The correlation space of Gaussian latent tree models and model selection without fitting

N. Shiers[*1], P. Zwiernik[†2], J. A. D. Aston[‡3] and J. Q. Smith[§1]

[1]Department of Statistics, University of Warwick, Coventry, CV4 7AL, U.K.
[2]Department of Economics and Business, Pompeu Fabra University, 08005 Barcelona, Spain
[3]Statistical Laboratory, University of Cambridge, Cambridge, CB3 0WB, U.K.

April 13, 2016

Abstract

We provide a complete description of possible covariance matrices consistent with a Gaussian latent tree model for any tree. We then present techniques for utilising these constraints to assess whether observed data is compatible with that Gaussian latent tree model. Our method does not require us first to fit such a tree. We demonstrate the usefulness of the inverse-Wishart distribution for performing preliminary assessments of tree-compatibility using semialgebraic constraints. Using results from Drton et al. (2008) we then provide the appropriate moments required for test statistics for assessing adherence to these equality constraints. These are shown to be effective even for small sample sizes and can be easily adjusted to test either the entire model or only certain macrostructures hypothesized within the tree. We illustrate our exploratory tetrad analysis using a linguistic application and our confirmatory tetrad analysis using a biological application.

Keywords— Gaussian; latent tree model; tetrad analysis; tree constraint; tree quartets.

1 Introduction

Modelling with hidden variables is commonly performed within the framework of graphical models (Lauritzen, 1996; Koller & Friedman, 2009). When the observed variables are the leaves of a tree and the unobserved variables are interior nodes then the model is said to be a latent tree model (Choi et al., 2011; Wang et al., 2008). These models are used across disciplines including sociology, biology, and linguistics (Eisenstein et al., 2010; Mourad et al., 2013; Zwiernik, 2016). In the case of continuous data Gaussian latent tree models became a popular choice (Lawrence, 2004).

Standard latent tree model selection techniques often assume a priori that the data generating process is driven by some latent tree model. So in particular the appropriateness of any given tree model is often not assessed in absolute terms but only relative to other tree models. Knowing whether any latent tree model could adequately explain what is observed is pertinent in phylogenetic settings where, for example, the effect of a possible horizontal gene transfer (e.g. Hao & Golding (2008)) makes any underlying latent tree model hypothesis a dubious one.

By characterising the covariance space related to Gaussian latent tree models, we can better assess the suitability of trees or the fit of a particular tree for a data set. In this paper we present the complete description of this model class by relating this to the space of phylogenetic oranges (Engström et al., 2012; Gill et al., 2008; Kim, 2000; Moulton & Steel, 2004). Such a complete description had been known for a simple tree with only

---

[*]n.l.shiers@warwick.ac.uk

[†]piotr.zwiernik@gmail.com

[‡]j.aston@statslab.cam.ac.uk

[§]j.q.smith@warwick.ac.uk

four leaves (see Pearl & Xu (1987, Theorem 2)) or for a star tree (see Bekker & de Leeuw (1987)). For a general tree, only the defining equations have been derived; see Sullivant (2008, Corollary 6.5).

Our method will use the description of Gaussian latent tree models in two scenarios. In the first setting we are interested in whether any latent tree model is a possible explanation for a given data set. In the second situation we fix a latent tree model. In both situations the alternative hypothesis is given by the saturated model. We illustrate these methods in Section 5.4 where we test a previously hypothesized phylogenetic tree applied to certain yeast species. In both examples it is contentious whether the class of phylogenetic trees is appropriate. This is the question that we are able to address directly in our analyses and something that can be done without first fitting the model.

Let $Z = (Z_u)_{u \in U}$ be a random vector whose components are indexed by the vertices of an undirected tree $T = (U, E)$ with edge set $E \subset U \times U$. The tree T induces a Gaussian tree model $N(T)$ for Z, which is a Gaussian graphical model on T (Lauritzen, 1996, Section 5.2). For any two nodes $u, v \in U$, let $\mathrm{ph}(uv)$ denote the set of edges on the unique path between u and v in this tree. Then the model $N(T)$ is the collection of all multivariate normal distributions on $\mathbb{R}^{|U|}$ for which Z_u and Z_v are conditionally independent given a subvector Z_C whenever the set $C \subset U \setminus \{u, v\}$ contains a node on $\mathrm{ph}(uv)$. Note that for three nodes $u, v, w \in U$ the conditional independence of Z_v and Z_w given Z_u is equivalent to $\rho_{vw} = \rho_{uv}\rho_{uw}$. It follows that a normal distribution with correlation matrix $R = (\rho_{uv})$ belongs to $N(T)$ if and only if $\rho_{uv} = \prod_{e \in \mathrm{ph}(uv)} \rho_e$ for all $u, v \in U$, where $\rho_e = \rho_{uv}$ when e is the edge (u, v).

In this paper we study Gaussian latent tree models where only the observed random variables correspond to the tree's leaves. We henceforth denote the set of leaves of this tree by V. A typical such evolutionary tree, one of Romance languages, is displayed below in Fig. 1 where the observable, extant languages are represented as its leaves.
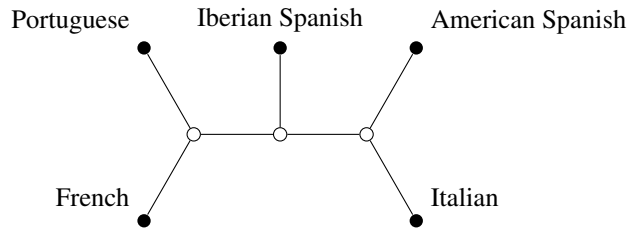

Figure 1: Quintet tree T_5 relating five Romance languages.

Definition 1. *The Gaussian latent tree model $M(T)$ for the subvector $X = (Z_v)_{v \in V}$ is the set of all V-marginal distributions of the distributions in $N(T)$, where the V-marginal distributions are those associated with leaf variables.*

The parameterization of $M(T)$ is induced from the parameterization of $N(T)$ and given by

$$\rho_{ij} = \prod_{e \in \mathrm{ph}(ij)} \rho_e \tag{1}$$

for all $i, j \in V$. As the variances σ_{uu} for $u \in U \setminus V$ never appear in this parameterization, without loss of generality, we can assume they are equal to 1.

2 Semialgebraic description of the latent tree model

2.1 Tree metrics and phylogenetic oranges

Let $T = (U, E)$ be a tree with leaf set $V \subseteq U$. Associate to each edge a non-negative number d_e, which we interpret as the length of this edge. Then for any two leaves $i, j \in V$ we can compute the distance between them as $d_{ij} = \sum_{e \in \mathrm{ph}(ij)} d_e$. It is easy to check that a collection of such distances for all pairs $u, v \in V$ forms a metric. The set of all metrics that arise in this way for all T with leaves labelled by V is called the space of tree metrics. We recall the following result.

Theorem 1 (Buneman (1974)). *A collection of positive numbers d_{ij} for $i, j \in V$ forms a tree metric if and only if for all (not necessarily distinct) $i, j, k, l \in V$ we have*

$$\max(d_{ik} + d_{jl}, d_{il} + d_{jk}) \geq d_{ij} + d_{kl}.$$

Equivalently, for any three sums $d_{ik} + d_{jl}$, $d_{il} + d_{jk}$, $d_{ij} + d_{kl}$ two are equal and not less than the third. Moreover, if the above inequalities hold, then generically T is uniquely identified.

In the above theorem the term generically means that the statement holds outside a set of measure zero corresponding to the vanishing of some edge lengths d_e. We note that a more precise statement is also possible if we allow semi-labelled trees, see Semple & Steel (2003, Section 7). A careful analysis shows that this generic tree is always a binary tree, i.e. a tree with all its inner nodes of degree three. The usual triangle inequality follows from setting i, j, k distinct and $k = l$ in Theorem 1, which in turn implies that every tree metric is a metric on V.

Corollary 1. *The space of tree metrics on a fixed tree T is given as a set of all metrics on V satisfying: for any four distinct leaves i, j, k, l such that $\mathrm{ph}(i, j) \cap \mathrm{ph}(k, l) = \emptyset$, we have*

$$d_{ik} + d_{jl} \quad = \quad d_{il} + d_{jk} \quad \geq \quad d_{ij} + d_{kl}.$$

We emphasize that $\mathrm{ph}(i, j)$ is the set of edges and hence, for example, for a star tree any four leaves i, j, k, l that satisfy $\mathrm{ph}(i, j) \cap \mathrm{ph}(k, l) = \emptyset$. The condition $\mathrm{ph}(i, j) \cap \mathrm{ph}(k, l) = \emptyset$ implies that the induced subtree over i, j, k, l, that is, the smallest connected subgraph of T containing i, j, k, l, looks like a quartet tree in Figure 2. This also explains the conditions of Corollary 1.
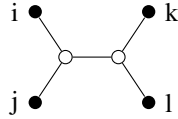

Figure 2: A quartet tree $ij|kl$

Another closely related space defined over a tree is the space of phylogenetic oranges; see e.g. Kim (2000); Moulton & Steel (2004). For a fixed tree T this is given by the same parameterization (1) as the Gaussian latent tree model but where in addition the edge correlations ρ_e are non-negative. The set of all points in $\mathbb{R}^{m(m-1)/2}$ that arise in this way is denoted by $\mathrm{PO}(T)$ and it forms a toric cube as defined in Engström et al. (2012). The union of all $\mathrm{PO}(T)$ is denoted by $\mathrm{PO}(V)$.

Denote by $\mathrm{PO}_+(T)$ and $\mathrm{PO}_+(V)$ the subsets of $\mathrm{PO}(T)$ and $\mathrm{PO}(V)$ respectively, where all coordinates are assumed to be strictly positive. This implies in particular that the corresponding edge correlations ρ_e must be strictly positive. The space of tree metrics on a fixed tree T is isomorphic to $\mathrm{PO}_+(T)$, with the isomorphism given by $d_{ij} = -\log(\rho_{ij})$.

Theorem 2. *Let $R = (\rho_{ij})_{i,j \in V}$ and suppose that $\rho_{ij} \geq 0$ for all $i, j \in V$. The following two statements hold:*

(1) $R \in \mathrm{PO}(V)$ if and only if for every four not necessarily distinct elements i, j, k, l in V at least two out of three products $\rho_{ik}\rho_{jl}$, $\rho_{il}\rho_{jk}$, $\rho_{ij}\rho_{kl}$ are equal and less than or equal to the third. Moreover, if this holds then T with the property $R \in \mathrm{PO}(T)$ is generically identified uniquely.

(2) For a fixed T, the space $\mathrm{PO}(T)$ has dimension $|E|$. This is described by the following set of constraints. For any four distinct elements i, j, k, l of V such that $\mathrm{ph}(i, j) \cap \mathrm{ph}(k, l) = \emptyset$, we have that

$$\rho_{ik}\rho_{jl} \quad = \quad \rho_{il}\rho_{jk} \quad \leq \quad \rho_{ij}\rho_{kl}. \tag{2}$$

Moreover, for any three distinct leaves i, j, k we have the triangle inequality $\rho_{ij}\rho_{ik} \leq \rho_{jk}$.

2.2 Latent tree models and phylogenetic oranges

We are now ready to derive the semialgebraic description of the model $M(T)$. Let $\mathcal{S}_+(V)$ denote the space of all symmetric positive definite $|V| \times |V|$-matrices.
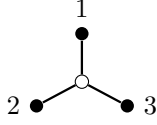
Figure 3: Tripod tree.

Theorem 3. *Let T be a tree and let $R = [\rho_{ij}] \in \mathcal{S}_+(V)$ be a correlation matrix. Then $R \in M(T)$ if and only if $R' = [|\rho_{ij}|] \in \mathrm{PO}(T)$ and $\rho_{ij}\rho_{ik}\rho_{jk} \geq 0$ for any three distinct $i, j, k \in V$.*

The proof is given in the appendix.

Example 1. *Let T be the tripod tree in Fig.3. The space of correlation matrices in $M(T)$ is described by*

$$\rho_{12}\rho_{13}\rho_{23} \geq 0, \quad |\rho_{12}\rho_{13}| \leq |\rho_{23}|, \quad |\rho_{12}\rho_{23}| \leq |\rho_{13}|, \quad |\rho_{13}\rho_{23}| \leq |\rho_{12}|.$$

If $\rho_{12}, \rho_{13}, \rho_{23} \geq 0$ then by Theorem 2(2) the space described by the above inequalities corresponds to $\mathrm{PO}(T)$. There are three other sign patterns for $\rho_{12}, \rho_{13}, \rho_{23}$ that ensure that $\rho_{12}\rho_{13}\rho_{23} \geq 0$. For every such pattern we obtain a copy of $\mathrm{PO}(T)$. Quite remarkably, the space of the correlation matrices in $M(T)$ looks exactly like the three-dimensional slice of the corresponding binary latent class model; see Allman et al. (2015, Figure 1). It is interesting to note that such constraints cannot, in general, be neglected. For example, simple calculations show that the ratio of the volume of $M(T)$ to the volume of all 3×3 correlation matrices is only $\frac{2}{\pi^2} \approx 0.2$.

Based on Theorem 2(2) and Theorem 3 we formulate the following result.

Proposition 1. *If T is a fixed tree then the space $M(T)$ has dimension $|V| + |E|$. Let Σ be a covariance matrix with no zeros. Then $\Sigma \in M(T)$ if and only if for any three distinct leaves i, j, k*

$$(\sigma_{kk}\sigma_{ij} - \sigma_{ik}\sigma_{jk})(\sigma_{jj}\sigma_{ik} - \sigma_{ij}\sigma_{jk})(\sigma_{ii}\sigma_{jk} - \sigma_{ij}\sigma_{ik}) \geq 0, \tag{3}$$

and for any four distinct elements i, j, k, l of V such that $\mathrm{ph}(i,j) \cap \mathrm{ph}(k,l) = \emptyset$

$$\frac{\sigma_{ik}\sigma_{jl}}{\sigma_{ij}\sigma_{kl}} \quad = \quad \frac{\sigma_{il}\sigma_{jk}}{\sigma_{ij}\sigma_{kl}} \quad \leq \quad 1. \tag{4}$$

This full algebraic and semialgebraic description can be viewed as a generalization from star trees to general trees of the main results in Bekker & de Leeuw (1987); Pearl & Xu (1987). An analogous description of the second order moments for binary latent tree models was given in Zwiernik & Smith (2011). The similarity of both descriptions comes from the fact that the parameterization of correlations in the binary latent tree model is precisely (1); see Zwiernik & Smith (2011, Lemma 4.1).

2.3 Necessary constraints for non-Gaussian tree models

Fix an inner node r in a tree T and direct all edges away from r to obtain a rooted tree T^r. Set

$$Z_r = \epsilon_r \quad \text{and} \quad Z_v = \lambda_v Z_u + \epsilon_v \quad \text{for} \quad u \to v \text{ in } T^r, \tag{5}$$

where $\lambda_v \in \mathbb{R}$ and ϵ_v for $v \in U$ are independent random variables with mean zero. It is known that the vector Z follows the Gaussian tree model on T if and only if all ϵ_v are Gaussian. Its subvector X corresponding to the leaf nodes will then follow the latent tree model on T.

It is natural to ask what happens if ϵ_v are not jointly Gaussian. In a nonparametric setting we could instead assume that ϵ_v have distributions in the family of all univariate distributions with mean zero and finite variance. In this case it is easily shown that all second order moments of Z exist, $\lambda_v = \mathrm{cov}(Z_u, Z_v)/\mathrm{var}(Z_u)$, and the correlations satisfy the parameterization in (1). In particular, we have the following result that provides a set of necessary constraints on the correlations of a non-Gaussian tree model.

Theorem 4. *Suppose that Z is a random vector satisfying the recursive equations in (5) for a rooted tree. If all ϵ_v have finite variance, then the correlation matrix of X must satisfy the constraints of the Gaussian latent tree model.*

So it appears that our results, whilst focused on Gaussian modes, apply to and could be extended beyond this setting.

4

3 Utilising semialgebraic constraints

We now describe how the semialgebraic constraints can be used more formally to give an indication of tree-compatibility. Here the constraints in (3), which hold for every tree topology, will be called tree-compatibility constraints. A test based on these constraints can be used as an effective preliminary assessment tool to inform whether it is legitimate to proceed to a more complex tetrad analysis. For a fixed T we can further extend our test by including the inequality constraints in (4). The resulting constraints are called T-compatibility constraints. A test of fit based on such constraints is called a tree-compatibility or T-compatibility test as appropriate.

A straightforward but effective assessment of T-compatibility constraints can be obtained from the posterior probabilities by applying an inverse-Wishart prior on the sample covariance. More precisely, if $\hat{\Sigma}$ is a sample covariance matrix based on a sample X of size n from $\mathcal{N}_m(0, C)$, then the estimated scatter matrix is calculated as $S = n\hat{\Sigma} = XX^T$ and it is well known that the scatter matrix is Wishart distributed $S \sim \mathcal{W}_m(n, C)$ (Wishart, 1928). A common prior distribution for unknown covariance C is the inverse-Wishart $\mathcal{W}_m^{-1}(n_0, C_0)$, e.g. Gelman et al. (2013); Carlin & Louis (2008); Roverato (2002). The inverse-Wishart is a conjugate prior and so the posterior density $p(C \mid X)$ is inverse-Wishart $\mathcal{W}_m^{-1}(n_0 + n, C_0 + S)$. As in Roverato (2002), for C_0 the identity matrix $I_{|V|}$ can be used and by letting $n_0 = m$ ensure that the prior density is well defined. Then $C \mid X$ can be sampled with each draw being translated to a correlation and then tested against the constraints. After N such draws from the posterior distribution an estimate of the posterior probability that C satisfies the positivity constraint can be obtained. Of course other choices of families of priors could be chosen instead (for example the scaled inverse-Wishart (O'Malley & Zaslavsky, 2008)) or we could use a strategy that models correlation and covariance separately (Barnard et al., 2000). However, these alternatives bring additional computational cost and complexity. Alternatively, it may be possible to adapt the work on inequality-constrained hypotheses to this framework using Bayesian methods, see Van de Schoot et al. (2012); Gu et al. (2014); Gardner et al. (2014).

In Example 1, an estimate of the probability of C satisfying the semialgebraic structure of $M(T)$ can be constructed using indicator functions. For each draw l from the relevant inverse-Wishart posterior distribution for $\hat{\Sigma}$, the following identity is evaluated:

$$r_{123}^l(\hat{\Sigma}) = \mathbb{1}\{(\tilde{\sigma}_{33}\tilde{\sigma}_{12} - \tilde{\sigma}_{13}\tilde{\sigma}_{23})(\tilde{\sigma}_{22}\tilde{\sigma}_{13} - \tilde{\sigma}_{12}\tilde{\sigma}_{23})(\tilde{\sigma}_{11}\tilde{\sigma}_{23} - \tilde{\sigma}_{12}\tilde{\sigma}_{13}) \geq 0\} \tag{6}$$

where $\tilde{\sigma}_{ij}, i, j = 1, 2, 3$ are the covariances corresponding to covariance draw l of the posterior, the index l being dropped to keep the notation clean. The posterior probability of tree-compatibility is thus estimated using:

$$R_{123}(\hat{\Sigma}) = \frac{1}{N}\sum_{l=1}^{N} r_{123}^l(\hat{\Sigma}) \tag{7}$$

For a tree with four variables such that $\mathrm{ph}(1, 2)$ and $\mathrm{ph}(3, 4)$ do not intersect, the final test of the inequality constraints is then:

$$R_{12|34}(\hat{\Sigma}) = \frac{1}{N}\sum_{l=1}^{N} \mathbb{1}\{\tilde{\sigma}_{14}\tilde{\sigma}_{23} - \tilde{\sigma}_{12}\tilde{\sigma}_{34} \leq 0\}\mathbb{1}\{\tilde{\sigma}_{13}\tilde{\sigma}_{24} - \tilde{\sigma}_{12}\tilde{\sigma}_{34} \leq 0\} \prod_{\substack{1 \leq i < j \\ < k \leq 4}} r_{ijk}^l(\hat{\Sigma}) \tag{8}$$

These sampling approaches do not extend to the algebraic constraints because the set of draws from the posterior satisfying an equality constraint will have zero probability. Thus an alternative approach is needed that uses sample distributions of the minors of a covariance matrix.

4 The sample distribution of algebraic constraints

From the previous section and Theorem 2(2), the signs of tetrad constraints $\sigma_{ik}\sigma_{jl} - \sigma_{il}\sigma_{jk}$ and other quadratic binomials of the form $\sigma_{ii}\sigma_{jk} - \sigma_{ij}\sigma_{ik}$ provide essential information about whether a Gaussian distribution lies in $M(T)$. This type of constraints can be realized as minors of the covariance matrix Σ, that is

$$\det(\Sigma_{ij,kl}), \qquad \det(\Sigma_{ij,ik}), \tag{9}$$

where $\Sigma_{ij,kl}$ denotes the 2×2 sub-matrix of Σ with rows i and j and columns k and l. Let $\left\{ {m \atop 2} \right\}$ denote the set of all subsets of $\{1, \ldots, m\}$ of cardinality two. We now propose the following estimator of the value of $\det(C_{I,J})$ for $I, J \in \left\{ {m \atop 2} \right\}$

$$Q_{I,J} \quad = \quad \frac{1}{n(n-1)} \det(S_{I,J}). \tag{10}$$

We note from Drton et al. (2008, Corollary 4.2) that $Q_{I,J}$ is an unbiased estimator of $\det(C_{I,J})$.

In what follows we provide the covariances between different $Q_{I,J}$. It is convenient to introduce the following notation. For an $m \times m$ matrix A let $A^{(2)}$ denote the matrix with rows and columns indexed by elements $\left\{ {m \atop 2} \right\}$ whose (I, J)-th element is the corresponding minor $\det(A_{I,J})$. With this notation, the matrix, whose elements are the estimators $Q_{I,J}$, is $S^{(2)}/(n(n-1))$.

Sadly, there is no simple explicit formula for covariances of various 2-minors. However, these can be computed if the true distribution C is known. From Drton et al. (2008, Proposition 3.3)

$$\text{cov}(S^{(2)}) = \{(C^{1/2})^{(2)} \otimes (C^{1/2})^{(2)}\}\{\text{cov}(W^{(2)})\}\{(C^{1/2})^{(2)} \otimes (C^{1/2})^{(2)}\}, \tag{11}$$

where W has standard Wishart distribution $\mathcal{W}_m(n, I)$ and $\otimes$ is the Kronecker product.

In the rest of this section we provide a complete description of the covariance matrix $\text{cov}(W^{(2)})$. Our discussion follows Drton et al. (2008, Example 4.6). This gives the same derivation for the case $m = 4$. We show below that the generalization to $m \geq 4$ is straightforward.

The matrix $\text{cov}(W^{(2)})$ has many symmetries that we want to exploit. Note that for all $I, J \in \left\{ {m \atop 2} \right\}$, $\det W_{I,J} = \det W_{J,I}$ and hence

$$\text{cov}\{\det(W_{I,J}), \det(W_{K,L})\} = \text{cov}\{\det(W_{J,I}), \det(W_{K,L})\} = \text{cov}\{\det(W_{K,L}), \det(W_{I,J})\}.$$

We can therefore, without loss, consider only unordered pairs of sets (I, J), where $I = \{i, j\}$ and $J = \{k, l\}$ with $i < j$, $k < l$ and either $i < k$ or $i = k$ and $j \leq l$.

Let $A \Delta B = (A \setminus B) \cup (B \setminus A)$ be the symmetric difference of two sets. We split the rows and the columns of $\text{cov}(W^{(2)})$ into blocks according to the value of $I \Delta J$ and $K \Delta L$. With this convention, by Drton et al. (2008, Corollary 4.2 and Proposition 3.4), $\text{cov}(W^{(2)})$ is a block-diagonal matrix. Therefore, it is enough to describe its diagonal blocks. Since $|I \Delta J| \in \{0, 2, 4\}$, we have three types of blocks. We first describe the block corresponding to $I \Delta J = K \Delta L = \emptyset$, or equivalently $I = J$, $K = L$. This block forms a $\binom{m}{2} \times \binom{m}{2}$-matrix that satisfies:

$$\text{cov}\{\det(W_{I,I}), \det(W_{K,K})\} = \begin{cases} 0 & |I \cap K| = 0 \\ 2n(n-1)^2 & |I \cap K| = 1 \\ 2n(2n+1)(n-1) & I = K. \end{cases}$$

We now have $\binom{m}{2}$ blocks corresponding to $I \Delta J = K \Delta L = \{i, j\}$ for $1 \leq i < j \leq m$. Every such block is an $(m-2) \times (m-2)$-matrix, where $I = \{i, k\}$, $J = \{j, k\}$, $K = \{i, l\}$, $L = \{j, l\}$ for some $k \leq l \in \{1, \ldots, m\} \setminus \{i, j\}$. All of these matrices have two types of elements. The diagonal entries ($k = l$) are equal to $n(n+2)(n-1)$. The off-diagonal elements ($k < l$), up to a sign, are $n(n-1)^2$. The sign depends on the relative order of i, j, k, l. By Drton et al. (2008, Theorem 4.5), the sign is positive if $k < i < j < l$. Now a simple sign analysis shows that the sign is negative only if either $i < k < j < l$ or $k < i < l < j$. This yields that:

$$\text{cov}\{\det(W_{ik,jk}), \det(W_{il,jl})\} = \begin{cases} n(n+2)(n-1) & k = l \\ -n(n-1)^2 & i < k < j < l \text{ or } k < i < l < j \\ n(n-1)^2 & \text{otherwise.} \end{cases}$$

Finally, there are $\binom{m}{4}$ blocks corresponding to $I \Delta J = \{i, j, k, l\}$, where $1 \leq i < j < k < l \leq m$. Each such block is a 3×3 matrix of the form

$$\begin{array}{ccc} ij,kl & ik,jl & il,jk \end{array}$$
$$\begin{bmatrix} a & b & -b \\ \cdot & a & b \\ \cdot & \cdot & a \end{bmatrix}$$

where $a = 2n(n-1)$ and $b = n(n-1)$.

5   Quartets and applications of tetrad analyses

5.1   The method of quartets

For any four distinct leaves $i, j, k, l \in V$ we say that $q_{ij,kl} = ij|kl$ forms a quartet of T if the paths $\mathrm{ph}(i,j)$ and $\mathrm{ph}(k,l)$ are disjoint, c.f. Figure 2. A binary tree T displays the set of quartets $\mathcal{Q}$ if each quartet $q \in \mathcal{Q}$ is a quartet of T. A set of quartets $\mathcal{Q}$ is said to determine T if T displays $\mathcal{Q}$ and T is the unique tree displayed by $\mathcal{Q}$ (Semple & Steel, 2003); the set of all quartets displayed by T is denoted by $\mathcal{Q}_T$. Quartets can be considered as fundamental components of binary trees; see Dress et al. (2012) for more details. A set $\mathcal{Q}_T$ is said to be minimal if there exists no element $q \in \mathcal{Q}_T$ such that $\mathcal{Q}_T \setminus \{q\}$ determines T. Grünewald et al. (2008, Theorem 2.4) provides the minimum size of any $\mathcal{Q}_T$ (i.e. the size of the smallest minimal defining quartet set), which for a binary tree is just the number of internal edges of T. Furthermore, Semple & Steel (2003, Theorem 6.8.8) provide a quick method for constructing minimal defining sets of quartets that define binary phylogenetic trees.

Let $V \subset U$ be such that $V = \{i, j, k, l\}$, where these elements are distinct. Consider three random variables $Q_{ik,jl}$, $Q_{il,jk}$ and $Q_{ij,kl}$ as defined in (10). By Theorem 2, if a tree model holds, then the mean of one of the three will be zero and the other two means will be equal up to sign. So these $Q_{I,J}$ can be used to test the algebraic constraints in Proposition 1.

Here we focus on testing the vanishing tetrads, i.e. testing whether the quartet $q_{ij,kl}$ is displayed in T given the data. To test a particular binary tree T, a set $\mathcal{Q}_T$ is required, i.e. a set of quartets $\mathcal{Q}$ that determines T. The number of edges of T is $2m - 3$ and so the Gaussian latent tree model on T has codimension $\binom{m}{2} - (2m - 3)$. This means that to test a model, we need to work with quartet systems $\mathcal{Q}_T$ of size quadratic in m. On the other hand if we believe that the data come from a latent tree model, then to only find the corresponding tree T we can work with any minimal quartet system determining T, and these are of size $m - 3$ (the number of internal edges). This makes a big difference for larger trees.

In practice, one may wish to select $\mathcal{Q}_T$ such that it is minimal (of size $\binom{m}{2} - (2m - 3)$), i.e. it contains no redundant quartets, because otherwise the covariance of minors matrix may be close to being singular; see Bollen & Ting (1993). However, there may not always be an obvious reason for selecting one minimal defining quartet set $\mathcal{Q}_T$ over another. In such cases one approach is to randomly select a number of sets to assess the robustness of the results; see Bollen & Ting (1993). For each $q_{ij,kl} \in \mathcal{Q}_T$ consider the corresponding $Q_{ij,kl}$ as in (10) and define $Q_T = [Q_{ij,kl}]$ to be the vector of these $Q_{ij,kl}$. We write $\hat{Q}_{ij,kl}$ for the sample means of the observations of $\hat{Q}_{ij,kl}$. Since $\hat{Q}_T$ is a consistent estimator of Q_T (see Drton et al. (2007)), as the sample size n tends to infinity any tree T is uniquely identified by the i, j, k, l such that $E(\hat{Q}_{ij,kl}) = 0$.

To standardize the data we use the sample covariance matrix $\hat{\Sigma}_{Q_T}$ which has dimension $p = |Q_T|$. Alternatively we can use its proxy $\tilde{\Sigma}_{Q_T}$. This can then be obtained by recycling $\mathrm{cov}(W^{(2)})$ computed in Section 4 and using (11) substituting C for the sample covariance of original variables $\hat{\Sigma}$. The matrix $\tilde{\Sigma}_{Q_T}$ can be obtained much more efficiently than $\hat{\Sigma}_{Q_T}$. An appropriate simultaneous test statistic (12) is provided in Bollen & Ting (1993),

$$\mathcal{T} \;=\; \hat{Q}_T^t \, \hat{\Sigma}_{Q_T}^{-1} \, \hat{Q}_T, \tag{12}$$

where A^t is the transpose of A. Given $\mathcal{T}$ is constructed with p algebraically independent quartets, the asymptotic distribution of this test statistic is χ^2-distribution with p degrees of freedom. Compare (12) with Bollen & Ting (1993, (20)) where their Σ_{tt} is the covariance of $\sqrt{n}\hat{Q}_T$. Here the sample size n is incorporated implicitly through $\hat{\Sigma}_{Q_T}^{-1}$ so (12) provides a significance test for the equality constraints in (4), where the required moments of $Q_{I,J}$ are given in Section 4. This provides a quick method for assessing whether a Gaussian data set appears consistent with the algebraic constraints associated with any tree model.

In deriving the asymptotic distribution in (12) we implicitly assume that the true covariance matrix is a sufficiently regular point of the given tree model. In practice, it is enough to assume that the true covariance matrix contains no zeros; see Section 5 in Drton et al. (2016b) and Drton et al. (2016a).

Hypothesis testing for vanishing tetrads can be used for both confirmatory tetrad analysis and for exploratory tetrad analysis. There are many algorithms for obtaining candidate trees, for instance see Junker & Schreiber (2011); Sung (2009) for surveys of methods. However, often there is no way to assess the suitability of the finally chosen tree. Confirmatory tetrad analysis takes a candidate tree and provides an absolute rather than relative value as to how well the data supports the purported tree.

In the case of a large tree it is infeasible to test all quartets at once. On the other hand, it is straightforward and very stable to test single quartets or a small subset of them. One advantage of this approach is that it allows us to identify easily certain macrostructures of the tree which may lead to more robust techniques for finding the underlying tree. We now illustrate confirmatory and exploratory techniques both for simulated data and some linguistics data sets.

5.2 Basic simulations for the method of quartets

In this section we provide a basic analysis of the methods discussed in the previous section. The only difference from the previous applications of this method in other contexts is that in (12) we explicitly replaced the sample covariance of the tetrads with $\tilde{\Sigma}_{Q_T}$ as explained in Section 5.1. The data in our simulations come from the quintet tree model 12|5|34; c.f. Figure 1.

We first randomly choose the true covariance matrix C by sampling the edge correlations uniformly from the interval $[1/2, 1]$. Given this random true covariance matrix, we can now repeat the following evaluation procedure 10,000 times. We sample $n = 60$ (5 times the dimension of the model) points from the given distribution C. In this scenario, standard packages that might be used to find the maximum likelihood estimate, such as the sem package (Fox et al., 2014), are unstable. On the other hand any set of quartets can be easily tested, and this does not even require any fitting of the model. Moreover, the sample distribution of the test statistic is already very close to the asymptotic distribution. Quite surprisingly this proximity remains true even when the sample size is only about twice the dimension of the model. Of course, in this case, the power of the test will be much lower.

In Figure 5 we show the simulated values of test statistics of the form (12) compared with their theoretical asymptotic distributions. Figure 5(a) depicts the statistic built on a single tetrad constraint for the quartet 12|34. This constraint holds for the data generating distribution and therefore the test statistic is expected to have asymptotic χ^2-distribution with one degree of freedom. We see that the histogram is very close to the theoretical distribution. For comparison, in Figure 5(b) we show the sample distribution of the same test statistic for the quartet 13|24. This constraint does not hold for the data generating distribution and we see that the sample distribution of the corresponding test statistic is very far from χ_1^2. The test statistic can be easily set up for any subset of quartets. In Figure 5(c) we plot the test statistic to test two quartets 12|35 and 15|34. This is the minimal set of quartets that identifies the quintet tree 12|5|34. This means that these two particular quartet will not be simultaneously satisfied for any other tree model . Again, the sample distribution lies very close to the asymptotic distribution, which in this case is χ_2^2.

In Figure 5(d) we test simultaneously a minimal set of quartets defining the quintet tree model; these are: 12|34, 12|35, 15|34. In this case the sample distribution of the test statistic also lies close to χ_3^2 with a slightly smaller variance. The reason for that is that the true distribution is closer to a mixture of χ^2-distributions. As a result, the test based on our statistic is typically more conservative. To obtain a better understanding of its performance we compare it with the structural expectation-maximization algorithm Friedman et al. (2002) as applied to Gaussian latent tree models. This algorithm tries to find the tree that gives the maximum value of the likelihood function. However, like the standard expectation-maximization algorithm, it often gets stuck in a local maximum. In our simulations we generated 100 data sets from the given quintet model. If the sample size $n = 60$, then for our particular choice of a correlation matrix with all edge correlations equal to 0.7, we obtained the correct tree only 68 out of 100 times. On the other hand, our tetrad method always confirms the correct tree on any significance level smaller than 0.1. If $n = 200$ then the structural expectation-maximization algorithm was correct 99 out of 100 times, and again our quartet method was always correct. We emphasize that in a less ideal situation, for example, when some edge correlations are small, or in the presence of some partial misspecification, the structural expectation-maximization algorithm will perform poorly because the likelihood function is less stable. In contrast, our computations show that the quartet method tends to be much more robust.

5.3 Exploratory tetrad analysis example: linguistics

Consider now the linguistic data set from Shiers et al. (2014). This comprises phonetic functional spectrogram data from five Romance languages: French, Italian, Portuguese, and two forms of Spanish, namely American and Iberian. Acoustic data have provided new insights into language developments (e.g. Bouchard-Côté et al. (2013), Aston et al. (2010)). Here the evolutionary dependencies between spoken numbers is studied with each extant

language treated as a leaf vertex. The high dimensional spectrogram data is projected from 8100 dimensions to 15 dimensions using a variation of canonical variate analysis, the full details of which can be found in Shiers et al. (2014). Each of the 15 canonical components projects the mean word data to obtain 15 new data sets referred to as canonical scores. Each canonical component accounts for a particular combination of phonetic variation and each set of canonical scores is considered independently. This gives us the flexibility to hypothesize different evolutionary relationships for different aspects of the speech. For each set of canonical scores a 5×5 covariance matrix is calculated between the five languages. Royston's multivariate normality test (Royston, 1983) does not reject Gaussianity at the 0.01 level for any of these 15 sets of scores.

We sampled 10^5 covariance matrices from the inverse-Wishart posterior for each of the sample covariances $\hat{\Sigma}_1, \ldots, \hat{\Sigma}_{15}$. We then performed a tree-compatibility test with respect to the positivity constraint implied by the triangle inequalities in Theorem 2(2) for each canonical component. We identify four such components, the first, fourth, sixth, and second, with high posterior probabilities, respectively 1, 0.89, 0.77 and 0.74, which warrant further investigation.

Considering the quintet tree in Fig. 1, there are 15 different labelled binary trees to test. In order to test a particular configuration of labels we construct a set of minimal defining quartets $\mathcal{Q}$ for the quintet tree as referenced in Section 5.1; in the case of the quintet tree this smallest minimal set is two.

For each of the four dimensions of interest, using the sampling distributions given in Section 4 and the test statistic (12) with two degrees of freedom, a p-value can be calculated for each of the 15 non-isomorphic trees with languages as leaves. To retain an overall significance rate of less than 0.05 a Bonferroni correction (Dunn, 1961) is applied such that the significance level is set at $0.05/15 \approx 0.0033$ per test. If more than one tree is not rejected then the candidate tree proposed by exploratory tetrad analysis is that with the highest p-value. We find that multiple trees exceed the threshold for all four components. The highest p-values for the first, second, fourth and sixth components were 0.524, 0.960, 0.775 and 0.902 respectively relating to the candidate trees: 12|4|35, 13|5|24 ,14|2|35, and 23|1|45 respectively with coding 1 = French, 2 = Italian, 3 = Portuguese, 4 = American Spanish, 5 = Iberian Spanish.

For illustration we focus on the candidate tree for the second component, which is displayed in Fig 1. It is known from the analysis and expert interpretation in Shiers et al. (2014) that this component is likely to relate to variation in vowel sounds, nasality, and the lip rounding of language speakers. By isolating these phonetic features and identifying an accompany tree that fits the data we can gain insights which may have otherwise been obscured. For example, from this particular analysis we could hypothesize that the differences in nasality of Italian and French evolved independently conditional on the common ancestor of Spanish and Portuguese. In combination with expert judgement, such statements can provide good starting points for further analyses of these features in relation to a specified tree.

5.4 Confirmatory tetrad analysis example: biology

We next consider a data set consisting of growth curves for five yeast species each observed in the same 96 environments, each species with at least two replicates. The growth was recorded at approximately six minute intervals over a period of just over 26 hours. These species have been studied before (Marcet-Houben & Gabaldón, 2009) and a phylogeny has been suggested as in Fig. 4. However, Libkind et al. (2011) hypothesize that yeast species S. bayanus is a hybrid involving S. cerevisiae. This alternative hypothesis would violate the tree assumption. Previous research has indicated that for studied yeast species there is positive correlation between growth-related phenotypic variation and genotypic phylogenetic relationships, e.g. Liti et al. (2009); Warringer et al. (2011). Thus, this leads us to consider the yeast growth-curve data to investigate evolutionary questions. We carried out a confirmatory tetrad analysis to assess whether the proposed tree structure in Marcet-Houben & Gabaldón (2009) was reflected in any aspects of the growth data.

To pre-process the data, a smoothed cubic spline basis was fitted to each growth vector resulting in a set of functional data objects which were then regularly evaluated to obtain comparable discretized representations. Mean vectors were then calculated for each species and environment and then these were standardized to remove mean environmental effects. We then performed a principal component analysis across species to identify the core variability of the growth curves. Note that the first four components account for over 99% of variability. More detailed analysis, not reported here, can help interpret these components. For example, the first component relates only to growth variation in hours 10 to 26, whereas the second component relates to growth variation peaking at 12 hours with opposite growth variation from 18 hours onwards.

For each of the mean species projections in these dimensions, the sample covariance matrix was constructed. As a first step, the inverse-Wishart approach specified in (7) was implemented. Recall that the tripod constraints are tree-compatibility constraints and thus require no tailoring to a specific T. Hence, these can be utilized very simply to narrow the list of components to test as part of a confirmatory tetrad analysis. The tree-compatibility for the first four components were 31%, 2%, 18%, and 3% respectively. Thus, we consider the first and third components worth investigating further via confirmatory tetrad analysis for T-compatibility.
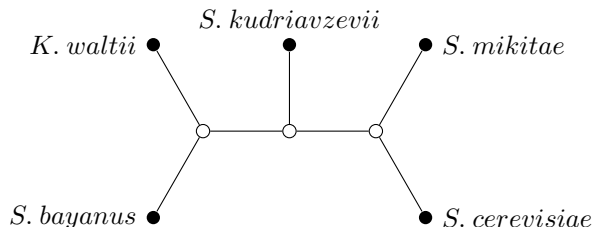
Figure 4: Quintet tree T_5 of yeast species (Marcet-Houben & Gabaldón, 2009).

The results of the confirmatory tetrad analysis for T_5-compatibility (see Fig. 4) gave p-values of 0.721 and 0.955 for the first and third components respectively. To double check these results we repeated the test using the bootstrapping strategy outlined in Bollen & Stine (1992). The results were very similar with p-values of 0.729 and 0.921 respectively. The confirmatory tetrad analysis and inverse-Wishart simulation results both gave upper bounds on T_5-compatibility, but on balance we concluded that the first and third components were T_5-compatible. Therefore, the class of Gaussian latent tree models did appear suitable for modelling some aspects of these yeast species' growth curves. However, for features relating to components 2 and 4, there is some evidence to support the exploration of a wider model class that could accommodate the hybrid hypothesis described in Libkind et al. (2011).

6 Discussion

Understanding the complete description of the correlation space associated with Gaussian latent tree models opens up a number of useful tools for assessing tree-compatibility either on a class basis or for a specified tree. Some of the methods described in this paper are particularly useful as part of an exploratory analysis for defining the relevant model search space, whereas others are ideal as a final step to check the conclusions of a model search. The complete semialgebraic structure of the correlation space has not been utilized elsewhere for assessing tree-compatibility of data, though the positivity constraint has been used previously, see Shiers et al. (2014). Incorporating a prior such as the inverse-Wishart and sampling from the posterior distribution allows for probabilistic conclusions about the model. This provides a more nuanced answer than a simple assessment of inequalities via the plugging in of covariance point estimates, and enables two or more incompatible but plausible trees to be compared.

One important practical consideration is the scalability of these methods. Techniques employing the semi-algebraic constraints can be adapted to larger number of variables reasonably well. For a confirmatory tetrad analysis the biggest computational cost is the calculation of the covariance of minors, which for p observed random variables has dimension of order p^4 which can become prohibitive. For example, if 8GB of RAM is allocated for a single matrix, the limit of p is approximately 25 even if redundant rows and columns are removed from the matrix. However, much larger p can be considered by calculating the relevant statistics for each quartet marginally. Then the covariance matrix of minors has dimension of only 36 and the memory can be released once each quartet has been tested. In either case, the final memory requirement could further be reduced with smart programming taking advantage of symmetries and sparseness. In a similar vein, to extend the scope of exploratory tetrad analysis to a greater number of variables, one strategy is to only assess single quartets in the first instance and use these results to reduce the set of possible trees worth considering. Given the effectiveness of the quartet testing as demonstrated with even small sample sizes, this approach seems sensible and to have significant advantage over methods that require a whole model to be tested at once.

Acknowledgement

A   Proofs

of Theorem 2. Assume first that all correlations ρ_{ij} are strictly positive, that is $R \in \mathrm{PO}_+(V)$ or $R \in \mathrm{PO}_+(T)$. We use the fact that $\mathrm{PO}_+(V)$ is isomorphic to the space of tree metrics, whose constraints are given in Theorem 1 and Corollary 1. Translating these constraints via $d_{ij} = -\log(\rho_{ij})$ gives exactly the constraints in the proposed theorem. These constraints describe a closed set, which is the smallest closed set containing $\mathrm{PO}_+(V)$. So it is enough to show that the closure of $\mathrm{PO}_+(T)$ is equal to $\mathrm{PO}(T)$. This follows from the fact that $\mathrm{PO}(T)$ is a toric cube and, by Engström et al. (2012, Theorem 1), every toric cube is equal to the closure of its interior. $\square$

of Theorem 3. If $R \in M(T)$ then each ρ_{ij} has representation (1). Thus $|\rho_{ij}| = \prod_{e \in \mathrm{ph}(ij)} |\rho_e|$ and hence R' also lies in $\mathrm{PO}(T)$. To show that $\rho_{ij}\rho_{ik}\rho_{jk} \geq 0$ consider the induced subtree over i, j, k, that is, the smallest connected subgraph of T containing vertices i, j, k. This subtree necessarily has a unique vertex v that lies on the intersection of paths $\mathrm{ph}(ij)$, $\mathrm{ph}(ik)$ and $\mathrm{ph}(jk)$. Moreover, by (1),

$$\rho_{ij}\rho_{ik}\rho_{jk} = \prod_{e \in \mathrm{ph}(ij)} \rho_e \prod_{e \in \mathrm{ph}(ik)} \rho_e \prod_{e \in \mathrm{ph}(jk)} \rho_e = \prod_{e \in \mathrm{ph}(iv)} \rho_e^2 \prod_{e \in \mathrm{ph}(jv)} \rho_e^2 \prod_{e \in \mathrm{ph}(kv)} \rho_e^2 \geq 0.$$

To prove the reverse implication, we note that every correlation matrix in $\mathrm{PO}(T)$ has (after permuting rows and columns) a block diagonal structure with strictly positive elements in each block. Consider first the case when all elements of R are non-zero, that is, R' has strictly positive entries. Distinguish one node in V and label it as 1. Let D be a diagonal matrix such that $D_{ii} = -1$ if $\rho_{1i} < 0$ and $D_{ii} = 1$ if $\rho_{1i} > 0$. If $R \in M(T)$ then also DRD lies in $M(T)$ because $M(T)$ is invariant with respect to all diagonal transformations. Moreover, it holds that $R' = DRD$ because $D_{11}D_{ii}\rho_{1i} = |\rho_{1i}|$ for all $i \in V \setminus \{1\}$ and $D_{ii}D_{jj}\rho_{ij} = |\rho_{ij}|$ for $i, j \in V \setminus \{1\}$. This last equality follows from our assumption that $\rho_{1i}\rho_{1j}\rho_{ij} \geq 0$ so that the sign of $\rho_{1i}\rho_{1j}$ is equal to the sign of ρ_{ij}. Now, since $R' \in \mathrm{PO}(T) \subset M(T)$ and $R = DR'D$ we also have that $R \in M(T)$. The analysis of the case when R is block diagonal will be omitted. $\square$

References

ALLMAN, E. S., RHODES, J. A., STURMFELS, B. & ZWIERNIK, P. (2015). Tensors of nonnegative rank two. *Linear Algebra and its Applications* **473**, 37 – 53. Special issue on Statistics.

ASTON, J. A. D., CHIOU, J.-M. & EVANS, J. P. (2010). Linguistic pitch analysis using functional principal component mixed effect models. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **59**, 297–317.

BARNARD, J., MCCULLOCH, R. & MENG, X.-L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* **10**, 1281–1312.

BEKKER, P. A. & DE LEEUW, J. (1987). The rank of reduced dispersion matrices. *Psychometrika* **52**, 125–135.

BOLLEN, K. A. & STINE, R. A. (1992). Bootstrapping goodness-of-fit measures in structural equation models. *Sociological Methods & Research* **21**, 205–229.

BOLLEN, K. A. & TING, K.-F. (1993). Confirmatory tetrad analysis. In *Sociological Methodology*, P. Marsden, ed., vol. 23 of *Sociological Methodology*, chap. 5. Wiley, pp. 147–175.

BOUCHARD-CÔTÉ, A., HALL, D., GRIFFITHS, T. L. & KLEIN, D. (2013). Automated reconstruction of ancient languages using probabilistic models of sound change. *Proceedings of the National Academy of Sciences* **110**, 4224–4229.

BUNEMAN, P. (1974). A note on the metric properties of trees. *J. Combinatorial Theory Ser. B* **17**, 48–50.

CARLIN, B. & LOUIS, T. (2008). *Bayesian Methods for Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

CHOI, M. J., TAN, V. Y. F., ANANDKUMAR, A. & WILLSKY, A. S. (2011). Learning latent tree graphical models. *J. Mach. Learn. Res.* **12**, 1771–1812.

DRESS, A., HUBER, K. T., KOOLEN, J., MOULTON, V. & SPILLNER, A. (2012). *Basic phylogenetic combinatorics*. Cambridge: Cambridge University Press.

DRTON, M., LIN, S., WEIHS, L. & ZWIERNIK, P. (2016a). Marginal likelihood and model selection for Gaussian latent tree and forest models. *to appear in Bernoulli, also available on arXiv:1412.8285* .

DRTON, M., MASSAM, H. & OLKIN, I. (2008). Moments of minors of Wishart matrices. *Ann. Statist.* **36**, 2261–2283.

DRTON, M., STURMFELS, B. & SULLIVANT, S. (2007). Algebraic factor analysis: tetrads, pentads and beyond. *Probab. Theory Related Fields* **138**, 463–493.

DRTON, M., XIAO, H. et al. (2016b). Wald tests of singular hypotheses. *Bernoulli* **22**, 38–59.

DUNN, O. J. (1961). Multiple comparisons among means. *Journal of the American Statistical Association* **56**, 52–64.

EISENSTEIN, J., O'CONNOR, B., SMITH, N. A. & XING, E. P. (2010). A latent variable model for geographic lexical variation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

ENGSTRÖM, A., HERSH, P. & STURMFELS, B. (2012). Toric cubes. *Rendiconti del Circolo Matematico di Palermo* , 1–12.

FOX, J., NIE, Z. & BYRNES, J. (2014). *sem: Structural Equation Models*. R package version 3.1-5.

FRIEDMAN, N., NINIO, M., PE'ER, I. & PUPKO, T. (2002). A structural em algorithm for phylogenetic inference. *Journal of Computational Biology* **9**, 331–353.

GARDNER, J. R., KUSNER, M. J., XU, Z. E., WEINBERGER, K. Q. & CUNNINGHAM, J. (2014). Bayesian optimization with inequality constraints. In *ICML*.

GELMAN, A., CARLIN, J., STERN, H., DUNSON, D., VEHTARI, A. & RUBIN, D. (2013). *Bayesian Data Analysis, Third Edition*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.

GILL, J., LINUSSON, S., MOULTON, V. & STEEL, M. (2008). A regular decomposition of the edge-product space of phylogenetic trees. *Adv. in Appl. Math.* **41**, 158–176.

GRÜNEWALD, S., HUMPHRIES, P. J. & SEMPLE, C. (2008). Quartet compatibility and the quartet graph. *The Electronic Journal of Combinatorics* **15**, R103.

GU, X., MULDER, J., DEKOVIĆ, M. & HOIJTINK, H. (2014). Bayesian evaluation of inequality constrained hypotheses. *Psychological methods* **19**, 511.

HAO, W. & GOLDING, G. B. (2008). Uncovering rate variation of lateral gene transfer during bacterial genome evolution. *BMC genomics* **9**, 235.

JUNKER, B. & SCHREIBER, F. (2011). *Analysis of Biological Networks*. Wiley Series in Bioinformatics. Wiley.

KIM, J. (2000). Slicing hyperdimensional oranges: the geometry of phylogenetic estimation. *Molecular phylogenetics and evolution* **17**, 58–75.

KOLLER, D. & FRIEDMAN, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press.

LAURITZEN, S. L. (1996). *Graphical models*, vol. 17 of *Oxford Statistical Science Series*. Oxford University Press. Oxford Science Publications.

LAWRENCE, N. D. (2004). Gaussian process latent variable models for visualisation of high dimensional data. *Advances in neural information processing systems* **16**, 329–336.

LIBKIND, D., HITTINGER, C. T., VALÉRIO, E., GONÇALVES, C., DOVER, J., JOHNSTON, M., GONÇALVES, P. & SAMPAIO, J. P. (2011). Microbe domestication and the identification of the wild genetic stock of lager-brewing yeast. *Proceedings of the National Academy of Sciences* **108**, 14539–14544.

LITI, G., CARTER, D. M., MOSES, A. M., WARRINGER, J., PARTS, L., JAMES, S. A., DAVEY, R. P., ROBERTS, I. N., BURT, A., KOUFOPANOU, V. et al. (2009). Population genomics of domestic and wild yeasts. *Nature* **458**, 337–341.

MARCET-HOUBEN, M. & GABALDÓN, T. (2009). The tree versus the forest: the fungal tree of life and the topological diversity within the yeast phylome. *PLoS One* **4**, e4357.

MOULTON, V. & STEEL, M. (2004). Peeling phylogenetic 'oranges'. *Advances in Applied Mathematics* **33**, 710–727.

MOURAD, R., SINOQUET, C., ZHANG, N. L., LIU, T. & LERAY, P. (2013). A survey on latent tree models and applications. *J. Artif. Intell. Res. (JAIR)* **47**, 157–203.

O'MALLEY, A. J. & ZASLAVSKY, A. M. (2008). Domain-level covariance analysis for multilevel survey data with structured nonresponse. *Journal of the American Statistical Association* **103**, 1405–1418.

PEARL, J. & XU, L. (1987). Structuring causal tree models with continuous variables. In *Proceedings of the Third Annual Conference on Uncertainty in Artificial Intelligence, Seattle, WA, USA, July 10-12, 1987*.

ROVERATO, A. (2002). Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models. *Scandinavian Journal of Statistics* **29**, 391–411.

ROYSTON, J. (1983). Some techniques for assessing multivariate normality based on the Shapiro-Wilk W. *Applied Statistics* , 121–133.

SEMPLE, C. & STEEL, M. (2003). *Phylogenetics*, vol. 24 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford: Oxford University Press.

SHIERS, N., ASTON, J. A. D., SMITH, J. Q. & COLEMAN, J. S. (2014). Gaussian Tree Constraints Applied to Acoustic Linguistic Functional Data. *ArXiv e-prints* .

SULLIVANT, S. (2008). Algebraic geometry of Gaussian Bayesian networks. *Advances in Applied Mathematics* **40**, 482–513.

SUNG, W. (2009). *Algorithms in Bioinformatics: A Practical Introduction*. Chapman & Hall/CRC Mathematical and Computational Biology. CRC Press.

VAN DE SCHOOT, R., HOIJTINK, H., HALLQUIST, M. N. & BOELEN, P. A. (2012). Bayesian evaluation of inequality-constrained hypotheses in sem models using m plus. *Structural equation modeling: a multidisciplinary journal* **19**, 593–609.

WANG, Y., ZHANG, N. L. & CHEN, T. (2008). Latent tree models and approximate inference in Bayesian networks. *J. Artif. Intell. Res.(JAIR)* **32**, 879–900.

WARRINGER, J., ZORGO, E., CUBILLOS, F. A., ZIA, A., GJUVSLAND, A., SIMPSON, J. T., FORSMARK, A., DURBIN, R., OMHOLT, S. W., LOUIS, E. J. et al. (2011). Trait variation in yeast is defined by population history. *PLoS Genet* **7**, e1002111.

WISHART, J. (1928). The generalised product moment distribution in samples from a normal multivariate population. *Biometrika* , 32–52.

ZWIERNIK, P. (2016). *Semialgebraic statistics and latent tree models*, vol. 146 of *Monographs on Statistics and Applied Probability*. Chapman & Hall/CRC.

ZWIERNIK, P. & SMITH, J. Q. (2011). Implicit inequality constraints in a binary tree model. *Electron. J. Statist.* **5**, 1276–1312.
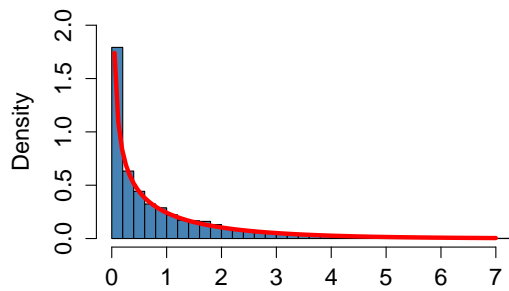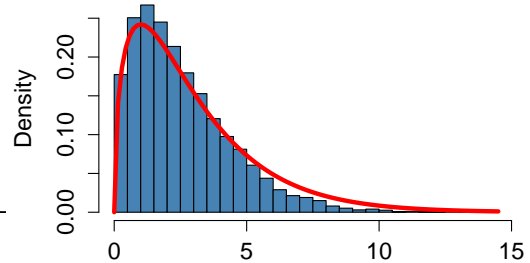
(a) single tetrad
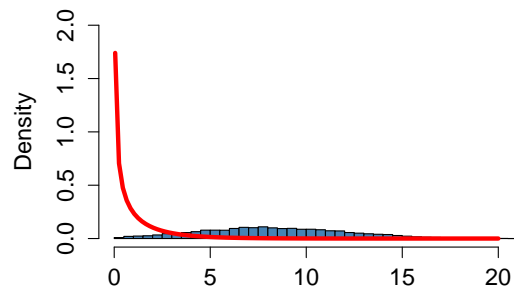
(b) single false tetrad
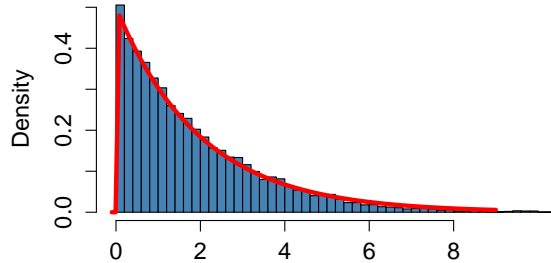
(c) two tetrads

(d) defining tetrads

Figure 5: Illustration of the simulations in Section 5.2. Sixty observations are generated from a random matrix in the tree model for the quintet tree $12|5|34$ and the corresponding test statistic is computed. This procedure is iterated 10,000 times. In each figure we compare the sample distribution of a test statistics against its theoretical distribution. In (a) we test a single tetrad $12|34$. In (b) we test a single false tetrad $13|24$. In (c) we test two tetrads $12|35$, $15|34$, and in (d) we test a minimal set of quartets defining the true quintet tree. The solid lines are densities of χ_1^2, χ_1^2, χ_2^2, and χ_3^2 respectively.