

PRINCIPAL COMPONENTS ANALYSIS IN THE SPACE OF PHYLOGENETIC TREES

BY TOM M. W. NYE

Newcastle University

Phylogenetic analysis of DNA or other data commonly gives rise to a collection or sample of inferred evolutionary trees. Principal Components Analysis (PCA) cannot be applied directly to collections of trees since the space of evolutionary trees on a fixed set of taxa is not a vector space. This paper describes a novel geometrical approach to PCA in tree-space that constructs the first principal path in an analogous way to standard linear Euclidean PCA. Given a data set of phylogenetic trees, a geodesic principal path is sought that maximizes the variance of the data under a form of projection onto the path. Due to the high dimensionality of tree-space and the nonlinear nature of this problem, the computational complexity is potentially very high, so approximate optimization algorithms are used to search for the optimal path. Principal paths identified in this way reveal and quantify the main sources of variation in the original collection of trees in terms of both topology and branch lengths. The approach is illustrated by application to simulated sets of trees and to a set of gene trees from metazoan (animal) species.

Introduction. Inference of evolutionary or *phylogenetic* trees is a fundamental task in many areas of biology, and tree estimation has developed over several decades into a mature statistical field [13]. On a phylogenetic tree, leaves correspond to existing observed taxa, internal vertices correspond to ancestral taxa, and branch lengths represent the degree of evolutionary divergence between taxa. A phylogenetic tree representing the division and divergence of different species is called a *species tree*. However, individual regions of DNA can evolve according to trees that differ from the underlying species tree, and an inferred phylogenetic tree from a particular gene or DNA region is called a *gene tree*. Gene trees can differ from the species tree for several reasons: random variation in the process of DNA letter substitution; population effects by which the evolutionary course of an individual gene does not match that of the species as a whole [10]; and even relatively rare events whereby genetic material is exchanged between species in a nontree-like manner [11]. Phylogenetic analysis of a number of different genes in a fixed set of species therefore generally gives rise to a collection of alternative phylogenetic trees. Collections of alternative phylogenetic trees also

Received July 2010; revised July 2011.

MSC2010 subject classifications. Primary 92D15; secondary 62H25.

Key words and phrases. Phylogeny, principal component, geodesic.

arise from inferential methods that involve simulation: bootstrap replication and MCMC sampling from Bayesian posteriors are widely used in the construction of phylogenetic estimates. Given such a collection of alternative trees, whether gene trees or a simulated sample, identifying differences and quantifying variation is a difficult problem, since we might potentially have several hundred trees on thousands of species. Standard multivariate statistical methods such as clustering [8, 23, 28] and Multi-Dimensional Scaling (MDS) [8, 19] have been used to address this problem. Principal Components Analysis (PCA), in contrast, cannot be applied directly since the space of phylogenetic trees on a fixed set of species is not a Euclidean vector space. This paper describes a geometric approach to PCA for sets of alternative phylogenetic trees. The aim is to identify which tree features are most variable within a given set of trees and to quantify this variation—just as the first few components in regular PCA pick out the most variable features of a Euclidean data set. Although PCA has been used to analyze different phylogenetic data previously (such as distance matrix data), the method presented here is the first to work intrinsically within the space of phylogenetic trees. The approach relies to a large extent on existing mathematical tools, and the main novel contribution comes from combining those elements into a computationally feasible method.

A key feature of our approach is the incorporation of both topological and geometrical information from the trees under analysis, via the so-called *geodesic metric* on the space of trees [5, 22, 25]. *Topological* information refers to the exact pattern of branching within a tree, while *geometrical* information refers to the distances between taxa induced by branch lengths on the tree. Topological features are generally more straightforward to characterize in a set of alternative trees, by counting the proportion of trees containing a given feature. For example, bootstrap replicate data sets are often represented by a single “consensus” tree annotated with a percentage support for each clade within the tree [12]. However, the geometry and topology of evolutionary trees are intimately related: we can continuously change the topology of a tree by shrinking down the length of any internal branch and expanding out an alternative branch in its place, as shown in Figure 1. Recent authors [22] have stressed the importance of using geometrical information to draw comparisons between trees on account of the interdependence of tree geometry and topology, and due to the increased distinguishability obtained by using continuous rather than discrete metrics. Moreover, tree geometry plays an important role in inference: it has been shown that long branches tend to “attract” each other, leading to mistakes in the topology of inferred trees [21].

Taking a set of alternative phylogenetic trees on some fixed set of taxa as input, our approach identifies a path L through tree-space that can be thought of analogously to the first principal component in regular Euclidean PCA. The path consists of a smoothly changing tree structure in which certain branches expand or shrink. Alternative topologies emerge when internal branches are shrunk to have

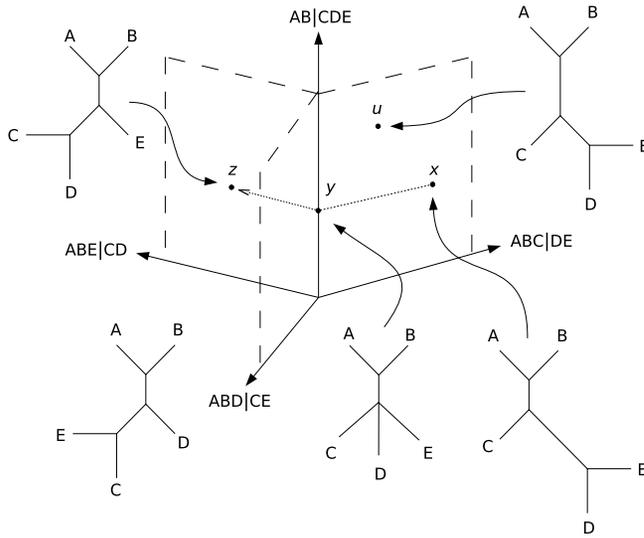


FIG. 1. A schematic view of a region of tree-space on five taxa: points in space correspond to unrooted trees. Trees with the same topology all lie in the same quadrant of tree-space (trees x , u , e.g.). Different quadrants are joined along their edges. Tree x can be continuously deformed into tree z by shrinking down the branch DE via tree y and replacing it with the branch CD . It follows that tree z is obtained from x via nearest neighbor interchange (NNI) of the split $ABC|DE$ into split $ABE|CD$. However, at y another NNI move is possible: $ABC|DE$ could be replaced by $ABD|CE$, corresponding to the lower left quadrant.

zero length and are then replaced with topologically distinct branches. The path L is constructed in such a way that the changing features—both in terms of topology and geometry—correspond to the most variable features within the data set. Just as for regular PCA, L also captures correlations in the data set: features that tend to occur together in the data are also represented together on L . A quantitative measure of variability can be assigned to L , in analogy with the proportion of variance contributed by each component in regular PCA. Unlike Euclidean vector spaces, there is no inner product on tree-space, and so the analysis cannot be extended in a straightforward manner to provide higher order principal paths by working orthogonally to L . Further discussion is given in Section 6.

Our approach—which we will refer to as Φ PCA (for “phylogenetic” PCA)—is motivated by geometrical analogy with regular vector space PCA. Construction of the first principal component in a Euclidean vector space can be thought of as follows:

- (1) Given a set of vectors x_1, \dots, x_n identify the centroid \bar{x} .
- (2) For a fixed line L through \bar{x} take the orthogonal projection of the points x_1, \dots, x_n onto L .

(3) Identify the line that maximizes the variance of the projected points along L , or, equivalently, which minimizes the sum of squared orthogonal distances of the points from the line.

For a Euclidean vector space, these steps can be re-expressed and solved in terms of simple linear algebra. However, the space of phylogenetic trees on a fixed set of taxa is not a Euclidean vector space, so these steps cannot be applied directly in the same way to sets of alternative phylogenies. Tree-space can be equipped with various metrics that allow geometry to be performed, and for reasons described below, we use the geodesic metric [5]. Φ PCA then follows a similar set of steps to those above, but working with the geodesic metric, $d(\cdot, \cdot)$. In step (2), the lines L become paths in tree-space with the property that, for any pair of points on the path, the path coincides with the geodesic between the points. Trees x_1, \dots, x_n are “projected” onto each path by finding points y_i on the path that minimize the distance $d(x_i, y_i)$ for $i = 1, \dots, n$. Pythagoras’ theorem does not hold in tree-space, so in the analog of step (3), paths which maximize the variance can be different from paths which minimize the sum of squared distances. We consider searching for both types of paths. Step (3) is potentially excessively computationally demanding, and so we describe (i) a greedy algorithm for constructing optimal paths and (ii) a Monte Carlo optimization approach. The methods we propose for steps (2) and (3) form the novel contribution of this paper. “Projection” of points onto a geodesic path L in step (2) is relatively simple to perform using existing methods for computing the geodesic metric, but a detailed algorithm has not been given previously. Searching over the set of possible paths is more technically demanding. Consideration of this particular problem and the solutions we present appear to be entirely novel.

The development of Φ PCA has been influenced by a recent paper by Wang and Marron [29]. Wang and Marron addressed a similar problem, developing a form of PCA for data sets with a tree-like structure. In a second paper [3], they applied their method to sets of trees obtained from medical imaging data. In particular, their reformulation of PCA in terms of the geometrical steps specified above motivated the corresponding steps in Φ PCA. Other authors have also developed analogs of PCA in nonstandard geometries [14, 18], and Wang and Marron give an excellent overview of this area of research [29]. However, it must be stressed that the method of Wang and Marron does not apply to sets of phylogenetic trees, and that Φ PCA is not simply a reworking of their approach. On account of the ostensible similarities between the approaches, we devote a section to explaining the relationship between them later in the paper.

The remainder of the paper is structured in the following way. We first describe the geometry of tree-space and set up necessary notation and mathematical background. Section 2 contains a description of the Φ PCA approach and proofs of its properties. We then explain more fully the relationship to the work of Wang and Marron, before evaluating Φ PCA on simulated sets of trees and a real set of gene trees from metazoan species.

1. Background: The geometry of tree-space.

1.1. *Splits and vector representation of trees.* We will work throughout with a fixed set of taxa $O = \{o_1, \dots, o_m\}$ and the set of unrooted phylogenetic trees \mathcal{T}_O on O . Given a tree $x \in \mathcal{T}_O$, cutting any branch on x partitions the taxa into two unordered nonoverlapping sets. Such a partition is called a *split*, and splits are usually denoted $X|X^c$ where $X \subset O$ and c denotes the complement in O . There are $M = 2^{m-1} - 1$ possible (nonempty) splits of the set O , and the set of these is denoted S . It is crucial to note that arbitrary sets of splits do not generally correspond to valid tree topologies—a compatibility condition must be satisfied. For example, if $O = \{A, B, C, D, E\}$, then the two splits $\{A, B\}|\{C, D, E\}$ and $\{A, C\}|\{B, D, E\}$ cannot both be represented on the same tree.

Any tree $x \in \mathcal{T}_O$ can be regarded as a weighted set of compatible splits, where the weight assigned to each split is given by the length of the corresponding branch on x . We only consider trees with positive branch lengths. We write T_x to denote the set of splits in x , and encapsulate the branch lengths via a function $\lambda_x : S \rightarrow \mathbb{R}^+$ defined by

$$\lambda_x(p) = \begin{cases} \text{branch length associated with } p, & \text{if } p \in T_x, \\ \text{zero,} & \text{otherwise.} \end{cases}$$

Tree-space \mathcal{T}_O can then be embedded in \mathbb{R}^M in the following way. Take the standard basis of \mathbb{R}^M and associate each split $p \in S$ with a different basis vector \mathbf{e}_p . Any tree $x \in \mathcal{T}_O$ can then be associated uniquely with the vector

$$(1.1) \quad \boldsymbol{\lambda}_x = \sum_{p \in S} \lambda_x(p) \mathbf{e}_p.$$

In fact, it is convenient to abuse notation slightly and write \mathbf{p} for the basis vector \mathbf{e}_p , identifying each split directly with the corresponding vector in \mathbb{R}^M . Equation (1.1) essentially associates every tree x with a vector of branch lengths, but due to the compatibility relations between splits, not every such vector corresponds to a tree. In fact, each tree contains at most $2m - 3$ splits, so as the number of taxa m increases, $2m - 3 \ll M$ and \mathcal{T}_O becomes an increasingly sparse subset of \mathbb{R}^M .

Since a collection x_1, \dots, x_n of trees can be regarded as a set of vectors $\boldsymbol{\lambda}_{x_1}, \dots, \boldsymbol{\lambda}_{x_n}$, why not just perform PCA on these vectors? In general, the principal components obtained in this way will not correspond to valid trees, and interpretation of the principal components becomes impossible. A form of PCA which operates intrinsically within \mathcal{T}_O and which produces interpretable “components” is required.

1.2. *Decomposition of tree-space by topology.* The geometry of \mathcal{T}_O was first comprehensively studied in a paper by Billera et al. [5], which included the definition and proof of existence of geodesics. Their description of \mathcal{T}_O amounts to a

decomposition into a set of overlapping component pieces, each piece corresponding to a different tree topology. In this section we recall aspects of this decomposition which are central to Φ PCA, most importantly for the definition of geodesics on \mathcal{T}_O .

The decomposition is easiest to understand by identifying \mathcal{T}_O with its image under the embedding in \mathbb{R}^M . Every tree in \mathcal{T}_O contains the set of splits corresponding to terminal edges (those that end in a leaf), denoted $S_{\text{term}} \subset S$. Since every tree contains every terminal split

$$\mathcal{T}_O \cong \text{span}_+ \{\mathbf{p} : p \in S_{\text{term}}\} \times \mathcal{T}_{O,\text{int}},$$

where span_+ denotes the span of vectors with nonnegative weights, and $\mathcal{T}_{O,\text{int}}$ is the part of tree-space corresponding to internal splits. Next consider a single unrooted tree x which is fully resolved, by which we mean every internal vertex has exactly 3 neighbors. Let t denote the topology of x or, more precisely, the set of nonterminal splits $t = T_x \setminus S_{\text{term}}$. Since x is fully resolved, it has $m - 3$ internal edges, so t contains $m - 3$ splits. The internal branch lengths of any tree with topology t are determined by a point in $Q_t = \text{span}_+ \{\mathbf{p} : p \in t\}$. We call Q_t the *topological orthant* containing x , and it is isomorphic to the positive orthant of \mathbb{R}^{m-3} . The faces of the orthant Q_t correspond to trees that have some zero length branches. Such trees are not fully resolved or, in other words, some internal vertices have more than 3 neighbors. This structure is illustrated in Figure 1.

Tree-space \mathcal{T}_O is formed from the union of the orthants Q_t over all possible fully resolved topologies t :

$$\mathcal{T}_O = \text{span}_+ \{\mathbf{p} : p \in S_{\text{term}}\} \times \bigcup_{\substack{\text{resolved} \\ \text{topologies } t}} Q_t.$$

The individual orthants Q_t are stitched together along their faces, since each unresolved tree occurs on the face of more than one orthant. To understand how the orthants are stitched together in more detail, consider a point on the face of an orthant Q_t at which a single branch length corresponding to a split p has been collapsed to zero. As illustrated in Figure 1, there are two ways in which this branch can be replaced with an alternative, thereby obtaining a fully resolved tree with a different topology. Each $(m - 4)$ -dimensional face of Q_t is therefore identified with corresponding faces in two other orthants $Q_{t'}$ and $Q_{t''}$. The operation illustrated in Figure 1 is called *Nearest Neighbor Interchange* (or NNI); we say that topologies t' and t'' are obtained by NNI of the split p within t . Faces of Q_t with co-dimension greater than 1 will be contained in more than two other orthants. Later, we will need to deal with paths in \mathcal{T}_O between such faces and so we need to extend the definition of NNI (which is usually taken as a relationship between strictly binary trees). Given a split p in a fixed tree, there are two or more subtrees hanging from each end of the associated edge in the tree. An *extended* NNI move (or XNNI) consists of swapping a subtree from one end of the branch with

a subtree from the opposite end. This operation removes split p from the tree and replaces it with an incompatible split p' . On a binary tree this definition coincides with the standard definition of NNI [1], and XNNI includes all NNI moves.

1.3. *Geodesics and the geodesic metric.* \mathcal{T}_O can be equipped with metrics via the embedding into \mathbb{R}^M described above. In particular, L_2 norm on \mathbb{R}^M defines a metric: $d_2(x, y) = |\lambda_x - \lambda_y|_2$. However, such metrics are not intrinsic to tree-space. For example, when x and y have different topologies, d_2 corresponds to the length of a straight line segment joining x to y through \mathbb{R}^M , but this line contains points outside the image of \mathcal{T}_O under the embedding.

Billera et al. [5] proved the existence of a metric that locally resembles the L_2 metric, but which is intrinsic to \mathcal{T}_O independent of the embedding in \mathbb{R}^M . This metric is called the *geodesic* metric d , and it is the canonical metric for Φ PCA due to its intrinsic nature. It is defined as follows. For two trees x and y with the same topology, $d(x, y) = d_2(x, y)$. When x and y have different topologies, $d(x, y)$ is defined as the length of the shortest continuous path joining x to y in \mathcal{T}_O which consists of a series of straight line segments through any feasible sequence of topological orthants. The length of such a path is defined to be the sum of the Euclidean lengths of each constituent line segment. The shortest such path joining x to y is called the *geodesic* between x and y . The proof that geodesics exist between points in \mathcal{T}_O and that geodesics define a valid metric is given in [5]. As part of the proof, Billera et al. [5] showed that tree-space is CAT(0) [16]. This means that triangles in \mathcal{T}_O are “skinny” in comparison to triangles in the Euclidean plane. More formally, given points $x, y, z \in \mathcal{T}_O$, consider the triangle between points x', y', z' in the Euclidean plane with the same edge lengths, so that $d(x, y) = d(x', y')$, etc. If $\gamma(t)$ is the path-length parameterized geodesic between x and y and $\gamma'(t)$ the corresponding geodesic in the Euclidean plane, then $d(z, \gamma(t)) \leq d(z', \gamma'(t))$ for all points $\gamma(t)$ between x and y .

Geodesics in \mathcal{T}_O have the following properties. First, if $x, y \in \mathcal{T}_O$ have the same topology t , then the geodesic joining them is the obvious Euclidean line segment in Q_t . Second, when x and y have some but not all splits in common, the splits in the intersection $T_x \cap T_y$ are all included at every point along the geodesic. The length of the branch associated to $p \in T_x \cap T_y$ changes in the obvious linear way from $\lambda_x(p)$ to $\lambda_y(p)$. Third, when x and y have different topologies, the geodesic may pass through other topological orthants than the two associated with x and y , as illustrated by Figure 2. This is the case for points x_1 and y_1 in the figure. However, trees along the geodesic only ever contain splits from $T_x \cup T_y$, albeit in different combinations. It follows that when x and y have different topologies, computing the geodesic distance $d(x, y)$ is nontrivial. However, an efficient polynomial-time algorithm has been developed for constructing geodesics [25], and we use this algorithm to calculate distances in Φ PCA.

A crucial feature of CAT(0) spaces is that paths which are everywhere locally geodesic are necessarily globally geodesic (see [25], Lemma 2.1). Geodesics like

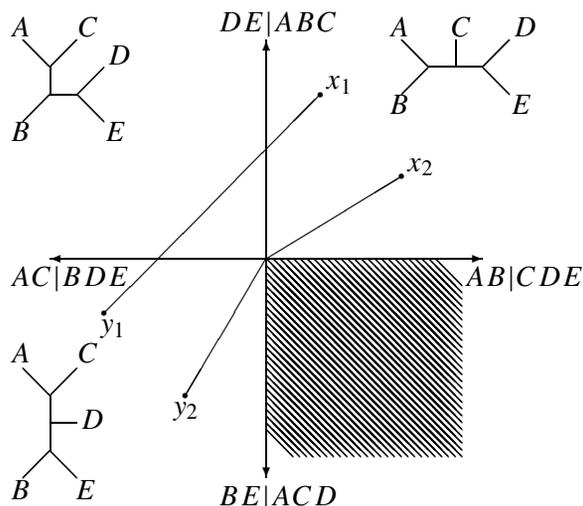


FIG. 2. Geodesics in tree-space consist of line segments through different topological quadrants. On five taxa there are 15 different quadrants, but only three are shown above, each with a sketch of the corresponding topology. Each axis corresponds to the length of a different split. The shaded region does not correspond to a valid quadrant since the splits $AB|CDE$ and $BE|ACD$ are incompatible. The geodesic between x_1 and y_1 passes through three quadrants, whereas the geodesic between x_2 and y_2 passes through just two quadrants. In this case the geodesic is the same as the cone path.

that between x_1 and y_1 in Figure 2 must therefore not “bend” as they cross between different orthants. For some pairs x, y , however, the shortest path is given by collapsing branch lengths for splits in $T_x \setminus T_y$ down to zero, so that the topology is then $T_x \cap T_y$ followed by expanding out branch lengths in $T_y \setminus T_x$ to obtain y . Any two points can be joined by such a path, and they are referred to as *cone paths*. In Figure 2 the cone path coincides with geodesic for points x_2 and y_2 ; geodesic between x_1, y_1 is clearly shorter than the cone path.

2. Methods.

2.1. *Existence of principal paths.* We now have the geometrical ingredients needed to define the Φ PCA procedure. Φ PCA seeks to construct a principal path from the set of \mathcal{T}_O -lines defined as follows.

DEFINITION 1. A path Γ in \mathcal{T}_O is a \mathcal{T}_O -line if:

- (i) every sub-path of Γ is the geodesic between its endpoints, and
- (ii) Γ extends to infinity in two directions.

We will often just use the term *line* to mean a \mathcal{T}_O -line where the context is obvious. Results in [5] show that any geodesic can be extended into a \mathcal{T}_O -line (though often not uniquely). The following proposition establishes existence and uniqueness of closest points on lines.

PROPOSITION 2.1. *Given a \mathcal{T}_O -line L and a point $x \in \mathcal{T}_O$, there is a unique closest point $y \in L$ to x .*

PROOF. The proof relies mainly on the CAT(0) property to enable comparison with Euclidean space. Let x_0 be any point on L and suppose $L(t)$ is a path-length parameterization of L such that $L(0) = x_0$. Defining $r = d(x_0, x)$, consider the triangle $x_0, x, L(t)$ for some $t > r$. The “skinny” triangle property implies that

$$d(x, L(r)) < d(x, L(t)).$$

The same bound applies to $L(-t)$. The closest point $y \in L$, if it exists, must therefore lie on $L(t)$ for $t \in [-r, r]$. Since this is a compact set and since the geodesic distance is a continuous function, $d(x, L(t))$ achieves its minimum on the interval. To prove uniqueness of the closest point y , suppose two distinct points $y, y' \in L$ achieve the same minimum distance ρ . Again, the “skinny” triangle property for the triangle y, x, y' implies that points on L between y and y' are closer to x than distance ρ . This is a contradiction, so y is unique. \square

Now suppose we are given a set of points $x_0, x_1, \dots, x_n \in \mathcal{T}_O$. For every line L through x_0 we can obtain the projection y_1, \dots, y_n of x_1, \dots, x_n onto L . This defines two functions, $f_{\parallel}(L)$ and $f_{\perp}(L)$, which are, respectively, defined as the sum of squared distance along L , $\sum d(x_0, y_i)^2$, and the sum of squared distances perpendicular to L , $\sum d(x_i, y_i)^2$.

PROPOSITION 2.2. *There is a \mathcal{T}_O -line through x_0 which maximizes f_{\parallel} . Similarly, there is a \mathcal{T}_O -line through x_0 which minimizes f_{\perp} .*

PROOF. We know from the proof of Proposition 2.1 that given any line L through x_0 , the points y_i are at most distance R from x_0 , where $R = \max\{d(x_0, x_i)\}$. Let S_R be the sphere $\{z \in \mathcal{T}_O : d(z, x_0) = R\}$. Each pair $(z, z') \in S_R \times S_R$ represents a pair of geodesics $\gamma(z, x_0)$ and $\gamma(x_0, z')$. If $d(z, z') = 2R$, then necessarily the geodesic between z and z' is exactly $\gamma(z, x_0)$ followed by $\gamma(x_0, z')$, and we say z, z' are *antipodal*. Every line L through x_0 determines an antipodal pair (z, z') , and since the projected points y_i all lie on the geodesic between that pair, f_{\parallel} and f_{\perp} only depend on the pair (z, z') . By continuity of the function $d : S_R \times S_R \rightarrow \mathbb{R}$, the set of antipodal pairs is a closed subset of $S_R \times S_R$ and is therefore compact. It follows that there is a geodesic between antipodal points on S_R which optimizes either f_{\parallel} or f_{\perp} . The geodesic can be extended into a line, and that establishes the proposition. \square

The optimal line may be nonunique for two reasons. First, different extensions of the geodesic between an antipodal pair (z, z') might exist. This would arise, for example, if all the points x_1, \dots, x_n lay in the same topological orthant. Second,

as in regular Euclidean PCA, the collection of points x_1, \dots, x_n can be isotropic, so that the optimal pair (z, z') is nonunique.

Given the existence of optimal \mathcal{T}_O -lines, we can now consider how to construct a principal line. As outlined in the [Introduction](#), construction of the principal line consists of the following steps:

- (1) Given trees x_1, \dots, x_n , construct a “central point” x_0 .
- (2) Given a line L through x_0 , “project” x_1, \dots, x_n onto L by finding the closest point y_i in L to x_i for $i = 1, \dots, n$.
- (3) Find the line L such which optimizes the particular choice of objective function f (either f_{\parallel} or f_{\perp}).

The details of each of these steps is described in turn, but step 3 forms the main challenge.

2.2. Centroids and consensus. Ideally, x_0 should be chosen so as to minimize the sum of squared distances:

$$(2.1) \quad x_0 = \arg \min_x \sum d(x, x_i)^2.$$

In a Euclidean vector space, this reduces to finding the mean of the data x_1, \dots, x_n . In tree-space, due to the lack of additive structure, the mean does not make sense, and there is no known closed solution to (2.1). Instead, Billera et al. [5] suggest using the *centroid*, which is defined via a recursive procedure based on finding the midpoint along the geodesic between any two points. However, for large data sets, this procedure is computationally demanding. We therefore propose taking x_0 to be the *majority consensus* tree [4]. Finding an “average” or consensus tree is a well-studied problem in phylogeny [7] and various forms of consensus tree exist. The majority consensus topology consists of splits which are found in strictly more than half the trees x_1, \dots, x_n . Branch lengths on x_0 are assigned their average value in the data set:

$$\lambda_{x_0}(p) = \frac{1}{|I(p)|} \sum_{i \in I(p)} \lambda_{x_i}(p)$$

for all $p \in T_{x_0}$ where $I(p)$ is the set $\{i : p \in T_{x_i}\}$. Results obtained later in this paper were obtained using this choice of x_0 . However, construction of the principal line L does not rely on any particular properties of the point x_0 , and Φ PCA works with any sensible choice.

2.3. “Projection” onto \mathcal{T}_O -lines. Given any line L and points x_1, \dots, x_n , Proposition 2.1 established the existence of closest points $y_1, \dots, y_n \in L$. Here we describe computational aspects of this “projection” onto L . Although this relies on existing mathematics, as presented in Section 1, the details of an algorithm for projection onto a geodesic path have not previously been given. We will assume

$L(t)$ is a path-length parameterization of L . For each point x_i , Euclidean projection under the embedding into \mathbb{R}^M described in Section 1.1 is used to obtain a first guess $L(s_i)$ for y . Amenta et al. [2] showed that the geodesic distance between two points is bounded by the Euclidean distance:

$$\|\lambda_x - \lambda_y\|_2 \leq d(x, y) \leq \sqrt{2} \times \|\lambda_x - \lambda_y\|_2.$$

It follows that if ε_i denotes the Euclidean distance between x_i and $L(s_i)$, then y_i lies on $L(s_i \pm \varepsilon_i)$. This bounding interval for y_i is used as the starting point for a golden-ratio search, which is iterated until some tolerance on y_i is achieved. It can be shown that finding y_i is a convex optimization, so the golden-ratio search is guaranteed to converge. The proof of convexity relies on the CAT(0) property and convexity of the equivalent Euclidean problem. The algorithm of Owen and Provan [25] is used to calculate geodesic distance during the golden-ratio search. However, it is not necessary to recompute geodesics from scratch at every iteration: the sequence of orthants for a geodesic at one iteration can often be reused in the next iteration, with an associated gain in computational efficiency.

In Euclidean vector spaces, Pythagoras’ theorem gives a decomposition of the total sum of squared distances $d_0^2 = \sum d(x_0, x_i)^2$ of a collection of points into contributions from directions perpendicular and parallel to any given line L . However, this decomposition does not apply in \mathcal{T}_O with the geodesic metric. Nonetheless, we can evaluate the quantities

$$d_{\perp}^2 = \sum d(x_i, y_i)^2 \quad \text{and} \quad d_{\parallel}^2 = \sum d(y_i, x_0)^2$$

for any metric. It can be shown that for the geodesic metric, unlike the Euclidean case, the sum of these two quantities depends on L . Despite this, when evaluated for a principal path L , the sums of squared distances provide a useful quantification of variability, as we demonstrate in the results sections.

2.4. *Lines through x_0 .* We need to construct \mathcal{T}_O -lines through x_0 and identify one which optimizes our choice of objective function, f . This is a challenging problem which has not previously been considered in the literature. In order to achieve computational tractability, we restrict to a particular class of \mathcal{T}_O -lines and then employ different optimization algorithms to search over the restricted class. To motivate this approach, we start by considering properties of lines through x_0 .

In the topological orthant containing x_0 , any \mathcal{T}_O -line L consists of a straight line segment. For brevity, we will write T_0 for the midpoint topology T_{x_0} and λ_0 for the branch length function λ_{x_0} . Let Q_0 denote the orthant containing x_0 , and for now assume that x_0 is a binary tree (so it contains the maximal number of splits). If p is a split contained in Q_0 , then the branch length at a point $y(s) \in Q_0$ on L has the form

$$(2.2) \quad \lambda_{y(s)}(p) = \lambda_0(p) + s \times w_p,$$

where w_p is a “weight” associated to split p , and s lies on some interval containing zero. The set of weights determines the direction vector of the line segment through x_0 . Given such a line segment, we need to know how it might extend beyond Q_0 into the rest of tree-space.

Where the segment meets a face of Q_0 , at least one split is assigned zero branch length. Generically, the line segment will meet a co-dimension 1 face of Q_0 , so that just one split p will have zero length. Solving equation (2.2) for this split gives

$$(2.3) \quad \lambda_{y(s)}(p) = 0 \quad \Rightarrow \quad s = -\lambda_0(p)/w_p.$$

The line then extends from this point into one of the neighboring alternative orthants. In a similar way, every other split whose length varies in the initial line segment containing x_0 is associated with a solution of equation (2.3) and, correspondingly, with an alternative split related to the first by NNI. If we restrict to the set of “generic” lines (those which always meet a co-dimension 1 face of every orthant), then finding the optimal line L therefore consists of a topological problem (namely, choosing a new split p' to replace each p) and a geometrical problem (finding the best set of weights w_p). However, these problems are not independent. We can order the solutions to (2.3) as we move out from x_0 in a particular direction along L . Suppose the first solution we come to is at $s = s_1$ and we replace split p_1 with p'_1 . At the next solution $s = s_2$, split p_2 is assigned zero length and we replace it via an NNI move. However, the choice of splits available as replacements for p_2 does not depend solely on p_2 but also on the rest of the tree topology just before $s = s_2$ —and therefore potentially on the choice of replacement p'_1 of p_1 . Thus, the topological aspect of construction depends on the relative order of the solutions to (2.3), which in turn depends on the weights w_p . Optimization over the set of possible weights and splits will be computationally demanding for trees with more than a few species—an exhaustive search will not be possible.

A key feature of the description above is the assumption that line segments meet the boundary of orthants in co-dimension 1 faces. We restrict our search space for L similarly, but take into account the possibility that x_0 might not be fully resolved. We make this more formal as follows.

DEFINITION 2. Suppose $p \in S$ is compatible with x_0 and $p' \in S$ is obtained by extended nearest neighbor interchange of p in $T_0 \cup \{p\}$. The *simple line* through x_0 associated with p, p' and weight w is the path $y(s) \in \mathcal{T}_O$ defined by

$$\begin{aligned} \lambda_{y(s)}(p) &= \lambda_0(p) + sw && \text{if } \lambda_0(p) + sw \geq 0 \\ &= 0 && \text{otherwise,} \\ \lambda_{y(s)}(p') &= -(\lambda_0(p) + sw) && \text{if } \lambda_0(p) + sw \leq 0 \\ &= 0 && \text{otherwise,} \\ \lambda_{y(s)}(q) &= \lambda_0(q) && \text{if } q \notin \{p, p'\}. \end{aligned}$$

Such a path moves through a single pair of orthants. The next definition extends simple lines to pass through more than two orthants.

DEFINITION 3. Suppose $x(s)$ is the simple line through x_0 defined by split pairs $(p_1, p'_1), \dots, (p_k, p'_k)$ and weights w_1, \dots, w_k , and suppose that the pair of splits (p_{k+1}, p'_{k+1}) and weight $w_{k+1} \in \mathbb{R}$ satisfy the following:

- (i) p_{k+1} is compatible with $x(s)$ for all s such that $\lambda_0(p_{k+1}) + sw_{k+1} \geq 0$,
- (ii) p'_{k+1} is compatible with $x(s)$ for all s such that $\lambda_0(p_{k+1}) + sw_{k+1} < 0$,
- (iii) p_{k+1}, p'_{k+1} are related by XNNI in $x(s_{k+1})$ where $s_{k+1} = -\lambda_0(p_{k+1})/w_{k+1}$.

Then the simple line $y(s)$ defined by $(p_1, p'_1), \dots, (p_{k+1}, p'_{k+1})$ and weights w_1, \dots, w_{k+1} is given by

$$(2.4) \quad \begin{aligned} \lambda_{y(s)}(p_i) &= \lambda_0(p_i) + sw_i && \text{if } \lambda_0(p_i) + sw_i \geq 0 \\ &= 0 && \text{otherwise,} \end{aligned}$$

$$(2.5) \quad \begin{aligned} \lambda_{y(s)}(p'_i) &= -(\lambda_0(p_i) + sw_i) && \text{if } \lambda_0(p_i) + sw_i \leq 0 \\ &= 0 && \text{otherwise,} \end{aligned}$$

$$(2.6) \quad \lambda_{y(s)}(q) = \lambda_0(q) \quad \text{if } q \notin \{p_i\} \cup \{p'_i\}, \quad i = 1, \dots, k + 1.$$

To prove that simple lines satisfy the conditions of Definition 1, Proposition 4.2 of [5] can be applied to any pair of points on a simple line in order to show that the subpath between those points is geodesic.

Simple lines through x_0 resemble the geodesic between x_1 and y_1 in Figure 2: they continue between orthants without bends, and hence are always locally geodesic. Moreover, any straight line segment through x_0 can be obtained as part of a simple line. Nonetheless, restriction to the class of simple lines removes many lines from consideration. Cone paths are ruled out, together with any geodesic for which some subset of the splits changes like a cone path. (This latter class of geodesics resembles $x_1 - y_1$ in Figure 2 for some splits and $x_2 - y_2$ for others.) This restriction is carried out for the sake of computational tractability. More discussion is given in Section 6.

Definition 3 describes how to extend a simple line on k split pairs to one on $k + 1$ split pairs. Our algorithms for finding an optimal simple line are based precisely on this operation. Suppose a simple line L is determined by sets of splits $P = \{p_1, \dots, p_k\}$, $P' = \{p'_1, \dots, p'_k\}$ and weights $W = \{w_1, \dots, w_k\}$. Conditions (i)–(iii) of Definition 3 place constraints on any proposed splits p, p' and weight w which might be used to extend L . The values $s_i = -\lambda_0(p_i)/w_i$ correspond to points at which L crosses the boundary between orthants, and we can assume they are ordered with $s_1 \leq s_2 \leq \dots \leq s_k$. They divide L up into $k + 1$ intervals $I_i = [s_i, s_{i+1}]$ for $i = 0, 1, \dots, k$ taking $s_0 = -\infty$ and $s_{k+1} = \infty$, such that the

topology of L is constant on each interval. Let t_i denote the tree topology on I_i . Now suppose that p is compatible with T_0 and p' is a proposed replacement for p . Suppose we also propose an interval I_i on which we require the XNNI move to be performed. Conditions (i)–(iii) are then equivalent to the following.

Geometrical constraint:

$$s_i \leq -\frac{\lambda_0(p)}{w} \leq s_{i+1}, \text{ so the XNNI move occurs on interval } I_i.$$

Topological constraint:

- If $w < 0$, then p must be compatible with t_0, \dots, t_i and p' must be compatible with $t_i \setminus p, \dots, t_{k+1} \setminus p$; or
- if $w > 0$, then p' must be compatible with $t_0 \setminus p, \dots, t_i \setminus p$ and p must be compatible with t_i, \dots, t_{k+1} .

When p is not contained in T_0 , but instead *extends* the midpoint topology, then $\lambda_0(p) = 0$ and the solution to equation (2.3) is $s = 0$. In this case, the geometric constraint corresponds to an unbounded interval for w , and the interval I_i on which the XNNI move $p \rightarrow p'$ takes place must necessarily contain $s = 0$.

2.5. *Greedy algorithm for finding an optimal simple line.* The following algorithm repeatedly extends a simple line by adding in a new split pair at each iteration. The pair chosen is the one which gives the best improvement in the objective f :

- (1) Let F be the set of feasible splits (see below).
- (2) Consider in turn every split p in F that is compatible with T_0 , and every possible replacement p' for p .
- (3) For each interval I_i , test whether p' is an XNNI replacement of p in t_i .
- (4) If (ii) holds on interval I_i , then next check whether the pair p, p' satisfies either topological constraint for that interval. Fix the sign of w depending on which constraint applies.
- (5) If either topological constraint holds, then find w that maximizes the variance of the projected points on L , subject to the geometrical constraint and sign of w . This is carried out using the golden ratio search for the optimum value of w .
- (6) Repeat for all feasible pairs $p, p' \in F$ and find the pair that gives the maximum projected variance.
- (7) Add p, p' and w to the lists P, P' and W , and reorder the lists according to the solutions of (2.3). Remove p, p' from F , and repeat from step 2.

The algorithm continues until no more splits can be added to L . This will occur in at most $m - 3$ iterations where m is the number of species, since every tree can contain at most $m - 3$ nontrivial splits.

The set of feasible splits F could be taken to be the entire set of possible splits S , but this is inefficient. If neither split p, p' is contained in any of the trees

x_1, \dots, x_n , then adding the pair to L will only increase the distances $d(x_i, y_i)$ so that L is a worse approximation to the data. We therefore take F to be the set of nontrivial splits found in at least one tree x_1, \dots, x_n . It is possible that at some stage the best improvement in f might be given by some $p \in F$ and a replacement $p' \notin F$ (e.g., consider the case that all the trees x_i lie in the same orthant). However, in such a situation, the data would not be informative about the choice of p' , and so we disregard this possibility.

The greedy algorithm terminates after at most $m - 3$ iterations. During each iteration $O(|F|^2)$ pairs of splits and $O(m)$ possible intervals for the move $p \rightarrow p'$ are considered. For each pair of splits and interval, n trees are projected onto the proposed line. Each projection requires $O(m^4)$ steps. The golden ratio search during projection is performed to a fixed tolerance, and so is independent of m, n and $|F|$. Overall, the algorithm therefore requires $O(m^6 \times n \times |F|^2)$ steps where F is at worst $O(nm)$.

2.6. Monte Carlo optimization algorithm. A Monte Carlo optimization algorithm was also implemented in order to provide comparisons with the greedy approach. A simulated-annealing type approach was adopted, where the state at each iteration comprised a simple line L through x_0 . At each iteration a “birth” or “death” move was randomly proposed from the current state. Birth moves consisted of adding a valid split pair to L , while death moves consisted of removing a split pair from L . Birth moves were obtained by selecting $p \in F$ uniformly at random, then selecting p' uniformly at random from the possible XNNI replacements of p satisfying the constraints defined above. The weight assigned to p, p' was obtained by the golden ratio search, as for the greedy approach. Death moves were carried out by choosing at random the split pair at either end of L (i.e., with largest positive or negative s_i) and removing it. Removing other split pairs results in incompatible sets of splits along L and is therefore forbidden. The relative probabilities of birth and death depended on the number k of split pairs in L and were designed to favor birth for small k and death when k was large. Proposals leading to improvement in the objective were always accepted. Other proposals were accepted with probability

$$\Pr(\text{accept}) = \left(1 - \frac{\delta}{D}\right)^{1/\tau},$$

where δ is the absolute difference of the objective for the proposed and current state, D is a bound for δ , and τ is the “temperature.” For $f = f_{\perp}$, D was taken to be $\sum d(x_0, x_i)^2$, while for $f = f_{\parallel}$, D was taken to be f_{\parallel} for the current state. The temperature τ was slowly decreased as the optimization progressed.

2.7. Branch length transformations. We investigated certain transformations of the data x_1, \dots, x_n prior to analysis with Φ PCA. Kupczok et al. [22] suggest scaling each tree x_1, \dots, x_n to have the same total branch length. In practice, this

seemed to make little difference to the examples we looked at in the results section below. Instead we considered the following branch length normalization. For each split p , branch lengths were scaled by a constant so that the average branch length associated with p across the whole data set was unity. This was repeated for each split in the data set. The idea behind this is to make Φ PCA measure variability relative to branch length and to amplify the variability in short branches. In regular PCA the correlation matrix can be analyzed instead of the covariance matrix, and this branch length transformation can be thought of as being analogous to the correlation matrix version. Principal geodesics obtained for branch-length normalized data can be back-transformed onto the original scale by scaling the weights W .

3. Relationship to the work by Wang and Marron. Wang and Marron [29] previously developed PCA in a space of trees, and on account of the similarities of our approach to theirs, in this section we look in detail at the relationship between the approaches. The steps underlying our approach specified at the start of Section 2 were taken directly from [29], but the details of how these steps are carried out are quite different on account of the different geometries under consideration.

In [29] rooted bifurcating trees are considered, but, unlike phylogenetic trees, the leaf vertices are not assigned taxon labels. Instead, each vertex can have a “left” and a “right” descendant, and trees in the data set can have different depths from root to leaf. Most importantly, branches do not have any associated length, but, instead, each vertex present in a tree has an associated real number (or vector). An example of such data consists of blood vessel information from medical imaging: vertices represent blood vessels, edges represent connections between blood vessels, and the data value associated to each vertex corresponds to some measurement at that point in the blood vessel structure. One crucial difference between the two spaces of trees is that in [29] there is no relationship between the values associated to vertices and the topological structure of the tree. This is different from the space \mathcal{T}_O , in which branch lengths can be shrunk down and replaced by an alternative topology.

This separation of “topological” and “geometrical” aspects of the problem in [29] results in principal components with separate topological and geometrical parts. In Wang and Marron’s terminology, a *structure tree line* is a sequence of vertices, each descended from the previous vertex, which can be thought of as (discontinuous) “growth” of a tree toward a leaf, by grafting on branches. In contrast, an *attribute tree line* consists of a fixed tree structure with “direction vectors” associated to vertices. This is clearly very different from the lines constructed by Φ PCA in which the principal path reflects both topological and geometrical variability in the data set.

Not surprisingly, given the different structures of the spaces considered, the metrics used in the two approaches differ. The metric in [29] is a linear combination of the (unweighted) Robinson–Foulds metric [27] and a Euclidean distance between the vectors associated to each vertex. This metric is inexpensive to compute, in

contrast to the geodesic metric which we consider, and this reduces the computational burden of their approach relative to ours. The midpoint x_0 in [29] is taken to have the majority consensus topology [4], as used here, since this minimizes the sum of the Robinson–Foulds distances of the midpoint from x_1, \dots, x_n . However, while Wang and Marron obtain an exact form of Pythagoras’ theorem with their metric, that is not the case for Φ PCA (see Section 2.3).

In summary, the method of Wang and Marron cannot be applied directly to phylogenetic trees, since the trees they consider lack taxon labels and branch lengths and so cannot represent phylogenies. While our approach builds on the same framework as that laid out in [29], differences in the geometries of the spaces under consideration make the mathematical details of the implementation of PCA substantially different. In particular, Φ PCA relies heavily on the geometry of \mathcal{T}_O described by Billera et al. [5]. It is interesting to note how a seemingly small difference in the geometry of the space under consideration can substantially change the way PCA is implemented.

4. Simulation studies.

4.1. *Simple mixtures.* Φ PCA was used to analyze collections of randomly generated trees with two (or more) known underlying topologies. These simulations were not intended as a model of a specific process giving rise to alternative trees, but were performed in order to verify the methodology and demonstrate how it works on simple examples. We describe the simulations very briefly here, but give more details in the supplementary material [24]. Two sets of simulations were performed. In the first set, trees were simulated such that each had one of two possible topologies t_1 or t_2 . The underlying topologies t_1, t_2 were related by an NNI move and represented alternative positions for a clade within the tree. Topology t_1 was adopted with probability θ and t_2 with probability $1 - \theta$. Apart from branches affected by the change in topology, all other branch lengths were kept fixed. For each value of θ , 100 trees were randomly generated in this way. Additional variability was added by simulating a DNA alignment for each tree, and then replacing the tree with the maximum likelihood (ML) tree estimated from the alignment. Φ PCA was used to analyze these estimated trees. A second set of simulations was performed in which there were two correlated changes in topology. Each tree consisted of two subtrees, and each subtree had either topology t_1 or t_2 as in the first set of simulations. The alternative topologies in each half of the tree were simulated to arise with correlation ρ . Again, additional variability was added by simulating alignments and replacing each tree with an ML estimate. 100 trees were generated for each pair of values θ, ρ and Φ PCA was used to analyze each set of estimated ML trees.

The results indicated that optimization of f_{\parallel} gave the best performance: paths obtained by optimizing f_{\perp} sometimes missed the changes in topology imposed in the data sets. In this non-Euclidean setting the sum $d_{\parallel}^2 + d_{\perp}^2$ is generally less than

the total sum of squared distances d_0^2 , so optimization of f_{\perp} may result in principal paths that fail to capture variability in the data by finding paths in which both sums d_{\parallel}^2 and d_{\perp}^2 are small.

In both sets of simulations, Φ PCA with f_{\parallel} gave principal paths corresponding to the change between the imposed alternative topologies. In the first set of simulations, based on a single pair of alternative topologies, as θ increased the change between the underlying topologies t_1, t_2 dominated the principal path (the corresponding splits received a higher weight) as variability due to tree estimation from alignments was dominated by the imposed variability in topology. In the second set of simulations, for small values of ρ the principal path corresponded to change between the alternative topologies on one part of the tree, with the other pair of alternative topologies receiving a low weight. For larger ρ , the correlated alternatives in both parts of the tree were identified by the principal path. More details are given in the supplementary material [24].

4.2. Long branch attraction. In order to demonstrate a potential application of Φ PCA, a simple study of long branch attraction (LBA) was performed. LBA is a feature of phylogenetic methods in which species on long branches are often grouped together erroneously on estimated trees. We took a tree from the literature [6] representing a deep phylogeny of eukaryote species which includes two long branches and a distant out-group, as shown in Figure 3. 100 trees were simulated by first simulating amino acid alignments from the base tree (300 base pairs, WAG+4 Γ model) using the seq-gen software [26] and then obtaining an ML estimate tree from each alignment using phyML [17]. Φ PCA was used to analyze the set of trees estimated from the simulated alignments.

Analysis of the simulated trees was carried out first with un-normalized data and then again with the normalization procedure described in Section 2.7. Optimization was carried out using the f_{\parallel} objective function, and results were obtained with both the greedy and Monte Carlo algorithms. Despite long runs, the Monte Carlo algorithm failed to improve on the results obtained with the greedy algorithm. The results obtained with the greedy algorithm are shown in Table 1 and Figure 3 shows the principal path obtained with normalized branch lengths. The “proportion of variance” d_{\parallel}^2/d_0^2 was greater with the normalized data, and so we suggest that normalization is preferable for this data set.

As explained in [6], estimated trees tend to place the long branches (*Guillardia* and Microsporidia) next to the out-group (Archaea). The analyses of both the un-normalized and normalized data show this effect with, respectively, *Guillardia* and Microsporidia “floating” round the tree to be placed closer to the Archaea. The fact that each principal path captures a single such effect suggests that the attraction of the two long branches to the Archaea is uncorrelated in the data. Φ PCA exactly captures the expected LBA artefact in the simulated data.

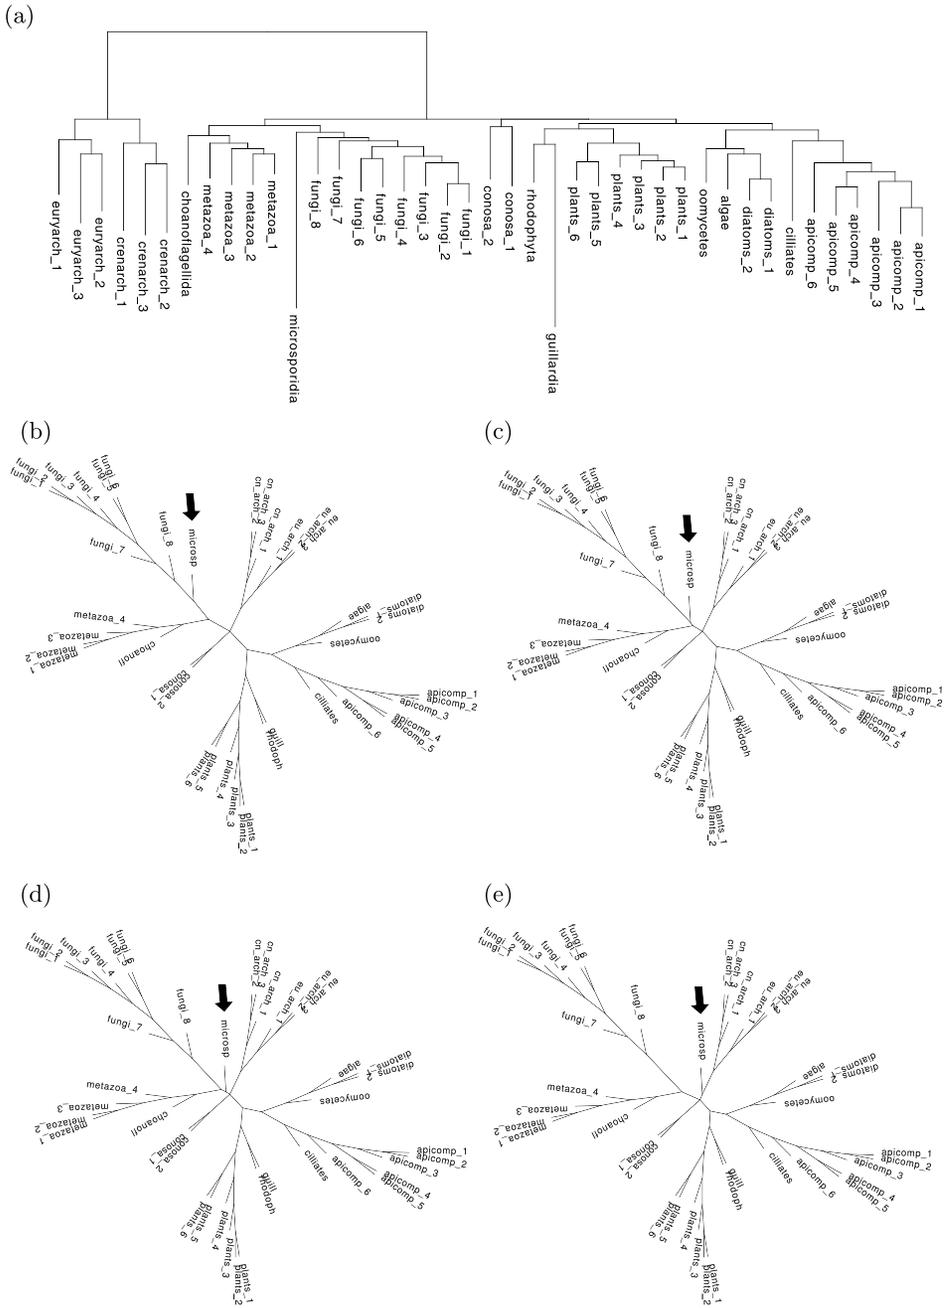


FIG. 3. Simulation study of LBA. (a) Underlying tree with two long branches and distant out-group (archaea, on the left). (b)–(e) Trees along the principal path. Branches were normalized to have unit mean. No back-transform to the original scale was performed, since this obscured the visual effect. Arrows highlight the microsporidia group (labeled “microsp”) moving round to join the outgroup (labeled “cn_arch” and “eu_arch”).

TABLE 1

Results of the LBA simulations. The largest two weights w for split pairs and a description of the corresponding changes in topology are given. Weights were normalized to have unit Euclidean norm

w	Change in topology
<i>Without normalization ($d_{\parallel}^2/d_0^2 = 2.6\%$)</i>	
0.859	Guilardia moves past pairing with Rhodophyta to top of clade with plants
0.152	Guilardia moves from top of clade with plants to position closer to Archaea
<i>Branch lengths normalized ($d_{\parallel}^2/d_0^2 = 10.3\%$)</i>	
0.706	Microsporidia moves from grouping with fungi to top of clade with Metazoa
0.482	Microsporidia grouped with Archaea

5. Analysis of metazoan data. Φ PCA was applied to a set of 118 gene trees from 21 metazoan (animal) species, previously analyzed in [22]. Φ PCA was performed on both unscaled and branch length normalized data using the f_{\parallel} objective function. The Monte Carlo optimization algorithm obtained principal paths with slightly higher f_{\parallel} scores than the greedy algorithm, and so we refer to that set of results here. The principal paths obtained with the two algorithms were similar, and shared the majority of split pairs in common. The “proportion of variance” d_{\parallel}^2/d_0^2 was 1.8% for the unscaled data and 4.6% for the normalized data—relatively low in both cases. However, the simulation studies produced similarly low scores (between 3% and 5% on artificial data), suggesting that low scores might be common even when Φ PCA is successfully capturing aspects of variability in the data. Further comments about the low proportion of variance are made in Section 6.

The principal path obtained for the unscaled data corresponded to uncertainty in the positioning of the out-group, yeast. It moves from being placed next to the worms to being grouped with sea squirt. This might be an LBA effect since sea squirt and yeast lie on relatively long branches. Results of the analysis using data with normalized branch lengths are shown in Figure 4. The principal path indicates uncertainty in the placement of Human: it is either grouped with Chimpanzee or Macaque. The position of Human relative to its neighbors was a longstanding problem in phylogenetics [15]. Uncertainty in the positioning arises from the presence of a relatively short internal branch in the species tree joining Human and Chimpanzee to the other primates. Although well known in evolutionary biology, this simple example illustrates how Φ PCA can be used to identify and visualize alternatives within a set of trees.

6. Conclusion. We have presented a procedure for identifying principal paths in the space of phylogenetic trees which best approximate a set of alternative phylogenies in an analogous way to standard PCA in Euclidean vector spaces. A key feature of the approach is the use of metrics that combine geometric and topological information about trees. The principal paths constructed coincide with the

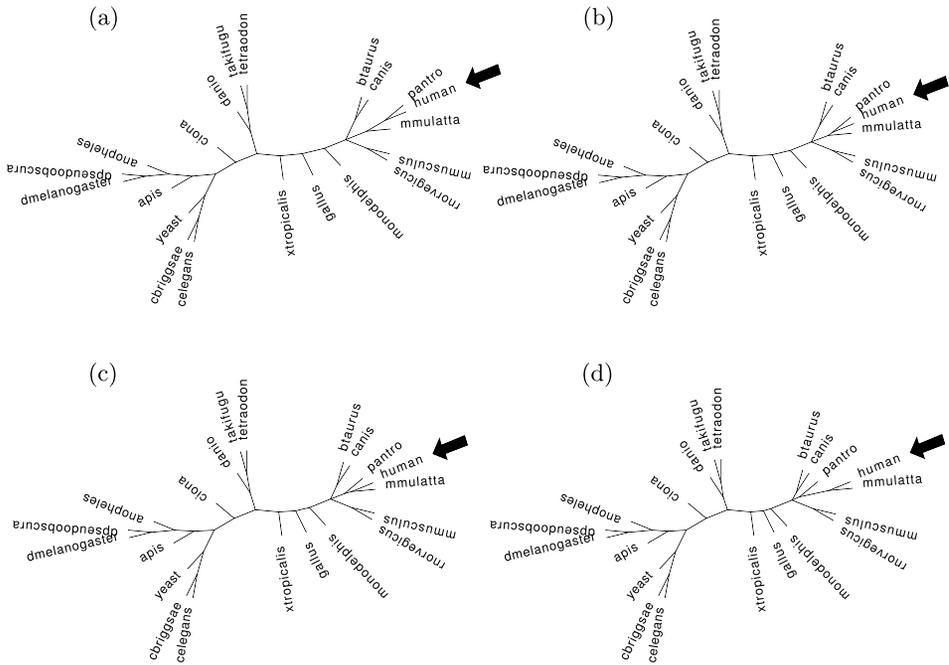


FIG. 4. Trees along the principal path for the normalized metazoan data in order (a)–(d). (a) corresponds to the majority consensus topology. Human moves from being grouped with Chimp (labeled “panthro”) to Macaque (labeled “mmulatta”), as highlighted by the solid arrows.

geodesic between every pair of points on the path. Each principal path is equipped with a summary statistic analogous to the Euclidean proportion of variance which quantifies variability along the path.

Results obtained from simulated and experimental data sets gave values for the “proportion of variance” d_{\parallel}^2/d_0^2 which were relatively low in comparison to typical values for standard PCA (e.g., about 5% for the normalized metazoa data set). This is the result of two features of the problem. First, the data sets analyzed in this paper are high dimensional (containing over 100 different splits), and in the same way as for standard Euclidean data, this tends to lead to lower proportions of variance. To illustrate this, consider, for example, a multivariate normal distribution with dimension 100 and covariance matrix $\text{diag}(5\sigma^2, \sigma^2, \dots, \sigma^2)$. Standard PCA would give a proportion of variance of roughly 5% even though the variance along the principal component is substantially higher than in other directions. Second, the failure of Pythagoras’ theorem in tree-space means that for any analysis variance “leaks out,” that is, the sum of squares $d_{\parallel}^2 + d_{\perp}^2$ is less than the total sum of square distances for the original data d_0^2 , further decreasing d_{\parallel}^2/d_0^2 .

In order to construct principal paths, two approximations have been imposed:

(1) A greedy algorithm or Monte Carlo search is carried out in the configuration space of paths in order to find the optimal path. There is no guarantee that the optimal path will always be found.

(2) The configuration space itself is restricted to a subclass of paths (referred to as simple lines). By restricting in this way we might rule out capturing types of variability in the data under analysis.

We consider the second of these approximations to be more limiting, and it is difficult to generalize the approach we have described to overcome it. Another area requiring further research is the construction of higher-dimensional approximations to data, analogous to second and third, etc. principal components. Construction of higher order principal components in Euclidean vector spaces is carried out by working orthogonally to the first principal component. Tree-space is not equipped with an inner product, and so this procedure cannot carry over directly to \mathcal{T}_O . Our algorithms for constructing the principal path L do not therefore readily generalize to give higher order paths. Instead, we would need to consider two-dimensional subsets of \mathcal{T}_O which approximate the data x_1, \dots, x_n as closely as possible. In analogy to the definition of \mathcal{T}_O -lines, we would require that such subsets Π contain the geodesic between any two points in Π , and so Π would locally resemble a plane in each orthant. However, in contrast to the theory of geodesics, the theory of higher-dimensional surfaces in tree-space is not well developed. We have not attempted to advance this theory in this paper, but have focused on the already considerable problem of identifying lines which best approximate the data.

The Φ PCA procedure has been presented as an empirical analysis of sampled trees without reference to any underlying distribution that generated the trees. Distributions such as sampling distributions, bootstrap distributions and Bayesian posteriors are of fundamental importance in phylogenetic inference, but the geometrical properties of these distributions have received little study. Billera et al. [5] considered spherically symmetric distributions with density decaying exponentially away from a central point. A second form of isotropic distribution consists of the limit of a random walk in tree-space from a fixed central point. By simulating samples from a suitable random walk and carrying out Φ PCA on the samples, an empirical p -value could be assigned to the proportion of variance of a principal line constructed from experimental data, as a test for significant departure from isotropy. Holmes [20], however, suggests that the assumption of spherical symmetry is not realistic for most distributions of interest. One area where distributions on tree-space have been defined more precisely is the study of population-genetic effects on gene phylogenies [9]. Such distributions could be studied in the context of tree-space geometry, and it might be possible to obtain the sampling theory of principal lines under Φ PCA in this case.

This paper has presented the results of applying Φ PCA to some relatively simple examples, and demonstrated the type of information principal paths reveal. The method can be applied to larger data sets and it has the potential to provide new

insights into a range of problems in evolutionary biology. Software for performing Φ PCA and for visualizing principal paths as animations of trees is available in the supplementary material [24], together with the data sets analyzed in this paper.

Acknowledgments. The author would like to thank Anne Kupczok and the group of Arndt von Haeseler for generously providing the set of metazoan gene trees. I am also grateful to two anonymous reviewers for their very helpful comments.

SUPPLEMENTARY MATERIAL

Principal components analysis in the space of phylogenetic trees: Supplementary information (DOI: [10.1214/11-AOS915SUPP](https://doi.org/10.1214/11-AOS915SUPP); .pdf). This contains further information about the simulation studies in Section 4.1.

REFERENCES

- [1] ALLEN, B. L. and STEEL, M. (2001). Subtree transfer operations and their induced metrics on evolutionary trees. *Ann. Comb.* **5** 1–15. [MR1841949](#)
- [2] AMENTA, N., GODWIN, M., POSTARNAKEVICH, N. and ST. JOHN, K. (2007). Approximating geodesic tree distance. *Inform. Process. Lett.* **103** 61–65. [MR2322069](#)
- [3] AYDIN, B., PATAKI, G., WANG, H., BULLITT, E. and MARRON, J. S. (2009). A principal component analysis for trees. *Ann. Appl. Stat.* **3** 1597–1615. [MR2752149](#)
- [4] BARTHÉLÉMY, J.-P. (1986). The median procedure for n -trees. *J. Classification* **3** 329–334. [MR0874243](#)
- [5] BILLERA, L. J., HOLMES, S. P. and VOGTMANN, K. (2001). Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.* **27** 733–767. [MR1867931](#)
- [6] BRINKMANN, H., VAN DER GEIZEN, M., ZHOU, Y., PONCELIN DE RAUCOURT, G. and PHILIPPE, H. (2005). An empirical assessment of long-branch attraction artefacts in deep eukaryotic phylogenomics. *Syst. Biol.* **54** 743–757.
- [7] BRYANT, D. (2003). A classification of consensus methods for phylogenetics. In *Bioconsensus (Piscataway, NJ, 2000/2001)*. DIMACS Series in Discrete Mathematics and Theoretical Computer Science **61** 163–183. Amer. Math. Soc., Providence, RI. [MR1982426](#)
- [8] CHAKERIAN, J. and HOLMES, S. (2010). Computational tools for evaluating phylogenetic and hierarchical clustering trees. Available at [arXiv:1006.1015](https://arxiv.org/abs/1006.1015).
- [9] DEGNAN, J. H. and SALTER, L. A. (2005). Gene tree distributions under the coalescent process. *Evolution* **59** 24–37.
- [10] DONNELLY, P. and TAVARÉ, S. (1995). Coalescents and genealogical structure under neutrality. *Annu. Rev. Genet.* **29** 401–421.
- [11] DOOLITTLE, W. F. (1999). Lateral genomics. *Trends Genet.* **15** M5–M8.
- [12] FELSENSTEIN, J. (1985). Confidence limits on phylogenies: An approach using the bootstrap. *Evolution* **39** 783–791.
- [13] FELSENSTEIN, J. (2004). *Inferring Phylogenies*. Sinauer, Sunderland, MA.
- [14] FLETCHER, P. T., LU, C., PIZER, S. M. and JOSHI, S. (2004). Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans. Medical Imaging* **23** 995–1005.
- [15] GOODMAN, M., PORTER, C. A., CZELUSNIAK, J., PAGE, S. L., SCHNEIDER, H., SHOSHANI, J., GUNNELL, G. and GROVES, C. P. (1998). Toward a phylogenetic classification of primates based on DNA evidence complemented by fossil evidence. *Mol. Phyl. Evol.* **9** 585–598.

- [16] GROMOV, M. (1987). Hyperbolic groups. In *Essays in Group Theory. Mathematical Sciences Research Institute Publications* **8** 75–263. Springer, New York. [MR0919829](#)
- [17] GUINDON, S. and GASCUEL, O. (2003). A simple, fast, and accurate algorithm to estimate large phylogenies by maximum likelihood. *Syst. Biol.* **52** 696–704.
- [18] HASTIE, T. and STUETZLE, W. (1989). Principal curves. *J. Amer. Statist. Assoc.* **84** 502–516. [MR1010339](#)
- [19] HILLIS, D. M., HEATH, T. A. and ST. JOHN, K. (2005). Analysis and visualization of tree space. *Syst. Biol.* **54** 471–482.
- [20] HOLMES, S. (2005). Statistical approach to tests involving phylogenies. In *Mathematics of Evolution and Phylogeny* (O. Gascuel, ed.) 91–120. Oxford Univ. Press, Oxford.
- [21] HUELSENBECK, J. and HILLIS, D. (1993). Success of phylogenetic methods in the four-taxon case. *Syst. Biol.* **42** 247–264.
- [22] KUPCZOK, A., VON HAESELER, A. and KLAERE, S. (2008). An exact algorithm for the geodesic distance between phylogenetic trees. *J. Comput. Biol.* **15** 577–591. [MR2425443](#)
- [23] NYE, T. M. W. (2008). Trees of trees: An approach to comparing multiple alternative phylogenies. *Syst. Biol.* **57** 785–794.
- [24] NYE, T. M. W. (2011). Supplement to “Principal components analysis in the space of phylogenetic trees.” [DOI:10.1214/11-AOS915SUPP](#).
- [25] OWEN, M. and PROVAN, J. S. (2010). A fast algorithm for computing geodesic distances in tree space. *IEEE/ACM Trans. Comp. Biol. and Bioinf.* **8** 2–13.
- [26] RAMBAUT, A. and GRASSLY, N. C. (1997). Seq-Gen: An application for the Monte Carlo simulation of DNA sequence evolution along phylogenetic trees. *Comput. Appl. Biosci.* **13** 235–238.
- [27] ROBINSON, D. F. and FOULDS, L. R. (1981). Comparison of phylogenetic trees. *Math. Biosci.* **53** 131–147. [MR0613619](#)
- [28] STOCKHAM, C., WANG, L.-S. and WARNOW, T. (2002). Statistically based postprocessing of phylogenetic analysis by clustering. *Bioinformatics* **18** S285–S293.
- [29] WANG, H. and MARRON, J. S. (2007). Object oriented data analysis: Sets of trees. *Ann. Statist.* **35** 1849–1873. [MR2363955](#)

SCHOOL OF MATHEMATICS AND STATISTICS
NEWCASTLE UNIVERSITY
NEWCASTLE UPON TYNE
TYNE AND WEAR
NE1 7RU
UNITED KINGDOM
E-MAIL: tom.nye@ncl.ac.uk
URL: <http://www.mas.ncl.ac.uk/~ntmwn>