

A COVARIANCE MATRIX INVERSION PROBLEM ARISING FROM THE CONSTRUCTION OF PHYLOGENETIC TREES

TOM M. W. NYE, BRAD J. C. BAXTER AND WALTER R. GILKS

Abstract

We describe an efficient algorithm for the inversion of covariance matrices that arise in the context of phylogenetic tree construction. Phylogenetic trees describe the evolutionary relationships between species, and their construction is computationally demanding. Many approaches involve the symmetric matrix of evolutionary distances between species. Regarding these distances as random variables, the corresponding set of variances and covariances form a rank-4 tensor, and the inner-product defined by its inverse can be used to assign statistical scores to candidate trees. We describe a natural set of assumptions for the phylogenetic tree under construction, and show how under these assumptions the covariance tensor for a tree with n leaves can be inverted in $O(n^2)$ operations. In addition to presenting the inversion algorithm, we hope this article will open algebraic and computational problems from the field of phylogeny to a wider audience.

1. Introduction

Suppose that we are given a set of n species and an $n \times n$ symmetric matrix of random variables (d_{ij}) representing the evolutionary distances between them. In this paper we show how simple assumptions on the tree of evolutionary relationships between species give rise to a covariance matrix which essentially has the form

$$\text{Cov}(d_{ij}, d_{kl}) = \frac{1}{2}B(\delta_{ik}\delta_{jl} + \delta_{jk}\delta_{il}) + C^2 + \frac{1}{2}Ca_j(\delta_{jk} + \delta_{jl}) + \frac{1}{2}Ca_i(\delta_{ik} + \delta_{il}), \quad (1)$$

for $i \neq j$ and $k \neq l$, where B and C are positive constants, $a \in \mathbb{R}^n$ has every component non-negative, and δ_{ij} is the Kronecker delta. One advantage of our model is that covariance matrices of this form can be inverted in $O(n^2)$ operations, via the Sherman–Morrison–Woodbury formula (see, for instance, [8, p. 51]).

In order to motivate this inversion problem it is necessary to present some background from the field of phylogeny. While this area might be unfamiliar to the reader, we hope that the mathematical problems it raises will be of interest. The rest of this section gives a very brief introduction to phylogeny, before we describe the origin of equation (1) in Section 2. We then show how the covariance tensor can be inverted algebraically in Section 3, and present the inversion algorithm in Section 4. Readers principally concerned by the linear algebraic details of the inversion algorithm rather than the phylogenetic background may skip directly to Section 3.

Received 22 August 2005, revised 2 February 2007; published 23 March 2007.

2000 Mathematics Subject Classification 15A09, 15A51, 92B15

© 2007, Tom M. W. Nye, Brad J. C. Baxter and Walter R. Gilks

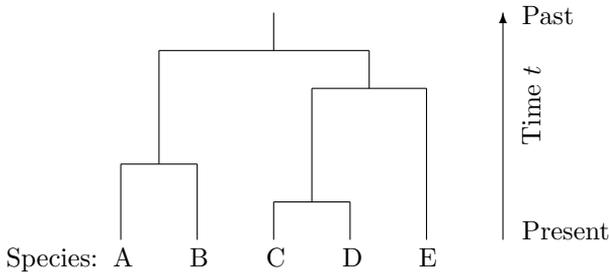


Figure 1: A typical phylogenetic tree

Evolutionary relationships between species can be represented by a tree: the leaf nodes represent extant species, interior nodes represent ancestral species, and the branch lengths indicate the extent to which species have diverged. Such trees are referred to as *phylogenies*. One kind of tree is illustrated in Figure 1. Here the branch lengths specify the time since divergence of species, and it is this type of tree we will consider throughout this paper.

There are a range of different statistical methods available for inferring the phylogeny of a set of species given their DNA sequences. The DNA sequence for each species may be a single gene, a set of genes, or even an entire genome. All of these methods are based on the fundamental idea that species with similar DNA sequences are more closely related than species for which the sequences have diverged, as mutations in sequence accumulate with time. One class of methods, the so-called *distance-based* approaches, constructs a matrix of evolutionary distances, or distance functionals, between species that, to a certain degree, summarizes the information contained in the full set of sequences [3, 11, 9, 6, 2]. The evolutionary distance between two species typically measures the number of letter changes in the DNA sequence, and several different distance functionals exist. Distance-based methods take the matrix of distances between extant species and from that infer the topology and branch lengths of the underlying phylogenetic tree. They have the advantage of being relatively fast, and therefore suitable for large problems, but are less suitable when the set of sequences under investigation has diverged widely.

While the biological literature is extensive, there is an increasing mathematical and statistical literature relating to phylogenetic reconstruction. The book by Gascuel *et al.* [7] is a comprehensive reference to the mathematics of phylogeny; the first chapter concerns distance-based methods. A more practical approach is adopted in the book by Felsenstein [5], which describes the principal statistical and computational approaches to phylogenetic reconstruction. Current developments and future challenges in phylogeny have been described in a recent review [4].

2. Origin of the covariance tensor

The inversion problem studied in this paper arises from a novel distance-based phylogenetic method which we briefly describe in this section. The method arose as a prototype for a more complete statistical treatment of variances and covariances in

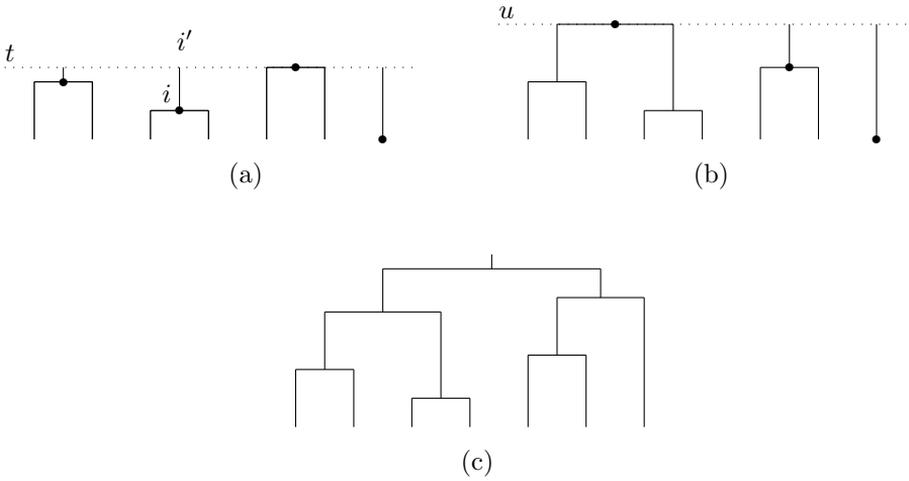


Figure 3: Tree construction. (a) A generic point in construction: the tree is complete up to time t . The set of hanging nodes H_t is highlighted with circles. (b) The next step in the construction obtained from (a). Construction is complete up to a time $u > t$. The set of hanging nodes H_u is highlighted. (c) The finished tree.

process on each branch; however, such probabilistic details are not our main concern here.

Equation (3) defines the covariance structure when the underlying phylogenetic tree T is known. However, our estimate of T is built up as a sequence of partially constructed trees, as illustrated by Figure 3. A generic stage of the construction is shown in Figure 3(a). This consists of an estimate of T constructed back as far as some time t , which we denote T_t . The covariance tensor defined in equation (1) arises from considering the set of nodes descended directly from time t with no bifurcation. We refer to these as ‘hanging nodes’, as suggested by the appearance in Figure 3, and the set of such nodes is denoted H_t . These nodes are highlighted by black circles in Figure 3(a). Given a node $i \in H_t$ we use the notation i' to denote the ancestor of i at time t , and let t_i denote the time of node i . The length of the line segment between i and i' in Figure 3(a) is therefore $t - t_i$.

Given $i, j, k, l \in H_t$ such that $i \neq j$ and $k \neq l$, the additivity condition (4) gives

$$\text{Cov} [d_{ij}, d_{kl} \mid T_t] = \text{Cov} [d_{ii'} + d_{i'j'} + d_{j'j}, d_{kk'} + d_{k'l'} + d_{l'l} \mid T_t]. \quad (5)$$

The distances $d_{ii'}$ and $d_{kk'}$ have a covariance of zero (when $i \neq k$) because the corresponding tree branches have no overlap (that is, the right-hand side of equation (3) is zero). This applies to the other indices as well, so expanding the right-hand side of equation (5) gives

$$\begin{aligned} \text{Cov} [d_{ij}, d_{kl} \mid T_t] &= (\delta_{ik} + \delta_{il})\text{Var}[d_{ii'} \mid T_t] \\ &\quad + (\delta_{jk} + \delta_{jl})\text{Var}[d_{jj'} \mid T_t] + \text{Cov}[d_{i'j'}, d_{k'l'} \mid T_t]. \end{aligned} \quad (6)$$

The first two terms of this expression can be obtained using equation (3):

$$\text{Var}[d_{ii'} \mid T_t] = \nu(t - t_i), \quad \text{and} \quad \text{Var}[d_{jj'} \mid T_t] = \nu(t - t_j).$$

The final term of equation (6) depends on the paths joining i' to j' and k' to l' . However, these paths lie in the part of the tree that has not yet been estimated (that is, above the dotted line in Figure 3(a)). By symmetry, since we condition only on T_t , the final term of equation (6) depends only on whether the two paths $t_{i'j'}$ and $t_{k'l'}$ share common terminal nodes: it adopts three different values according to whether the paths share zero, one, or two terminal nodes. This gives

$$\begin{aligned} \text{Cov}[d_{ij}, d_{kl} \mid T_t] &= \nu(t - t_i)(\delta_{ik} + \delta_{il}) + \nu(t - t_j)(\delta_{jk} + \delta_{jl}) \\ &+ c_0 + c_1(\delta_{ik} + \delta_{il} + \delta_{jk} + \delta_{jl}) + c_2(\delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk}) \end{aligned} \quad (7)$$

for some positive constants c_0, c_1, c_2 . This equation is now in exactly the same form as equation (1), and it is this tensor that is used to score different partially constructed trees.

Construction of the estimated phylogeny is carried out in the following way. Note that our intention here is not to provide detailed statistical justification for the procedure, but simply to provide motivation for the inversion problem. Given a partially constructed tree T_t , at each stage we propose a tree T_u ($u > t$) by joining together two nodes from H_t with a new node at time u . One such tree is shown in Figure 3(b). This tree T_u is assigned a χ^2 statistic using the inverse of the covariance tensor (7). First, the ordinary least squares estimates d_u of the distances between the hanging nodes H_u are obtained given d_t (the matrix of distances between nodes in H_t). Then, the χ^2 statistic assigned to T_u is defined by

$$\chi^2(T_t, T_u) = \langle (d_u - \mathbb{E}[d_u \mid T_u]), \text{Cov}[d_u \mid T_u]^{-1} (d_u - \mathbb{E}[d_u \mid T_u]) \rangle$$

where the inner product is the Frobenius inner product on matrices. The time u is chosen to optimize this score. Every possible tree T_u derived from T_t by joining two nodes together is scored and assigned a time in this way, so that if H_t contains m nodes then $m(m - 1)/2$ candidate trees are considered. The candidate T_u for which $u - t$ is smallest is then taken as the new estimate of the underlying tree, and the construction procedure continues recursively from T_u . The process starts from the set of leaf nodes and ends when the entire tree has been estimated. Of course, in order to construct the tree it is necessary to invert the covariance tensor many times, and so an efficient inversion algorithm is highly desirable. Examples of some trees constructed using this algorithm are given in Section 5.

3. Inverting the covariance tensor

We start by defining some notation. Let $V = \mathbb{R}^n$ be equipped with the standard basis $\{e_i : i = 1, \dots, n\}$ and define the vector of all ones

$$e = \sum_{i=1}^n e_i.$$

Let W denote the vector space of symmetric $n \times n$ real matrices, so that

$$W = \text{span}\left\{\frac{1}{2}e_i \otimes e_j + \frac{1}{2}e_j \otimes e_i : i \leq j\right\}. \quad (8)$$

We shall move between the equivalent notations $u \otimes v$ and uv^T , for $u, v \in \mathbb{R}^n$, as appropriate. If we view $\mathbb{R}^{n \times n}$ as the tensor product space $\mathbb{R}^n \otimes \mathbb{R}^n$, then the inner product

$$\langle a \otimes b, c \otimes d \rangle = \langle a, c \rangle \langle b, d \rangle, \quad a, b, c, d \in \mathbb{R}^n,$$

implies that

$$\langle M_1, M_2 \rangle = \sum_{j,k=1}^n M_1(j, k)M_2(j, k).$$

In other words, our inner product on W is simply the Frobenius inner product on matrices.

Using the basis of W implicit in equation (8), the covariance tensor (1) defines a linear map $H : W \rightarrow W$ given by

$$H = B \times \text{id} + C^2\theta \otimes \theta + C\alpha \otimes \theta + C\theta \otimes \alpha$$

where $\theta = e \otimes e$ and $\alpha = \text{diag}(a_1, \dots, a_n)$. Thus H is a perturbation of the identity matrix, namely $H = B \times \text{id} + A$, where

$$A = C^2\theta \otimes \theta + C\alpha \otimes \theta + C\theta \otimes \alpha.$$

Since θ projects onto the direction e , $\theta(v) = 0$ for any vector v perpendicular to e . Similarly, it can be seen that

$$\begin{aligned} A(u \otimes v + v \otimes u) &= C^2(\theta u) \otimes (\theta v) + C^2(\theta v) \otimes (\theta u) \\ &\quad + C(\alpha u) \otimes (\theta v) + C(\alpha v) \otimes (\theta u) \\ &\quad + C(\theta u) \otimes (\alpha v) + C(\theta v) \otimes (\alpha u) \\ &= 0 \quad \text{for all } u, v \text{ perpendicular to } e. \end{aligned}$$

On the other hand, A is non-zero (in general) on vectors of the form $v \otimes e + e \otimes v$ for any $v \in \mathbb{R}^n$. This suggests that we define the subspace

$$U = \text{span}\{w \otimes e + e \otimes w : w \in \mathbb{R}^n\}.$$

It is easy to verify that

$$U^\perp = \text{span}\{u \otimes v + v \otimes u : u, v \in e^\perp\}.$$

We have already shown that $U^\perp \subset \ker A$ and, since A is symmetric, we deduce that $\text{im } A \subset U$, which is also evident by direct calculation. Thus A has the following decomposition.

$$A = \left[\begin{array}{c|c} U & U^\perp \\ \hline * & 0 \\ 0 & 0 \end{array} \right] \begin{array}{l} \updownarrow U \\ \updownarrow U^\perp \end{array}$$

The $*$ symbol represents the non-zero part of A . At this stage it is apparent that away from U , H is trivial, so the inversion problem reduces to the problem of inverting H on the n -dimensional subspace U .

However, up to this point we have ignored a crucial point: the random variables d_{ij} satisfy

$$d_{ii} \equiv 0 \quad \text{for } i = 1, \dots, n.$$

(In other words the distance of a species from itself is zero.) Instead of working on the space W , we are really dealing with the restriction of H to the space

$$\hat{W} = \text{span}\{\frac{1}{2}e_i \otimes e_j + \frac{1}{2}e_j \otimes e_i : i < j\}.$$

Note that the inequality in the indices is strict here. If we define $P : W \rightarrow W$ by

$$P\left(\frac{1}{2}e_i \otimes e_j + \frac{1}{2}e_j \otimes e_i\right) = \begin{cases} \frac{1}{2}e_i \otimes e_j + \frac{1}{2}e_j \otimes e_i & \text{if } i \neq j, \\ 0 & \text{when } i = j, \end{cases}$$

then $\hat{W} = P(W)$ and the restriction of H to \hat{W} , denoted \hat{H} , is given by

$$\begin{aligned} \hat{H} &= PHP \\ &= PAP + P(B \times \text{id})P. \end{aligned}$$

It is really the map \hat{H} that we need to invert; $P(B \times \text{id})P$ is simply a multiple of the identity on \hat{W} .

Essentially we want to identify a decomposition for \hat{H} equivalent to the one above. Let $\hat{U} = P(U)$ and let \hat{U}^\perp be the orthogonal complement of \hat{U} in \hat{W} , so that

$$\hat{W} = \hat{U} \oplus \hat{U}^\perp.$$

As we showed above, $\text{im } A \subset U$, so

$$\text{im } PAP \subset P(U) = \hat{U}.$$

Since A and P are symmetric, we therefore have the desired decomposition.

$$\begin{array}{ccc}
 & \hat{U} & \hat{U}^\perp \\
 & \longleftrightarrow & \longleftrightarrow \\
 PAP = & \left[\begin{array}{c|c} * & 0 \\ \hline 0 & 0 \end{array} \right] & \begin{array}{l} \updownarrow \hat{U} \\ \updownarrow \hat{U}^\perp \end{array}
 \end{array}$$

It follows that

$$\hat{H}^{-1}(x) = (P(A + B)P|_{\hat{U}})^{-1}(x_{\hat{U}}) + B^{-1}(x_{\hat{U}^\perp}) \tag{9}$$

where

$$x = x_{\hat{U}} \oplus x_{\hat{U}^\perp} \tag{10}$$

represents the \hat{U} and \hat{U}^\perp decomposition of $x \in \hat{W}$. The first term of (9) denotes the restriction of $P(A + B)P$ to the n -dimensional space \hat{U} . The inversion problem therefore reduces to the question of inverting this component of the map.

To address this problem it is useful to work with an orthonormal basis of \hat{U} . It can be shown that

$$w_k = \phi\left(\frac{1}{2}e_k \otimes e + \frac{1}{2}e \otimes e_k - e_k \otimes e_k\right) + \psi\left(e \otimes e - \sum_l e_l \otimes e_l\right)$$

for $k = 1, \dots, n$ defines an orthonormal basis of \hat{U} when

$$\phi = \left(\frac{2}{n-2}\right)^{1/2}$$

and

$$\psi = -\frac{1}{n} \left(\left(\frac{2}{n-2}\right)^{1/2} + \left(\frac{1}{n-1}\right)^{1/2} \right).$$

We assume that $n > 2$, since the inversion problem is trivial when $n = 2$, so that the constants are well-defined. In this basis \hat{H} is represented by the matrix

$$\begin{aligned} \Lambda_{ij} &= B\delta_{ij} + \langle w_i, PAP(w_j) \rangle \\ &= (C(n-2)a_i + B)\delta_{ij} + \frac{1}{2}C\phi(n-2)(\phi + 2(n-1)\psi)(a_i + a_j) \\ &\quad + \frac{1}{2}C(\phi + 2(n-1)\psi)^2 \left(\sum_l a_l \right) + C^2(n-1)^2(\phi + n\psi)^2. \end{aligned}$$

The matrix defines a linear map $\Lambda : \mathbb{R}^n \rightarrow \mathbb{R}^n$

$$\Lambda = M + \omega e^T + e\omega^T$$

where

$$M = \text{diag}(C(n-2)a_k + B : k = 1, \dots, n),$$

and $\omega \in \mathbb{R}^n$ is defined by

$$\begin{aligned} \omega &= \frac{1}{2}C\phi(n-2)(\phi + 2(n-1)\psi)a \\ &\quad + \left(\frac{1}{2}C^2(n-1)^2(\phi + n\psi)^2 + \frac{1}{4}C(\phi + 2(n-1)\psi)^2 \sum_l a_l \right) e. \end{aligned} \tag{11}$$

The matrix Λ is a rank-2 perturbation of M and its inverse is given by the Sherman–Morrison–Woodbury formula [8]:

$$\Lambda^{-1} = M^{-1} + \frac{c_{\omega\omega}\hat{e}\hat{e}^T + c_{ee}\hat{\omega}\hat{\omega}^T - (1 + c_{\omega e})(\hat{\omega}\hat{e}^T + \hat{e}\hat{\omega}^T)}{(1 + c_{\omega e})^2 - c_{\omega\omega}c_{ee}} \tag{12}$$

where

$$\hat{e} = M^{-1}e, \quad \hat{\omega} = M^{-1}\omega$$

and

$$c_{\omega\omega} = \langle M^{-1}\omega, \omega \rangle, \quad c_{ee} = \langle M^{-1}e, e \rangle, \quad \text{and} \quad c_{\omega e} = \langle M^{-1}\omega, e \rangle.$$

We now have all the elements in place for a complete algebraic inverse to the map \hat{H} defined by the covariance tensor (1). The next section puts these elements together and specifies the inversion algorithm.

4. The inversion algorithm

Suppose that we want to compute $\hat{H}^{-1}x$ for some symmetric matrix x_{ij} that is zero on the diagonal. The inversion algorithm has the following steps.

1. On account of the decomposition (9), the \hat{U} and \hat{U}^\perp components of x can be dealt with independently. The \hat{U} component of x is obtained by taking the inner product of x with the vectors w_k that form our basis of \hat{U} . We therefore define the vector $\xi \in \mathbb{R}^n$ by

$$\xi_k = \langle x, w_k \rangle = \frac{1}{2} \phi \left(\sum_i x_{ik} + \sum_j x_{kj} \right) + \psi \sum_{ij} x_{ij}. \tag{13}$$

2. The inverse of \hat{H} on \hat{U} is given by equation (12). Let η be the inverse of ξ :

$$\eta = \Lambda^{-1} \xi.$$

3. The final step is to combine this result with the \hat{U}^\perp component of the inverse. Using the decompositions (9) and (10), and since $B^{-1}(x_{\hat{U}^\perp}) = B^{-1}(x) - B^{-1}(x_{\hat{U}})$ it follows that

$$\hat{H}^{-1}(x) = (P(A + B)P|_{\hat{U}})^{-1}(x_{\hat{U}}) - B^{-1}(x_{\hat{U}}) + B^{-1}(x). \tag{14}$$

The first term of (14) is given using the result of step 2:

$$\begin{aligned} (P(A + B)P|_{\hat{U}})^{-1}(x_{\hat{U}}) &= \sum_{k=1}^n \eta_k w_k \\ &= \frac{\phi}{2} P(\eta \otimes e + e \otimes \eta) + \psi \langle e, \eta \rangle P(e \otimes e). \end{aligned}$$

By replacing η with $(\eta - B^{-1}\xi)$ in this equation we obtain the first two terms of (14). The full inverse is then given by

$$\begin{aligned} \hat{H}^{-1}(x) &= B^{-1}x + \frac{\phi}{2} P((\eta - B^{-1}\xi) \otimes e + e \otimes (\eta - B^{-1}\xi)) \\ &\quad + \psi \langle e, \eta - B^{-1}\xi \rangle P(e \otimes e). \end{aligned}$$

The computational complexity of the algorithm can be obtained by analysing each step. It is easy to see that overall there are $O(n^2)$ multiplications and $O(n^2)$ additions. Moreover, in our application we actually use the inverse covariance tensor to evaluate inner products like $\langle x, \hat{H}^{-1}x \rangle$ where x is a symmetric matrix with zero diagonal. If the partial sums in equation (13) are already known for the matrix x , then the inner product $\langle x, \hat{H}^{-1}x \rangle$ can be evaluated in $O(n)$ operations.

5. Examples

Although the main aim of this paper is to present the details of the inversion algorithm, some examples of phylogenies reconstructed with our approach may provide some insight. This section presents the results of applying the tree reconstruction algorithm described in Section 2 to some small simulated distance matrices. Simulated data offer an advantage over real biological data as the ‘true’ evolutionary history is known, and features of the simulated data can be varied systematically. The distance matrices were generated by taking a fixed tree T of divergence times, then (1) distorting each branch length, and (2) adding noise to the resulting distance matrix. Distortion was achieved by replacing each branch length t with a sample from a Gamma distribution with mean equal to the original branch length t , and

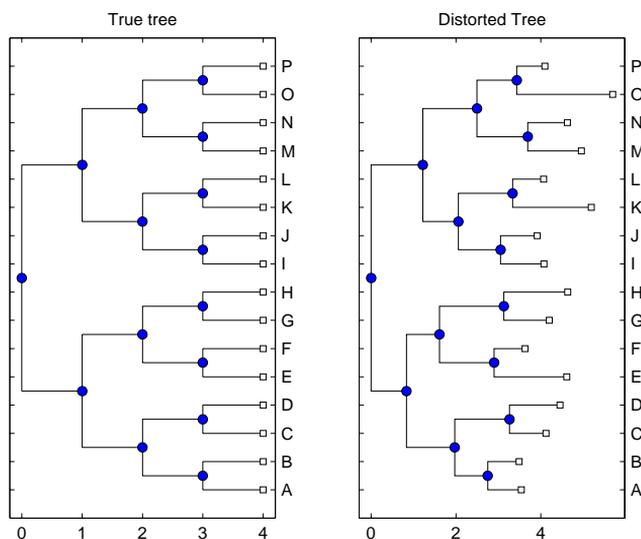


Figure 4: The fixed underlying tree T (left) and a distorted version with $p = 10\%$ (right).

standard deviation a fixed proportion of t (which we will specify as a percentage $p\%$). This distortion process replicates the assumptions of equations (2) and equations (3). Noise was added to the distance matrices in a similar way, by replacing each distance d with an independent sample from a Gamma distribution with mean d and standard deviation a fixed proportion of d specified by a percentage $q\%$. Noise of this form is not explicitly included in our model but will be present in any real biological data.

Figure 4 shows a typical tree T on 16 taxa and a distorted version of the same tree with $p = 10\%$. Figure 5 shows our reconstruction of T under three conditions: (1) $p = 10\%$ with no noise, (2) $q = 10\%$ with no distortion, and (3) $p = 10\%$, $q = 10\%$. As might be anticipated, since our approach explicitly models distortion, the correct topology is reconstructed when there is no noise. Even with $q = 10\%$ noise the reconstructed phylogeny is relatively accurate: a single topological error is present in the tree with $q = 10\%$ and no distortion. The success of the method depends on the topology of the underlying tree T as well as the amount of distortion and noise. Figure 6 shows some results for a different topology. Reconstruction for the topology shown in Figure 6 is generally more prone to error than that for the topology in Figure 4. Real biological phylogenies will generally lie between the two extremes.

6. Discussion

The inversion algorithm presented in Section 4 relies principally on the covariance matrix being a perturbation of the identity, which in turn is a result of the symmetry assumptions made to obtain equation (7). An alternative approach to tree construction is to propose and score a complete tree at each stage rather than

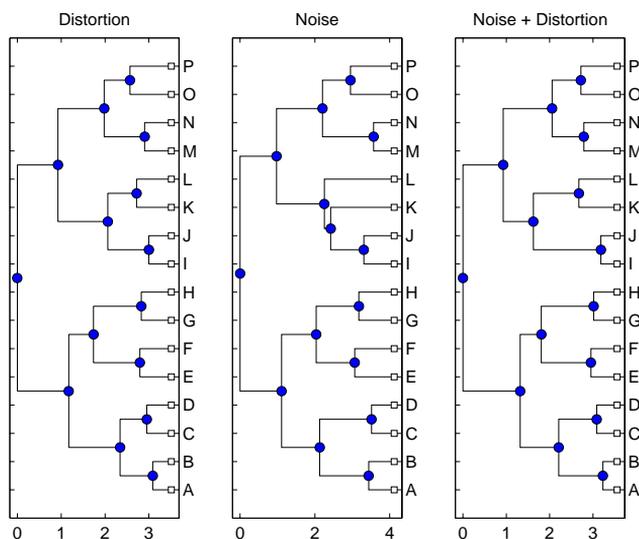


Figure 5: Results for phylogenetic reconstruction for the tree in Figure 4. Reconstruction from a distance matrix generated with $p = 10\%$ distortion and no noise (left); reconstruction for $q = 10\%$ noise and no distortion (centre); and reconstruction for $p = 10\%$, $q = 10\%$ (right).

build up the tree piece by piece as described at the end of Section 2. The space of such trees can be searched for the tree with the best score via a number of algorithms described in the phylogenetic literature. Given a proposed tree T that is complete, the tree would be assigned a χ^2 statistic using the inverse of a covariance matrix like that in equation (3). This approach raises a generalization of the inversion problem addressed in this paper: given any tree T that is complete and the set of paths t_{ij} between its leaf nodes, how can the covariance matrix defined by equation (3) be inverted efficiently?

A crucial advantage of the probabilistic model chosen in this paper is that it generates covariance matrices which are mild perturbations of the identity, so ensuring inexpensive inversion. However, the basis chosen is highly reminiscent of of similar bases chosen to elucidate the structure of Euclidean distance matrices, and we mention this interesting connection. We recall that an $n \times n$ matrix A is a Euclidean distance matrix if there exist vectors $u_1, \dots, u_n \in \mathbb{R}^n$ for which

$$A_{ij} = \|u_i - u_j\|^2, \quad 1 \leq i, j \leq n,$$

where $\|\cdot\|$ denotes the Euclidean norm. Such matrices were characterized by I. J. Schoenberg [10], who proved that a symmetric matrix M , whose diagonal elements vanish, is a Euclidean distance matrix if and only if $v^T M v \leq 0$ when v is orthogonal to e , the vector of all ones defined at the beginning of Section 3. The theory of such matrices is highly relevant to the linear of radial basis functions and learning theory; see, for example, [1].

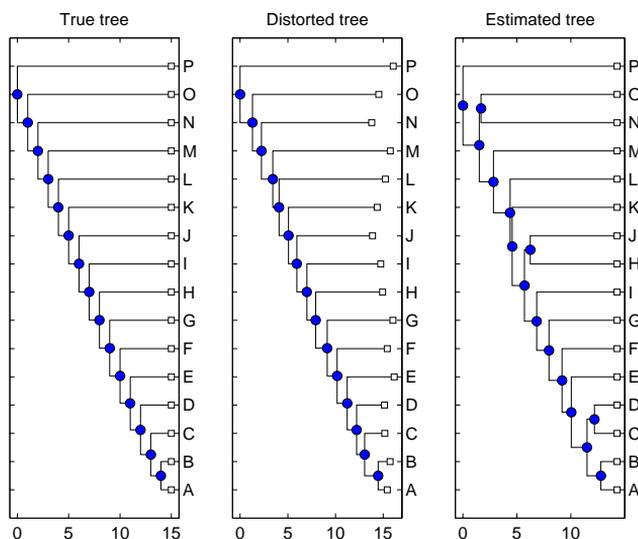


Figure 6: Results for phylogenetic reconstruction for a second underlying topology. The underlying tree T (left); a distorted version of T with $p = 5\%$ (centre); and reconstruction for $p = 5\%$, $q = 5\%$ (right).

References

1. B. J. C. BAXTER, ‘Conditionally positive functions and p-norm distance matrices’, *Constr. Approx.* 7 (1991) 427–440. [129](#)
2. W. J. BRUNO, N. D. SOCCI and A. L. HALPERN, ‘Weighted neighbour-joining: A likelihood-based approach to distance-based phylogeny reconstruction’, *Molecular Biology and Evolution* 17 (2000) 189–197. [120](#)
3. L. L. CAVALLI-SFORZA and A. W. F. EDWARDS, ‘Phylogenetic analysis: Models and estimation procedures’, *American Journal of Human Genetics* 19 (1967) 233–257. [120](#)
4. F. DELSUC, H. BRINKMANN and H. PHILIPPE, ‘Phylogenomics and the reconstruction of the tree of life’, *Nature Reviews Genetics* 6 (2005) 361–375. [120](#)
5. J. FELSENSTEIN, *Inferring phylogenies* (Sinauer, 2004). [120](#)
6. O. GASCUEL, ‘Bionj: An improved version of the neighbour-joining algorithm based on a simple model of sequence data’, *Molecular Biology and Evolution* 14 (1997) 685–695. [120](#)
7. O. GASCUEL (ed.), *Mathematics of evolution and phylogeny* (Oxford University Press, 2005). [120](#)
8. G. H. GOLUB and C. F. VAN LOAN, *Matrix computations* (Johns Hopkins University Press, 1996), 3rd edn. [119](#), [126](#)

9. N. SATOU and M. NEI, 'The neighbour-joining method: A new method for reconstructing phylogenetic trees', *Molecular Biology and Evolution* 4 (1987) 406–425. 120
10. I. J. SCHOENBERG, 'Remarks to Maurice Fréchet's article "Sur la définition axiomatique d'une classe d'espace distanciés vectoriellement applicable sur l'espace d'Hilbert"', *Ann. of Math.* 36 (1935) 724–736. 129
11. P. H. A. SNEATH and R. R. SOKAL, *Numerical taxonomy* (W. K. Freeman, 1973). 120

Tom M. W. Nye tom.nye@ncl.ac.uk
<http://www.mas.ncl.ac.uk/~ntmwn/>

School of Mathematics and Statistics
Newcastle University
Newcastle Upon Tyne NE1 7RU
United Kingdom

Brad J. C. Baxter b.baxter@bbk.ac.uk

Birkbeck
University of London
Malet Street
London WC1E 7HX
United Kingdom

Walter R. Gilks wally@maths.leeds.ac.uk

Department of Statistics
University of Leeds
Leeds LS2 9JT
United Kingdom