for the first time (Görgen and Smith, 2018). We now extend this formal study by 1. proving that in general staged tree models form curved exponential families and by 2. expressing their natural parameters, sufficient statistic,

*Christiane Görgen and Manuele Leonelli were supported by the programme “Oberwolfach Leibniz Fellows” of the Mathematisches Forschungsinstitut Oberwolfach in 2017. Orlando Marigliano was supported by International Max Planck Research School and Brummer & Partners MathDataLab.

and cumulant-generating function as a function of the underlying graphical representation. We also give a graphical criterion under which they form regular exponential families.

Because exponential families exhibit a multitude of desirable inferential properties (Kass and Vos, 1997), similar characterisations for other graphical models have been studied. For instance, undirected graphical models with no hidden variables are regular exponential families (Lauritzen, 1996), Bayesian networks (e.g. Koller and Friedman, 2009) are curved exponential families, and directed graphical models with hidden variables are stratified exponential families (Geiger et al., 2001). These results are critical for model selection techniques. In particular, Haughton (1988) proved that for curved exponential families the Bayesian information criterion (Schwarz, 1978) is an asymptotically valid rule. To this date model selection for staged trees has usually been carried out in a Bayesian fashion by selecting the maximum a posteriori model (Freeman and Smith, 2011; Barclay, Hutton and Smith, 2013; Cowell and Smith, 2014). Now, our results make it theoretically sound to use the Bayesian information criterion as well, and its implementation in the R package `stagedtrees` (Carli et al., 2020) is thus justified. This criterion has already been employed by Silander and Leong (2013) although its asymptotic geometric validity was not assured at the time.

By expressing every staged tree model explicitly as a curved exponential family, we in particular achieve such an explicit description for every discrete Bayesian network. To the best of our knowledge, this is still missing in the literature. We demonstrate on small-scale examples how a given Bayesian network's exponential form can be stated in staged-tree language. We can then also specify new graphical criteria under which the model is a regular exponential family.

2. Staged trees

2.1. Probability trees as graphical statistical models

Probability trees are highly intuitive depictions of unfoldings of discrete events (Shafer, 1996) and have been used in a variety of real-world applications (Smith, 2010; Collazo, Görden and Smith, 2018).

Let \mathcal{T} be a directed, rooted tree graph where every vertex has either no or at least two emanating edges. For simplicity, we number the inner (non-leaf) vertices v_1, \dots, v_k , with v_1 denoting the root, and count the edges emanating from each vertex with indices $j = 1, \dots, \kappa_i$ for all $i = 1, \dots, k$. To every edge we assign a positive probability label $\theta_{ij} \in (0, 1)$ such that the total sum of all labels belonging to the same vertex is always equal to one, $\sum_{j=1}^{\kappa_i} \theta_{ij} = 1$. Such a labelled tree graph is called a *probability tree*.

For all root-to-leaf paths x of \mathcal{T} we define $p_\theta(x) = \prod_{i=1}^k \prod_{j=1}^{\kappa_i} \theta_{ij}^{\alpha_{ij}(x)}$ where $\alpha_{ij}(x)$ is one if x passes through the j th edge emanating from v_i and zero otherwise. By the constraints on the θ_{ij} , the function p_θ is a probability distribution on the set of all root-to-leaf paths. We denote the size of this set by n . A *probability tree model* is then the set of all such probability distributions p_θ on n

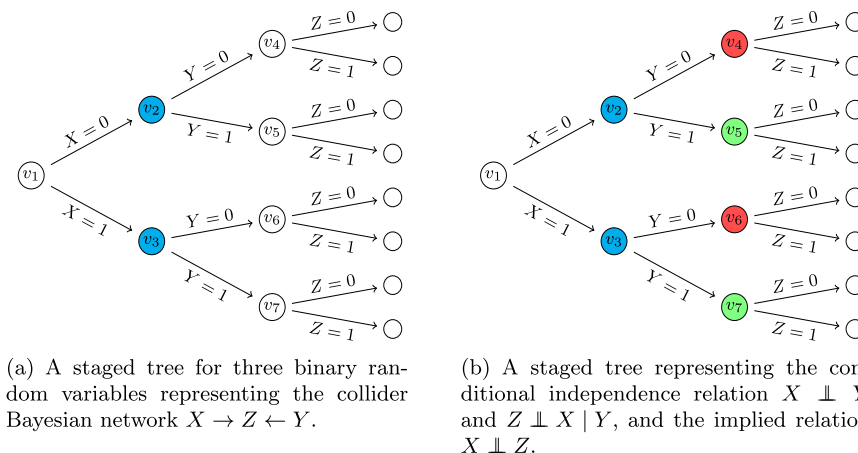


FIG 1. Two discrete graphical models. Vertices which are in the same stage have been assigned the same colour. In this picture, white-coloured vertices are not in the same stage with any other vertex.

atoms, for varying labels θ . The set of parameters θ satisfying the constraints above is the *parameter space* $\Theta_{\mathcal{T}}$ of the model. It equals the product $\times_{i=1}^k \Delta_{\kappa_i-1}^{\circ}$ of the k open probability simplices $\Delta_{\kappa_i-1}^{\circ} = \{(t_1, \dots, t_{\kappa_i}) \in \mathbb{R}^{\kappa_i} \mid \sum_{r=1}^{\kappa_i} t_r = 1 \text{ and } 0 < t_r < 1 \text{ for all } r = 1, \dots, \kappa_i\}$. Thus, the parameter space $\Theta_{\mathcal{T}}$ has dimension $d = \sum_{i=1}^k (\kappa_i - 1)$.

2.2. Probability trees and conditional independence

An event in a probability tree is a set of root-to-leaf paths, often specified by shared vertices or edges. Conditional independence relationships between events can be visualized using a simple type of colouring of these vertices in the following way.

A *staged tree* is a probability tree together with an equivalence relation on the vertex set such that two vertices are in the same stage if and only if their outgoing edges have the same attached probabilities. Figure 1 depicts a first example. A *staged tree model* is a submodel of a probability tree model where probability simplices in the parameter space have been identified with each other according to the equivalence relation.

For example, if a tree depicts the product state space of a vector of discrete random variables Y_i , then every inner vertex represents a random variable $Y_i \mid Y_{[i-1]} = y_{[i-1]}$ conditional on specific values taken by its ancestors, every edge corresponds to a state y_i of that variable, and every edge can be labelled by the conditional probability $P(Y_i = y_i \mid Y_{[i-1]} = y_{[i-1]})$ of the random variable being in that state given the particular ancestor configuration. Here, $[i - 1]$ is the set of indices $\{1, \dots, i - 1\}$. The stage relations on the tree can then be used to model conditional independence relations such as

$$P(Y_i = y_i | y_{[i-1]}) = P(Y_i = y_i | y'_{[i-1]}).$$

To illustrate this, in Fig. 1(a) the upward edge emanating from v_2 can be labeled $P(Y = 0 | X = 0)$ and the upward edge from v_4 as $P(Z = 0 | X = 0, Y = 0)$. Here, the vertices v_2 and v_3 are in the same stage. This models the independence $X \perp\!\!\!\perp Y$ via the equalities $P(Y = y | X = 0) = P(Y = y | X = 1)$ for $y = 0, 1$. Similarly in Fig. 1(b), the additional stages together entail $Z \perp\!\!\!\perp X | Y$.

Thus, staged tree models include as a special case both discrete Bayesian networks and context-specific discrete Bayesian networks, modelling conditional independence relations that hold only for a subset of the states of a discrete random variable (Boutilier et al., 1996). In particular, all conditional independence relations given by the Bayesian network can be reflected in the staging as described above. Varando, Carli and Leonelli (2021) give an algorithm that for any Bayesian network G outputs a staged tree \mathcal{T}_G representing the same model.

Because probability trees and staged trees may have root-to-leaf paths of different lengths and because there are no constraints on where stages can be imposed, the class of staged tree models is much wider than the one of discrete Bayesian networks. Working within this class unlocks an immediate graphical representation of modeling assumptions such as the space of events, the intricate conditional independence relationships within this space, and the local sum-to-one conditions. Because staged trees grow quickly even for small problems, whenever they exhibit many symmetries in their subtrees they are sometimes represented more compactly by alternative graphical depictions called *chain event graphs*: see the textbook by Collazo, Görden and Smith (2018) for an in-depth discussion of these points.

3. Staged trees are curved exponential families

3.1. Exponential families

Following Kass and Vos (1997), a parametric statistical model $\{p_\theta | \theta \in \Theta\}$ on a space \mathcal{X} is called an *exponential family* if every distribution in the model can be written in the form

$$p_\theta(x) = h(x) \exp\left(\eta(\theta)^\top T(x) - \psi(\theta)\right) \quad \text{for all } x \in \mathcal{X} \quad (1)$$

where $\eta : \Theta \rightarrow \mathbb{R}^d$ is the canonical parameter, $T : \mathcal{X} \rightarrow \mathbb{R}^d$ is a minimally sufficient statistic, $^\top$ denotes the transpose operation, and $\psi : \Theta \rightarrow \mathbb{R}$ is the cumulant-generating function. The space $\mathcal{N} = \{\eta(\theta) \in \mathbb{R}^d | (1) \text{ is integrable}\}$ is called the *natural parameter space*. If \mathcal{N} is an open and non-empty subset of \mathbb{R}^d then the model is called a *regular* exponential family of dimension d .

For exponential families, moments of all orders exist and the maximum likelihood estimator exists and is unique. Regular exponential families are closed under linear constraints on the natural parameter space and the log-likelihood function on that space is concave. Under more general constraints on the parameters, these families are more technical to study but can often retain many

useful asymptotic properties. The two best-studied generalizations of regular exponential families are so-called curved and, more generally, stratified exponential families. The development of this paper requires only the former type of models which we formally introduce in the third subsection below.

3.2. Regular exponential families

In a first step, we focus on probability trees without imposing a stage structure. In this context, the probability tree whose graph is a root vertex connected to its leaves exclusively via single edges is a *star* as in the graph-theoretic sense.

Lemma 1. *Every probability tree model on n atoms is equal to the full probability simplex Δ_{n-1}° .*

In particular, the star with n leaves and with probability θ_{1r} attached to edge r , for $r = 1, \dots, n$, represents the Multinomial distribution $\text{Multi}(1, \theta)$ with one trial and parameters $\theta = (\theta_{11}, \dots, \theta_{1n})$.

Proof. Every probability tree with n root-to-leaf paths specifies a set of distributions inside the open $n - 1$ -dimensional probability simplex as outlined in Section 2.1. Conversely, given a fixed tree graph with n leaves, every point (p_1, \dots, p_n) inside the simplex can be interpreted as a vector whose r th component p_r is the probability of going down the r th root-to-leaf path in the tree, $r = 1, \dots, n$. We can pick the label of the j th edge out of vertex v_i to be the fraction $\sum_{b \in [ij]} p_b / \sum_{a \in [ij]} p_a$ where $[i], [ij] \subseteq \{1, \dots, n\}$ denote the indices of all root-to-leaf paths passing through vertex v_i or through the tail of its j th outgoing edge, respectively. These labels are conditional probabilities and their product along a root-to-leaf path is equal precisely to the atomic probability of that path. This proves the first claim.

As for the second claim, $p_\theta(x) = \prod_{r=1}^n \theta_{1r}^{\alpha_{1r}(x)}$ where $\sum_{r=1}^n \theta_{1r} = 1$ is the probability distribution induced by the star with edges numbered $r = 1, \dots, n$. This is the stated Multinomial distribution. \square

Wishart (1949) derives the exponential form (1) for the Multinomial distribution $\text{Multi}(1, \theta)$ as follows: $h \equiv 1$ is constant, the natural parameters are normalized log-probabilities $\eta_r(\theta) = \log(\theta_{1r} / (1 - \sum_{s=1}^{n-1} \theta_{1s}))$ for $r = 1, \dots, n$, the sufficient statistic is the vector of the first $n - 1$ edge indicators $T = (\alpha_{11}, \dots, \alpha_{1, n-1})$, and the cumulant-generating function is the logarithm of the normalizing constant $\psi(\theta) = \log(1 - \sum_{s=1}^{n-1} \theta_{1s})$.

By Lemma 1, all probability trees on the same number of root-to-leaf paths are statistically equivalent in the sense that they all represent the same model, namely the full probability simplex. Thus, *any* probability tree is a graphical representation of the Multinomial distribution, and Wishart's result gives a sufficient statistic, cumulant-generating function, and the natural parameters for any probability tree, possibly after reparametrization.

In Proposition 1, we give a parametrization of the exponential family of a probability tree model which is alternative to Wishart's. This new parametrization has the advantage of respecting the structure of trees having root-to-leaf




FIG 2. Illustration of the notation used in Proposition 1. The natural parameter η_{ij} belonging to the indicator α_{ij} of “passing through the edge (v_i, v_{ij}) ” is a function of the label θ_{ij} of that edge and of the products N_i and N_{ij} of the labels of the downwards pointing v_i - and v_{ij} -to-leaf paths, respectively. The edges involved in the definition of η_{ij} are depicted in bold.

paths longer than single edges. Retaining this extra structure allows us to generalize the exponential-family parametrization of probability trees to the case of staged trees in Section 3.3.

We thus derive the following result. Here for any probability tree \mathcal{T} we say that the κ_i th edge emanating from vertex v_i points downwards, $i = 1, \dots, k$. We then recursively define functions $N_i : \Theta_{\mathcal{T}} \rightarrow \mathbb{R}$ as $N_i(\theta) = 1$ for leaf vertices v_i and else as a product of labels pointing downwards $N_i(\theta) = (1 - \sum_{s=1}^{\kappa_i-1} \theta_{i,s}) N_{i\kappa_i}(\theta)$ for $i = 1, \dots, k$. The shorthand N_{ij} denotes N_r for the vertex v_r which is the tail of the j th edge coming out of vertex v_i . Figure 2 illustrates this notation.

Proposition 1. *Every probability tree \mathcal{T} with parameters θ represents a regular exponential family with the following attributes:*

- the indicators $T_{ij} = \alpha_{ij}$ of the first $j = 1, \dots, \kappa_i - 1$ edges of all inner vertices $i = 1, \dots, k$ are a sufficient statistic,
- the natural parameters η_{ij} are locally normalized log-probabilities defined by $\eta_{ij}(\theta) = \log(\theta_{ij} N_{ij}(\theta) / N_i(\theta))$ for all $j = 1, \dots, \kappa_i - 1$ and $i = 1, \dots, k$, and the natural parameter space is \mathbb{R}^d with $d = \sum_{i=1}^k (\kappa_i - 1)$, and
- the cumulant-generating function is the negative log-sum of normalizing constants along the root-to-leaf path whose edges all point downwards, $\psi(\theta) = -\log(N_1(\theta))$.

Proof. The desired parametrization is equivalent to

$$p_{\theta}(x) = N_1(\theta) \prod_{i=1}^k \prod_{j=1}^{\kappa_i-1} \left(\theta_{ij} \frac{N_{ij}(\theta)}{N_i(\theta)} \right)^{\alpha_{ij}(x)} \quad \text{for all } x. \quad (2)$$

To prove this equality, we first rewrite the probability mass function introduced in Section 2. For any inner vertex $i = 1, \dots, k$ in the probability tree, the label of the κ_i th outgoing edge is a function of the first $1, \dots, \kappa_i - 1$ edges, namely $\theta_{i\kappa_i} = 1 - \sum_{j=1}^{\kappa_i-1} \theta_{ij}$. The indicator $\alpha_{i\kappa_i}$ of passing along that edge is a function of the indicators of those same edges and equals $\alpha_i(x)(1 - \sum_{j=1}^{\kappa_i-1} \alpha_{ij}(x))$ where $\alpha_i(x)$ is one if x reaches vertex v_i and zero otherwise.

Thus, the probability mass function of a probability tree can be written as

$$p_\theta(x) = \prod_{i=1}^k \prod_{j=1}^{\kappa_i} \theta_{ij}^{\alpha_{ij}(x)} = \prod_{i=1}^k \prod_{j=1}^{\kappa_i-1} \theta_{ij}^{\alpha_{ij}(x)} \left(1 - \sum_{j=1}^{\kappa_i-1} \theta_{ij}\right)^{\alpha_i(x)(1-\sum_{j=1}^{\kappa_i-1} \alpha_{ij}(x))}.$$

And as a consequence, the claim reduces to:

$$\prod_{i=1}^k \left(1 - \sum_{j=1}^{\kappa_i-1} \theta_{ij}\right)^{\alpha_i(x)(1-\sum_{j=1}^{\kappa_i-1} \alpha_{ij}(x))} = N_1(\theta) \prod_{i=1}^k \prod_{j=1}^{\kappa_i-1} \left(\frac{N_{ij}(\theta)}{N_i(\theta)}\right)^{\alpha_{ij}(x)} \tag{3}$$

for all root-to-leaf paths x . We now prove this claim by induction on the number of inner vertices k .

Let thus in a first step the graph be a star, $k = 1$. If x is the root-to-leaf path which is one of the first $1, \dots, \kappa_1 - 1$ edges coming out of the root then (3) states that $1 = N_1(\theta) \cdot 1/N_1(\theta) = 1$. If otherwise x equals the κ_1 st edge then $1 - \sum_{j=1}^{\kappa_1-1} \theta_{1j} = N_1(\theta)$ which is also true.

If $k > 1$ then x may have length greater one and we distinguish two analogous cases. Hereby the induction hypotheses holds for trees smaller than the one we consider and so, in particular, the claim is true for its subtrees. Without loss, we number the vertex at the tail of the first edge of x as v_2 . For simplicity, A denotes the left hand side of (3) and B the right hand side of that equation. Then:

Case 1: The edge (v_1, v_2) is one of the first $1, \dots, \kappa_1 - 1$ edges coming out of the root. Then

$$\begin{aligned} A &= 1 \cdot \prod_{i=2}^k \left(1 - \sum_{j=1}^{\kappa_i-1} \theta_{ij}\right)^{\alpha_i(x)(1-\sum_{j=1}^{\kappa_i-1} \alpha_{ij}(x))} \\ &= N_2(\theta) \prod_{i=2}^k \prod_{j=1}^{\kappa_i-1} \left(\frac{N_{ij}(\theta)}{N_i(\theta)}\right)^{\alpha_{ij}(x)} \\ &= N_1(\theta) \frac{N_2(\theta)}{N_1(\theta)} \prod_{i=2}^k \prod_{j=1}^{\kappa_i-1} \left(\frac{N_{ij}(\theta)}{N_i(\theta)}\right)^{\alpha_{ij}(x)} = B \end{aligned}$$

where the final step is true because $N_2(\theta) = N_{12}(\theta)$.

Case 2: The edge (v_1, v_2) is the κ_1 st edge coming out of the root. Then

$$\begin{aligned} A &= \left(1 - \sum_{j=1}^{\kappa_1-1} \theta_{1j}\right) \prod_{i=2}^k \left(1 - \sum_{j=1}^{\kappa_i-1} \theta_{ij}\right)^{\alpha_i(x)(1-\sum_{j=1}^{\kappa_i-1} \alpha_{ij}(x))} \\ &= \left(1 - \sum_{j=1}^{\kappa_1-1} \theta_{1j}\right) N_2(\theta) \prod_{i=2}^k \prod_{j=1}^{\kappa_i-1} \left(\frac{N_{ij}(\theta)}{N_i(\theta)}\right)^{\alpha_{ij}(x)} \\ &= N_1(\theta) \prod_{i=2}^k \prod_{j=1}^{\kappa_i-1} \left(\frac{N_{ij}(\theta)}{N_i(\theta)}\right)^{\alpha_{ij}(x)} = B \end{aligned}$$

where the final step is true because $\prod_{j=1}^{\kappa_1-1} (N_{1j}(\theta)/N_1(\theta))^{\alpha_{1j}(x)} = 1$.

This proves (2), and every probability tree represents an exponential family in this parametrization.

The regularity claim follows since the map η is a diffeomorphism between the space of model parameters Θ and \mathbb{R}^d . Indeed, it has a smooth inverse obtained by first computing p_η for a given $\eta \in \mathbb{R}^d$ according to the given parametrization and then computing $\theta \in \Theta$ such that $p_\eta = p_\theta$ as described in the proof of Lemma 1. \square

Proposition 1 enables us to simply read the parametrisation of the underlying exponential family directly from a given probability tree.

Example 1. Consider Fig. 1 and the probability tree which has the same graph as the staged trees in Fig. 1(a) and 1(b). For simplicity, in this binary tree we do not use double indices but label the upwards edges going out of vertex v_i as θ_i and the downward edges as $1 - \theta_i$ for $i = 1, \dots, 7$.

The indicator functions α_{i1} of the upwards edges $i = 1, \dots, 7$ are a sufficient statistic for this tree. The corresponding natural parameters are then log-ratios of the downwards labels derived as

$$\begin{aligned}\eta_1(\theta) &= \log(\theta_1(1 - \theta_2)(1 - \theta_5)/(1 - \theta_1)(1 - \theta_3)(1 - \theta_7)), \\ \eta_2(\theta) &= \log(\theta_2(1 - \theta_4)/(1 - \theta_2)(1 - \theta_5)), \\ \eta_3(\theta) &= \log(\theta_3(1 - \theta_6)/(1 - \theta_3)(1 - \theta_7)),\end{aligned}$$

and $\eta_i(\theta) = \log(\theta_i/(1 - \theta_i))$ for $i = 4, \dots, 7$. The cumulant-generating function is the logarithm of the product of the labels along the downward root-to-leaf path coming out of the root vertex $\psi(\theta) = -\log[(1 - \theta_1)(1 - \theta_3)(1 - \theta_7)]$.

3.3. Curved exponential families

As stated in Section 2.2, every discrete Bayesian network has a corresponding staged tree representation. Since for instance the collider graph does not represent a regular exponential family (Koller and Friedman, 2009), staged trees cannot in general be regular exponential families either. They rather form what is called a *curved* exponential family: a submodel of a regular exponential family whose parameter space is a smooth manifold (Efron, 1978).

Theorem 1. *Staged tree models are curved exponential families.*

Proof. By Proposition 1, every staged tree model is a submodel of a regular exponential family. We prove that the natural parameter space is always a smooth manifold of the right dimension by showing that it is the image of a certain linear subspace of Θ under the diffeomorphism η .

Let \mathcal{T}_0 denote a probability tree with parameter space $\Theta_0 = \times_{i=1}^k \Delta_{\kappa_i-1}^\circ$. Let \mathcal{T} denote the same tree graph together with an imposed stage structure and parameter space $\Theta_{\mathcal{T}} = \times_{r \in R} \Delta_{\kappa_r-1}^\circ$ for some index set $R \subseteq \{1, \dots, k\}$. The parameter space of the staged tree model is the kernel of the parameter space of

the saturated model under the linear function $h_{\mathcal{T}} : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^{d_0-d}$ which encodes the $d_0 - d$ identifications of edge labels as

$$h_{\mathcal{T}}(\theta) = (\theta_{ij} - \theta_{st} \mid \text{for all } j = t = 1, \dots, \kappa_j \text{ and all } v_i \text{ and } v_s \text{ in the same stage}) \tag{4}$$

where $d = \sum_{r \in R} (\kappa_r - 1)$ is the number of free parameters in the staged tree, and $d_0 = \sum_{i=1}^k (\kappa_i - 1)$ is the number of free parameters in the probability tree. By construction, $h_{\mathcal{T}}(\theta) = 0$ if and only if θ fulfills the stage constraints in \mathcal{T} . That is, the parameter space of the staged tree model equals the kernel of the map (4), so $h_{\mathcal{T}}^{-1}(0) = \Theta_{\mathcal{T}}$.

Since $h_{\mathcal{T}}$ is surjective, its kernel is d -dimensional. Thus, the parameter space of the staged tree model is a d -dimensional linear space in \mathbb{R}^{d_0} . As a consequence, the natural parameter space, obtained as the image of the diffeomorphism η given in Proposition 1, is a smooth manifold of the same dimension. The claim follows. \square

As an alternative proof strategy for Theorem 1 one may use the characterization of staged tree models as the solution set of a collection of polynomial equations and inequalities provided by Duarte and G\u00f6rgen (2020). Whenever these polynomials do not exhibit any algebraic singularities inside the probability simplex, the model is a curved exponential family: compare the discussion of implicit model representations given in Geiger et al. (2001) and of algebraic exponential families in Drton and Sullivant (2007).

In particular, whilst the equations coding stage constraints in the conditional probabilities as in (4) are always linear, those coding the same constraints formulated in terms of the natural parameters in general are not. In the following proposition we describe these constraints and explain when they are linear.

Proposition 2. *Let \mathcal{T} be a staged tree. For every vertex v_i and outgoing edge indexed by j , let η_{ij} denote the natural parameter of the exponential form derived in Proposition 1, where we additionally define $\eta_{i\kappa_i} = 0$. Let further $\xi_{ij} = \exp(\eta_{ij})$ and $P_{ij} = \sum_r \prod_{ab} \xi_{ab}^{\alpha_{ab}^{(r)}}$ where r ranges over the v_{ij} -to-leaf paths for every fixed double index ij .*

Then two vertices v_i and v_s are in the same stage if and only if they have the same number κ of emanating edges and for all $j = 1, \dots, \kappa$ the equality

$$\eta_{ij} + \log(P_{ij}) + \log(P_{s\kappa}) = \eta_{sj} + \log(P_{sj}) + \log(P_{i\kappa}) \tag{5}$$

is true.

Furthermore, let N_{ij} denote the product of the labels along the downwards-pointing v_{ij} -to-leaf path as in Proposition 1. Then (5) is a linear expression in the η if and only if $P_{ij}P_{s\kappa} = P_{sj}P_{i\kappa}$, or equivalently

$$N_{ij}N_{s\kappa} = N_{sj}N_{i\kappa}. \tag{6}$$

Proof. By definition, $\xi_{ij} = \theta_{ij}N_{ij}/N_i$. For any i and edge j emanating from v_i , the equalities

$$N_{ij}P_{ij} = \sum_r N_{ij} \prod_{ab} \xi_{ab}^{\alpha_{ab}^{(r)}} = \sum_r \prod_{ab} \theta_{ab}^{\alpha_{ab}^{(r)}} = 1$$

are true, where r ranges over the v_{ij} -to-root paths. In particular, the second equality is obtained by applying (2) to the subtree rooted at v_{ij} . The third equality holds because its left-hand side is the sum of all atomic probabilities associated to the subtree rooted at v_{ij} . Now let v_i and v_s be in the same stage. Equivalently, $\theta_{ij} = \theta_{sj}$ for all $j = 1, \dots, \kappa$. This is equivalent to $\xi_{ij}N_i/N_{ij} = \xi_{sj}N_s/N_{sj}$ which in turn is equivalent to $\xi_{ij}N_{i\kappa}/N_{ij} = \xi_{sj}N_{s\kappa}/N_{sj}$ simply because $\theta_{i\kappa} = \theta_{s\kappa}$. Using $N_{ij}P_{ij} = 1$, rearranging, and taking logarithms yields the statement in (5).

As for the second claim, if (6) holds then the log terms in (5) vanish, making it a linear expression in the η . Conversely, let (5) be a linear expression in the η . Then the expression

$$\log(P_{ij}P_{s\kappa}/P_{sj}P_{i\kappa})$$

is also linear in the η . This implies $P_{ij}P_{s\kappa}/P_{sj}P_{i\kappa} = c \prod_{ab} \xi^{\beta_{ab}}$ for some $c, \beta_{ab} \in \mathbb{R}$. Since the left-hand side of this equality is a rational function in the ξ , all β_{ab} must be integers. Indeed, otherwise there would exist a choice of ξ that would make the right-hand side evaluate to an irrational number whilst the left-hand side can only evaluate to rational numbers. So, we obtain an equation of the form

$$P_{ij}P_{s\kappa} \prod_{ab \in A} \xi^{\beta_{ab}} = P_{sj}P_{i\kappa} \prod_{cd \in B} \xi^{\beta_{cd}}$$

for some sets A, B of indices and $\beta_{ab}, \beta_{cd} \in \mathbb{N}$.

Suppose now that one of A and B is nonempty, without loss of generality assume it to be B . Then some ξ_{cd} divides P_{ij} . By the definition of P_{ij} , the only way this is possible is if ξ_{cd} divides all summands of P_{ij} . That is, if all paths of the subtree rooted at v_{ij} (relabelled with $\theta \mapsto \xi$) have the label ξ_{cd} in common. But this is impossible because in a staged tree we assume that each node has at least two children, and thus we can always find a path that avoids the label ξ_{cd} . Hence we see that both A and B must be empty, so $P_{ij}P_{s\kappa} = P_{sj}P_{i\kappa}$. And hence $N_{ij}N_{s\kappa} = N_{sj}N_{i\kappa}$ using $N_{ij}P_{ij} = 1$. \square

Let \mathcal{T} be a staged tree. Proposition 2 shows that the exponential family of \mathcal{T} proposed in this paper is a regular exponential family if and only if the equality of ratios in (6) is true for every pair of vertices v_i and v_s in the same stage and every edge j emanating from these vertices. For this reason, staged trees satisfying this property are of particular interest. We henceforth call these *regular staged trees*.

Equation (6) gives a simple criterion for regularity by comparing certain concatenations of downwards-pointing paths. Algebraically, it can be checked by comparing the monomials $\prod_{ab} \theta_{ab}$ obtained by multiplying all labels in the four downward-pointing paths starting from $v_{ij}, v_{i\kappa}, v_{sj}, v_{s\kappa}$, respectively. Graphically, it can be verified by checking that the concatenation of the first and fourth path on the above list is the same as the concatenation of the second and third, up to a permutation of the edges.

Next, we explore two prominent classes of regular staged trees. The notion of a balanced staged tree, first formalized in Duarte and Ananiadi (2021), is

an important notion for studying the toric geometry of staged trees (Görge, Maraj and Nicklasson, 2021). It makes use of the interpolating polynomials t_{ij} defined by $t_{ij} = \sum_r \prod_{ab} \theta_{ab}^{\alpha_{ab}(r)}$, where r ranges over the v_{ij} -to-leaf paths. This sum is to be read as a formal sum of labels, i.e. without using local sum-to-one conditions. (Including these conditions would imply $t_{ij} = 1$.) A staged tree is *balanced* if for all v_i and v_s in the same stage and all edges j emanating from v_i we have $t_{ij}t_{s\kappa} = t_{i\kappa}t_{sj}$, where again κ denotes the chosen downwards-pointing edge emanating from v_i . This definition for balanced staged trees differs slightly from the one found in the literature but is nevertheless equivalent to it. For a graphical interpretation of this condition, note that $t_{ij}t_{s\kappa}$ is the interpolating polynomial of the staged tree obtained by attaching a copy of the subtree rooted at $v_{s\kappa}$ to each of the leaves of the subtree rooted at v_{ij} . The condition for a balanced tree now is equivalent to this composite tree being statistically equivalent to the composite tree corresponding to the product $t_{i\kappa}t_{sj}$.

Simple staged trees are defined in Collazo, Görge and Smith (2018). A staged tree is *simple* if for all v_i and v_s in the same stage and all edges j emanating from v_i we have $t_{ij} = t_{sj}$.

Proposition 3. *All balanced staged trees are regular. In particular, all simple staged trees are regular.*

Proof. By definition, all simple trees are balanced. Now let \mathcal{T} be a balanced staged tree and v_i, v_s vertices in the same stage. Let j be an edge emanating from v_i . Splitting off the unique downwards-pointing path from v_{ij} , write $t_{ij} = t'_{ij} + N_{ij}$, and likewise for the other pairs of indices. Then

$$\begin{aligned} 0 &= t_{ij}t_{s\kappa} - t_{i\kappa}t_{sj} \\ &= (t'_{ij}t'_{s\kappa} - t'_{sj}t'_{i\kappa} + t'_{ij}N_{s\kappa} - t'_{sj}N_{i\kappa} + t'_{s\kappa}N_{ij} - t'_{i\kappa}N_{sj}) + (N_{ij}N_{s\kappa} - N_{sj}N_{i\kappa}) \\ &=: R + Q \end{aligned}$$

where we define R resp. Q to be the left resp. right outer summand of the middle expression. Consider this expression as a polynomial P in the θ_{ab} labels. Split this set of labels into two sets A and B of labels pointing downwards, resp. not pointing downwards. Thus P can be viewed as a polynomial in $\mathbb{R}[A][B]$. Now, all summands of P in R are divisible by some label in the set B . This means that $Q \in \mathbb{R}[A]$ is the constant term of the polynomial $P \in \mathbb{R}[A][B]$. Thus $P = 0$ implies $Q = 0$. \square

The following converse to Proposition 3 holds for all binary trees with three levels.

Conjecture 1. *All regular staged trees are balanced.*

As an indication for this conjecture, while (6) seems much weaker than the balanced condition for two nodes v_i and v_s , in a regular tree it can be applied recursively everywhere downstream of v_i and v_s , greatly limiting the possible staging structure.

Example 2. Figure 1(a) shows a staged tree representation of the collider Bayesian network $X \rightarrow Z \leftarrow Y$ which is not a regular exponential family. We can see here that indeed in our parametrization, the linear stage identifications $\theta_2 = \theta_3$ do not give rise to linear constraints on the natural parameters η_2 and η_3 . Instead, by Proposition 2,

$$\eta_2 + \log(\xi_4 + 1) + \log(\xi_7 + 1) = \eta_3 + \log(\xi_5 + 1) + \log(\xi_6 + 1)$$

where $\xi_l = \exp(\eta_l)$ for $l = 1, \dots, 7$. This implies the following equation

$$\exp(\eta_2)(1 + \exp(\eta_4))(1 + \exp(\eta_7)) = \exp(\eta_3)(1 + \exp(\eta_5))(1 + \exp(\eta_6))$$

which fully characterises the corresponding curved exponential family in the natural parameters.

The staged tree in Fig. 1(b) however fulfills the linearity criteria in Proposition 2. In particular, the stage constraints $\theta_2 = \theta_3$, $\theta_4 = \theta_6$, and $\theta_5 = \theta_7$ together imply that $(1 - \theta_4)(1 - \theta_7) = (1 - \theta_6)(1 - \theta_5)$ as in (6). They thus give rise to linear (equality) constraints $\eta_2 = \eta_3$ on the natural parameters. The same simplification would occur were v_4 and v_5 , and v_6 and v_7 in the same stage, respectively, rather than v_4 and v_6 , and v_5 and v_7 . These stagings all give rise to balanced and regular trees.

Equations on the natural parameters in a curved exponential family can be highly non-trivial. In this paper we found that for a staged tree they are functions of subgraphs, derived from the conditional independence relations coloured in the tree. Thus, the equations on its exponential family parametrization can be directly read from the graph. This formulation has only been possible thanks to the probability tree's expressiveness of the underlying space of events and of its parametrization.

Analogous results have to the best of the authors' knowledge not been derived in the literature of Bayesian networks. While sufficient statistics for these models are known, natural parameters have not been explicitly computed (e.g. Loh and Wainwright, 2013). In order to translate our formulae into the language of Bayesian networks, first a given directed acyclic graph needs to be transformed into the corresponding staged tree using the algorithm provided by Varando, Carli and Leonelli (2021), and then the staged-tree language can be used to infer a natural parametrization as in Proposition 1. Bayesian networks themselves provide a too-compact representation of the underlying modelling assumptions to be able to directly express their exponential form as a function of the graph. However, Proposition 3 yields the following insight.

Corollary 1. *Let G be a Bayesian network. The exponential family obtained by applying Proposition 1 to the associated staged tree \mathcal{T}_G is regular if G is decomposable or if there exists a topological ordering $(1, \dots, n_G)$ of the nodes of G such that the parents of the $(i + 1)$ th node are a subset of the union of the i th node and its parents for all i .*

Proof. The first statement is well known but also a direct consequence of the fact that \mathcal{T}_G is balanced if and only if G is decomposable (Duarte and Solus,

2021). The second statement corresponds to the definition of G being *simple* given in Leonelli and Varando (2022). These authors prove that G is simple if and only if \mathcal{T}_G is simple. By Proposition 3, simple staged trees form regular exponential families. \square

Acknowledgments

We are grateful to Giovanni Pistone and to Piotr Zwiernik for discussions during the early stages of this project and to Eva Riccomagno for comments on an earlier version of this paper. We also thank the two anonymous referees for their comments which led to a significant improvement of our results.

References

- BARCLAY, L. M., HUTTON, J. L. and SMITH, J. Q. (2013). Refining a Bayesian network using a chain event graph. *International Journal of Approximate Reasoning* **54** 1300–1309. [MR3115418](#)
- BOUTILIER, C., FRIEDMAN, N., GOLDSZMIDT, M. and KOLLER, D. (1996). Context-specific independence in Bayesian networks. In *12th Conference on Uncertainty in Artificial Intelligence (UAI 96)* (E. HORVITZ and F. JENSEN, eds.). *Uncertainty in Artificial Intelligence* 115 – 123. Morgan Kaufmann Publishers Inc., San Francisco. [MR1617129](#)
- CARLI, F., LEONELLI, M., RICCOMAGNO, E. and VARANDO, G. (2020). The R Package `stagedtrees` for Structural Learning of Stratified Staged Trees. Preprint available from [arXiv:2004.06459](#) [stat.ME].
- COLLAZO, R. A., GÖRGEN, C. and SMITH, J. Q. (2018). *Chain Event Graphs. Computer Science & Data Analysis Series*. Chapman & Hall. [MR3752634](#)
- COLLAZO, R. A. and SMITH, J. Q. (2015). A new family of non-local priors for chain event graph model selection. *Bayesian Analysis* **11** 1165 – 1201. [MR3577376](#)
- COWELL, R. G. and SMITH, J. Q. (2014). Causal discovery through MAP selection of stratified chain event graphs. *Electronic Journal of Statistics* **8** 965 – 997. [MR3263109](#)
- DRTON, M. and SULLIVANT, S. (2007). Algebraic statistical models. *Statist. Sinica* **17** 1273–1297. [MR2398596](#)
- DUARTE, E. and ANANIADI, L. (2021). Gröbner bases for staged trees. *Algebraic Statistics* **12** 1–20. [MR4251251](#)
- DUARTE, E. and GÖRGEN, C. (2020). Equations defining probability tree models. *Journal of Symbolic Computation* **99** 127–146. [MR4051945](#)
- DUARTE, E. and SOLUS, L. (2021). A new characterization of discrete decomposable models. Preprint available from [arXiv:2105.05907](#) [math.ST].
- EFRON, B. (1978). The geometry of exponential families. *Ann. Statist.* **6** 362–376. [MR0471152](#)
- FREEMAN, G. and SMITH, J. Q. (2011). Bayesian MAP model selection of chain event graphs. *Journal of Multivariate Analysis* **102** 1152–1165. [MR2805655](#)

- GEIGER, D., HECKERMAN, D., KING, H. and MEEK, C. (2001). Stratified exponential families: graphical models and model selection. *Ann. Statist.* **29** 505 – 529. [MR1863967](#)
- GÖRGEN, C. and SMITH, J. Q. (2018). Equivalence Classes of Staged Trees. *Bernoulli* **24** 2676–2692. [MR3779698](#)
- GÖRGEN, C., MARAJ, A. and NICKLASSON, L. (2021). Staged tree models with toric structure. Preprint available from [arXiv:2107.04516](#) [math.AC].
- HAUGHTON, D. M. A. (1988). On the choice of a model to fit data from an exponential family. *Ann. Statist.* **16** 342–355. [MR0924875](#)
- KASS, R. E. and VOS, P. W. (1997). *Geometrical foundations of asymptotic inference*. *Wiley Series in Probability and Statistics: Probability and Statistics*. John Wiley & Sons, Inc., New York A Wiley-Interscience Publication. [MR1461540](#)
- KEEBLE, C., THWAITES, P. A., BAXTER, P. D., BARBER, S., PARSLow, R. C. and LAW, G. R. (2017). Learning through chain event graphs: the role of maternal factors in childhood type I diabetes. *American Journal of Epidemiology* **186** 1204–1208.
- KOLLER, D. and FRIEDMAN, N. (2009). *Probabilistic graphical models. Adaptive Computation and Machine Learning*. MIT Press, Cambridge, MA Principles and techniques. [MR2778120](#)
- LAURITZEN, S. L. (1996). *Graphical models*. Clarendon Press, Oxford. [MR1419991](#)
- LEONELLI, M. and VARANDO, G. (2022). Structural learning of simple staged trees. Preprint available from [arXiv:2203.04390](#).
- LOH, P.-L. and WAINWRIGHT, M. J. (2013). Structure estimation for discrete graphical models: generalized covariance matrices and their inverses. *Ann. Statist.* **41** 3022–3049. [MR3161456](#)
- SCHWARZ, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics* **6** 461–464. [MR0468014](#)
- SHAFER, G. (1996). *The Art of causal Conjecture. Artificial Management*. MIT Press, Cambridge.
- SILANDER, T. and LEONG, T. Y. (2013). A Dynamic Programming Algorithm for Learning Chain Event Graphs. In *Discovery Science*, (J. Fürnkranz, E. Hüllermeier and T. Higuchi, eds.). *Lecture Notes in Computer Science* **8140** 201–216. Springer Berlin Heidelberg.
- SMITH, J. Q. (2010). *Bayesian decision analysis: principles and practice*. Cambridge University Press, Cambridge; New York. [MR2828346](#)
- VARANDO, G., CARLI, F. and LEONELLI, M. (2021). Staged trees and asymmetry-labeled DAGs. Preprint available from [arXiv:2108.01994](#) [stat.ML].
- WISHART, J. (1949). Cumulants of multivariate multinomial distributions. *Biometrika* **36** 47–58. [MR0033996](#)